Orders, Complexes, Cryptography and Quantum Yang–Mills

Arithmetic modules, tensor grammars, quantum learning, fluid dynamics and a verifiable spectral research program

Prepared for Amelie Schreiber

Research synthesis and computational companion

14 November 2026

Unified research monograph, version 1.5: integrated constructions and completed computational steps

Abstract and reading contract

Unified edition, 14 November 2026. The original monograph and both later research bodies now form this one source-built book, with continuous chapter and theorem numbering. The integrated chapters add certified surface computation, measured graph-learning controls, complete finite elliptic circuits, covering theory for orders and networks, trained quantum coarse graining, observed WMAP searches and executed uniform-arithmetic, operator-learning and calibration steps. Earlier dated verification records retain their original scope; the new integration has its own source-bound tests. The existing project homepage embeds the new experiments and links to this same current document.

This monograph develops a common mathematical language for integral orders, elliptic curves, complexes, quantum algorithms and gauge theory. Its starting point is the two-volume *Lattices over Orders*: Volume I by Klaus W. Roggenkamp and Verena Huber-Dyson, and Volume II by Roggenkamp. The accompanying GRH research program supplies a second starting point: surface orders, explicit resolutions, determinant constructions and a carefully separated arithmetic spectral comparison problem. The objective is to explain the strongest defensible connections between these structures and to turn a selected set of elementary interfaces into reproducible mathematics and Lean 4 proofs.

The central organizational principle is that algebraic structure, metric information and limiting analysis are different kinds of data. An order describes integral multiplication. A complex describes compatible maps and their cohomology. An elliptic symbol supplies analytic regularity. A positive metric supplies inequalities. A limiting construction supplies a continuum object. None of these can silently replace another. Their interaction is nevertheless substantial: elliptic-curve endomorphism rings are orders; module-lattice cryptography equips arithmetic modules with geometric and probabilistic data; varieties of complexes describe exactness and deformation; elliptic complexes turn symbol classes into Fredholm indices; gauge theory uses such complexes to control classical moduli while quantum reconstruction requires additional positivity and convergence.

The cryptographic chapters derive the finite abelian hidden-subgroup mechanism underlying Shor's algorithm, explain what changes in recent multidimensional and space-efficient variants, and isolate the assumptions needed to transfer arithmetic structure into computational savings. They also derive the elementary correctness mechanism of noisy linear encryption, distinguish cryptographic hardness from a lattice's mere algebraic presentation, and explain why error distributions obstruct a direct noiseless-period argument. Proposed improvements are specified as research tasks with correctness and cost obligations, not presented as measured advances.

The Yang–Mills chapters incorporate the October–November 2026 OpenAI submission and the independent working group's dated assessment. The assessment supports a quantum field theory construction from a four-dimensional Wilson lattice scaling limit and explicitly separates this from uniqueness of the infinite-volume limit and from a positive mass gap. Historical build records and imported theorem statements are identified as external evidence; this companion does not claim to have independently recompiled that million-line formalization. Its own smaller Lean package is separately built and audited. The first four modules contain 82 declarations; reduction and transport

contracts, including a nonnormal matrix witness, together with the weighted-gain propagation module, extend the selected package to 121 declarations across seven modules. Each has its own exact type, source hash and axiom report.

The fluid extension reviews the complete English transcript of BrainTruffle's microscopic fluid-simulation introduction, its named primary sources with explicit coverage, and OpenAI's September 2026 Navier–Stokes release. The released theorem constructs smooth-forcing breakdown from initial rest in three dimensions, addressing the forced whole-space and periodic alternatives of the Clay problem; it is not an unforced Navier–Stokes theorem. We develop its algebraic and analytical interfaces through the quadratic order $\mathbb{Z}[\sqrt{2}]$, exact toral averaging, differential complexes, constrained moment correction, covariance inversion and forbidden operator compositions. Separate proofs explain information loss, memory, quantum heat channels, Carleman truncation and a structurally dissipative learned closure. The original continuum proof remains an attributed external theorem, and a fresh upstream formal build did not complete because dependency artifacts could not all be retrieved.

The second mathematical audit extends this framework to typed tensor diagrams, graph grammars, quantum graph learning and conformal cyclic cosmology. Birdtrack contractions carry explicit wire spaces, duals and normalization factors. The finite rewrite theorems distinguish generation of a new model from preservation of an existing interpretation. A trainable quantum ToricGT prototype supplies an executed graph-learning experiment with a matched classical control, a frozen-feature ablation, three declared seeds and independent exact path certificates. Its circuit is simulated on the authorized NVIDIA GPU; these measurements do not establish quantum computational advantage.

The cosmological chapters review Penrose's cross-aeon proposals, the 2025 spinor construction, both positive and negative CMB studies, and the physical and statistical conditions required to infer a message. An explicit elliptic-curve quadrature, conformal-field calculations, information inequalities and calibrated synthetic-sky experiment provide mathematical connections without claiming observational detection of a prior-aeon message. The language chapter takes a direct counterstance to categorical prohibitions on neural understanding: it constructs exact neural recurrences and finite counterfactual mechanisms, while distinguishing these capabilities from learned generalization, causal identifiability and claims about consciousness.

What this document proves, cites and proposes

Elementary results stated with a proof are proved in the text under their displayed assumptions. Deep foundational results such as the classification of supersingular endomorphism algebras, the Atiyah–Singer index theorem, the Buchsbaum–Eisenbud exactness criterion and the contemporary Yang–Mills construction are used as explicitly attributed external results. A proof of every theorem in the cited literature is not contained here. The new Lean package verifies the exact declarations enumerated in its verification report; a citation to a theorem or an informal proof in this monograph does not by itself mean that theorem is formalized.

Neither GRH nor the Yang–Mills mass gap is claimed as a theorem of this work. The manuscript contains rigorous conditional implications and explicit counterexamples to overstrong inference rules. This is a constructive research boundary: a missing implication becomes a named mathematical problem, with the required hypotheses and possible falsification tests written down.

Source review and its limits

The attached books contain 309 and 392 scanned PDF pages. The OCR text of all 701 pages was read across the review team, with page coverage recorded and selected load-bearing statements checked against page images. Optical character recognition is a navigation aid, especially unreliable for subscripts, ideals and exact-sequence arrows. An all-page extraction record is not an independent line-by-line proof audit of 701 pages.

The inherited repository is pinned to commit `dbd954eede479a5129831e4e7c30537389ae973d`. Its tracked content includes thousands of research, computational and historical output files and hundreds of distinct outbound links. The companion preserves a machine-readable inventory and distinguishes inherited material, directly read passages, inspected theorem statements, reproduced calculations and newly compiled formal declarations. It does not label every outbound linked work as fully proof-audited. The source concordance and review records state the actual coverage so that future work can extend it without guessing what has been checked.

How to use the companion

The project page offers exact finite examples. Its quantum companion executes an 18-qubit toric encoder and a four-parameter variational circuit, while separately reporting larger NVIDIA GPU state-vector and Pauli-frame experiments. Higher-genus cases retain exact cellulations, dessin permutations and binary boundary matrices. A four-coordinate quadric represents two consecutive maps whose composition is zero. A torus view illustrates a complex lattice and changes of basis. A finite-field elliptic curve is displayed separately from its real counterpart. A small Fourier-sampling example shows the annihilator of a relation subgroup. Three added algebraic chapters develop wedge shuffle maps, tableaux, Schur factors and standard-monomial filtrations; the subsequent gentle-algebra chapters give explicit module decompositions, corrected matching-semigroup relations, determinant characters and a cohomological route to finite zeta functions. A spectral family holds the Euler index fixed while its positive eigenvalue tends to zero. Every visualization states its coordinates, parameter range and mathematical interpretation, and the code includes independent numerical or exact checks.

Three further interactive pages execute the full selected quantum ToricGT model, typed circuit rewrites and neural language constructions, and a fully labelled synthetic cosmological search. Parameter changes recompute their stated finite models. The source records distinguish GPU training, browser inference, exact certificates and formal declarations.

The repository retains the original GRH material and places the new work under `extension/`. Third-party books and downloaded research PDFs used during review are not republished as part of this extension. The original manuscript, bibliography, source ledger, executable examples, Lean sources and verification records can be inspected and rebuilt together.

Contents

Abstract and reading contract **i**

I Arithmetic structure and the inherited research program **1**

1 A common language with explicit interfaces **2**
1.1 The mathematical objects that must be kept distinct . . . 2
1.2 A small order that already contains the key issues . . . 3
1.3 What survives from the GRH program . . . 4
1.4 Determinants, cohomology and spectra . . . 5
1.5 A map of rigorous and prospective connections . . . 6

2 Orders as integral arithmetic structures **7**
2.1 The role of the two volumes . . . 7
2.2 Full lattices and orders . . . 8
2.3 Endomorphism orders and faithful models . . . 9
2.4 Localization, support, and what local data detects . . . 9
2.5 An explicit Morita equivalence . . . 10
2.6 Trace pairings, dual lattices, and discriminants . . . 11
2.7 From arithmetic lattices to Euclidean geometry . . . 12
2.8 Quaternion orders as four-dimensional arithmetic . . . 13
2.9 Maximality, separability, and their limits . . . 14
2.10 A finite certificate for maximal overorders . . . 14
2.11 Averaging and integral extension obstructions . . . 15

3 Elliptic curves through their endomorphism lattices **17**
3.1 Two arithmetic questions attached to the same curve . . . 17
3.2 Complex tori and scalar endomorphisms . . . 17
3.3 Quadratic orders and the conductor . . . 18
3.4 Ideal classes as a torsor of elliptic curves . . . 19
3.5 Finite-field endomorphisms and Frobenius . . . 20
3.6 Tate modules as arithmetic lattices . . . 20
3.7 Quaternion ideals and the Deuring correspondence . . . 21
3.8 The quaternion norm and the meaning of degree . . . 22
3.9 What these bridges establish for cryptography . . . 22
3.10 A worked comparison of three rank-two objects . . . 23
3.11 An exact arithmetic certificate for a quadratic order . . . 23
3.12 A research program that is actually implied . . . 24

II Complexes, topology and spectral information **25**

4 From integral modules to complexes and stable invariants **26**
4.1 What the second volume contributes . . . 26
4.2 The exact category of lattices . . . 26
4.3 Three different Grothendieck groups . . . 27
4.4 Euler classes of complexes . . . 29
4.5 Schanuel's lemma as a comparison mechanism . . . 30
4.6 What hereditary and Bass hypotheses change . . . 30
4.7 A Gram-matrix bridge already present in the book . . . 31
4.8 From projective modules to topological bundles . . . 32

5 Varieties of complexes, exactness, and deformation **34**
5.1 A parameter space with an equation . . . 34
5.2 Classification under changes of bases . . . 35
5.3 Dimension of a rank stratum . . . 35
5.4 First-order deformations and tangent equations . . . 36
5.5 The endomorphism complex and the meaning of gauge . . . 37
5.6 Why ranks alone fail over commutative rings . . . 38
5.7 A complete exactness proof: the two-generator Koszul complex . . . 39
5.8 Base change and the appearance of new homology . . . 40
5.9 The real four-dimensional quadric as an exact visual model . . . 40

6 Elliptic complexes, Hodge theory, and gauge deformations **42**
6.1 Ellipticity is a condition on symbols . . . 42
6.2 Exterior algebra gives a complete symbol calculation . . . 43
6.3 Finite-dimensional Hodge theory with a full proof . . . 43
6.4 The analytic Hodge theorem and its input . . . 44
6.5 The index as an additive invariant . . . 45
6.6 A complete counterexample to a uniform-gap inference . . . 46
6.7 Connections and curvature: when a differential squares to zero . . . 46
6.8 The anti-self-dual deformation complex . . . 47
6.9 Expected dimension and actual local dimension . . . 48
6.10 Classical action bounds do not establish a quantum mass gap . . . 48
6.11 A lower-dimensional bridge through holomorphic bundles . . . 49
6.12 Interfaces suitable for formal verification . . . 49

7 Exterior algebra, tableaux, and Schur constructions **51**
7.1 The algebra needed to describe a space of complexes . . . 51
7.2 Exterior powers and determinant coordinates . . . 51
7.3 Exterior coproducts and shuffle signs . . . 52
7.4 Partitions and a characteristic-free Schur presentation . . . 53
7.5 Rational highest weights and duals . . . 55
7.6 Schur modules applied to complexes . . . 55

8 Standard monomials on spaces of complexes **57**
8.1 What a standard monomial theorem must establish . . . 57
8.2 Rank-one matrices: the Segre example . . . 57
8.3 Two-row tableaux and the Plucker algebra . . . 58
8.4 A complete complex example: dimensions $(1, 2, 1)$ . . . 59
8.5 The general complex and its standard products . . . 61
8.6 Rank strata and orbit closures without a source shortcut . . . 62
8.7 The Schur decomposition with all vertices and duals . . . 63

9 Associated graded rings, semigroups, and toric limits **64**
9.1 Three different constructions . . . 64
9.2 A monomial degeneration for a quadratic straightening algebra . . . 65
9.3 Initial subalgebras and the diagonal semigroup . . . 66
9.4 Gelfand–Tsetlin patterns in the example . . . 67
9.5 Two flat families from the same Grassmannian . . . 68
9.6 How the algebra enters the larger research program . . . 69

10 Gentle algebras, matching semigroups and determinant coordinates **71**
10.1 From a relation to a variety of complexes . . . 71
10.2 A second example: Kronecker pencils and regular modules . . . 73
10.3 Homological qualifications: a gentle algebra of infinite global dimension . . . 73
10.4 Schofield determinants, with the transformation law derived . . . 74
10.5 Schur factors and the semi-invariant matching condition . . . 75
10.6 A matching semigroup with two explicit binomial relations . . . 77
10.7 Quotients and the information retained by semi-invariants . . . 78

11 Exterior traces, zeta functions and the spectral comparison problem **80**
11.1 Exterior powers supply determinant coefficients . . . 80
11.2 Traces, closed walks and a finite zeta model . . . 81
11.3 The elliptic-curve numerator as a cohomological determinant . . . 82
11.4 What standard monomials can contribute to the arithmetic program . . . 82
11.5 A concrete program for further formalization . . . 83

III Cryptography and quantum algorithms **84**

12 Elliptic-curve arithmetic and quantum period finding **85**
12.1 The group problem and its arithmetic representation . . . 85
12.2 The relation subgroup . . . 86
12.3 Fourier sampling with every normalization explicit . . . 86
12.4 What an endomorphism order contributes . . . 87
12.5 A faithful comparison of recent algorithmic developments . . . 88
12.6 What could count as a new improvement . . . 90

13 Order lattices, noise and post-quantum cryptography **91**
13.1 From a module to a geometry . . . 91
13.2 An algebraic equivalence need not preserve short vectors . . . 92
13.3 Noisy linear equations . . . 92
13.4 A complete elementary correctness calculation . . . 93

13.5 Why the noiseless hidden-subgroup argument does not transfer directly 94
13.6 Orders, local structure and a responsible design criterion 94
13.7 Other cryptographic families and the limits of the common framework 95
13.8 A quantitative requirement for any proposed generalization 96

14 Implemented cryptography and verifiable artifact contracts 97
14.1 From mathematical examples to executable interfaces 97
14.2 Authenticated serialization and expected context 98
14.3 Signatures, replay and scientific meaning 98
14.4 The post-quantum boundary 99
14.5 Recorded execution and formal coverage 99

15 Resource-certified elliptic arithmetic 100
15.1 A complete specification before a resource claim 100
15.2 An emitted gate construction with clean work 100
15.3 Field and variable-point oracles 101
15.4 A fully paid endomorphism comparison 101
15.5 Formal and numerical verification layers 102

16 Executable quantum codes and neural circuits on dessins 103
16.1 From a drawing to a quantum computation 103
16.2 Oriented cells and clean-dessin permutations 103
16.3 An explicit elliptic Belyi model 104
16.4 The cellular chain complex is the stabilizer specification 105
16.5 Logical operators, distance and syndrome 106
16.6 Higher genus by a verified connected sum 107
16.7 A deterministic encoder with a complete proof 108
16.8 A finite Hamiltonian with a provable gap 109
16.9 A surface-local variational neural circuit 109
16.10 Training, parameter shifts and the meaning of a held-out input 110
16.11 Execution evidence and scaling 111
16.12 Extensions with explicit mathematical obligations 115

IV Typed tensor diagrams and quantum graph learning 116

17 Tensor diagrams, invariant theory and quantum graph grammars 117
17.1 A diagram must have a type and an interpretation 117
17.2 Wedge powers, Young symmetry and determinants as diagrams 118
17.3 A color contraction with every normalization fixed 120
17.4 Channels as positive tensors 121
17.5 A precise grammar for the executable circuits 121
17.6 How the grammar preserves a complex 123
17.7 Expressivity: four different classifications 124
17.8 When the tensor grammar is also efficiently classical 126
17.9 A learned rule selector with a checked semantic core 127
17.10 Executable obligations 127

18 A trainable quantum interface for ToricGT **128**
18.1 Which architecture is being extended? . . . 128
18.2 Tokenization graphs and exact solution certificates . . . 129
18.3 A finite Weyl encoder with trainable gates . . . 131
18.4 A graph transformer whose symmetries can be stated precisely . . . 134
18.5 Toric path relations, graph rewrites, and steering regions . . . 135
18.6 Three source corrections with complete small counterarguments . . . 137
18.7 Completed learning experiment and numerical evidence . . . 139
18.8 Connections that follow, and research questions that remain . . . 141

V Quantum fields and a precise research synthesis **143**

19 Gauge fields, curvature, and the meaning of a complex **144**
19.1 The precise connection to the preceding algebra . . . 144
19.2 Gauge covariance and the variational equations . . . 145
19.3 Four dimensions and self-duality . . . 146
19.4 The anti-self-dual deformation complex . . . 147
19.5 A finite-dimensional algebraic model: monads and ADHM . . . 148
19.6 Why an index cannot establish a quantitative gap . . . 149

20 From finite gauge lattices to a quantum spectral gap **150**
20.1 Three different uses of the word lattice . . . 150
20.2 A finite Wilson model exists by elementary integration . . . 150
20.3 Reflection positivity is additional structure . . . 151
20.4 The Hilbert-space quotient behind reconstruction . . . 152
20.5 What the mass-gap statement actually requires . . . 153
20.6 A rigorous countermodel to ultraviolet-to-infrared inference . . . 154
20.7 Uniform estimates that would survive a limit . . . 154
20.8 Renormalization and what a running coupling can prove . . . 155
20.9 Pure phases, clustering, and the order of quantifiers . . . 156

21 The 2026 OpenAI Yang–Mills construction: statement and verification **158**
21.1 A dated research result, not the older amplitude result . . . 158
21.2 Immutable source identities and review coverage . . . 158
21.3 The headline theorem and its quantifiers . . . 159
21.4 A hierarchy with finite-degree local fields . . . 159
21.5 The two limits and the physical normalization . . . 160
21.6 Reconstruction, phases, and test spaces . . . 160
21.7 The trace law, non-freeness, and OPE estimates . . . 161
21.8 A stronger acceptance target and the scope of checking . . . 162
21.9 What can be built on the result now . . . 162

22 An explicit arithmetic–geometric bridge through elliptic tori **164**
22.1 Orders acting on a complex elliptic curve . . . 164
22.2 The same lattice controls flat unitary line bundles . . . 165
22.3 An elliptic operator with a completely explicit spectrum . . . 165
22.4 How this bridge relates to cryptography . . . 166

22.5 A precise synthesis and a testable research program . . . 166

23 A character family inside a variety of complexes 168
23.1 A universal complex over the character torus . . . 168
23.2 Characters, local coefficients and sign conventions . . . 169
23.3 A complete Hodge calculation . . . 169
23.4 An order action with an explicit cochain map . . . 170

VI Information reduction, fluid dynamics and the released singularity proof 172

24 Microscopic information, kinetic limits, and fluid closure 173
24.1 What the microscopic perspective establishes . . . 173
24.2 Measurement, phase, and the limits of tomography . . . 174
24.3 Exterior powers are a physical state constraint . . . 176
24.4 Born–Huang expansion before Born–Oppenheimer truncation . . . 176
24.5 When a quantum packet behaves approximately as a particle . . . 178
24.6 Molecular dynamics and conserved microscopic information . . . 179
24.7 Marginals, correlations, and the kinetic approximation . . . 180
24.8 The collision algebra and its five invariants . . . 181
24.9 The weak collision identity and entropy production . . . 182
24.10 Exact moment balances and the closure problem . . . 184
24.11 DSMC as a stochastic approximation, with auditable assumptions . . . 186
24.12 Reduced models and learned closures . . . 187
24.13 What passes from complexes to a continuum theorem . . . 190

25 The Navier–Stokes release: exact scope and proof interfaces 191
25.1 What the released theorem says . . . 191
25.2 An exact algebra for residual correction . . . 192
25.3 Concentration, energy, and the information lost by norms . . . 193
25.4 Similarity coordinates as an exact change of variables . . . 194
25.5 A quadratic relation that controls analytic derivative loss . . . 195
25.6 From formal orders to actual smooth fields . . . 196
25.7 The finite correction cycle as a verifiable dependency graph . . . 198
25.8 Localization and smooth force extension . . . 199
25.9 Comparison, quantitative residual checks, and proof strength . . . 200
25.10 Viscosity, periodization, and preservation of quantifiers . . . 201
25.11 Formal proof provenance and what this extension adds . . . 201

26 Orders, complexes and the algebra of a fluid singularity 203
26.1 What a useful reformulation must preserve . . . 203
26.2 The de Rham complex behind incompressible motion . . . 203
26.3 A quadratic order occurs in the actual construction . . . 204
26.4 The number-field norm supplies an analytic estimate . . . 206
26.5 Symmetric tensors, covariance and the stress correction . . . 206
26.6 Moment determinants and exact compatibility . . . 207
26.7 From varieties of complexes to compatible fluid discretization . . . 208
26.8 Gentle algebras, path relations and quantitative interfaces . . . 210

26.9 Determinants, spectra and what the reformulation adds . . . 210

27 Quantum information, exact reduction and learned fluid models 211
27.1 Information reduction is a map, not an autonomous dynamics . . . 211
27.2 Coarse-graining and the unresolved stress . . . 212
27.3 A quantum observable can have the same closure obstruction . . . 213
27.4 An explicit heat channel on a periodic surface grid . . . 213
27.5 What the Madelung transformation does and does not derive . . . 214
27.6 Carleman lifting and the symmetric algebra . . . 215
27.7 What the quantum-algorithm literature actually supplies . . . 216
27.8 Exact finite repairs for collision block encodings . . . 219
27.9 A global compatibility test for a spin representation . . . 221
27.10 A mathematically constrained quantum ToricGT extension . . . 222
27.11 The executable examples and their proof boundaries . . . 223

28 Exact reduction contracts and the information lost by tropicalization 225
28.1 The exact closure criterion . . . 225
28.2 A lower bound on every approximate closure . . . 226
28.3 Signed cancellation is invisible to orders alone . . . 227
28.4 A current analytic obstruction to removing the force . . . 227
28.5 Formal coverage and reproducibility . . . 228

29 Corrected transport, information and quantum contracts for tropical quivers 229
29.1 Source scope and a quantitative contract . . . 229
29.2 Suffix transport is the correct cyclic coboundary . . . 230
29.3 A positive Laplacian does not itself supply a complex . . . 231
29.4 What a prefix-derived certificate can add . . . 231
29.5 Damping and stochastic limits require their own hypotheses . . . 232
29.6 Quantum and gentle-algebra versions of the transport contract . . . 233
29.7 Tropical approximation is not an automatic universality transfer . . . 234
29.8 Independent repairs to the tropical-attention reference . . . 235
29.9 A nonnormal obstruction to a Jacobian power certificate . . . 236
29.10 Path multiplicity separates tropical growth from operator stability . . . 237
29.11 Weighted small-gain certification at a fixed rate . . . 238
29.12 Proper scoring must include its probability boundary . . . 241
29.13 Predictive synergy is not determined by geometric overlaps . . . 241
29.14 Implications for the research program . . . 242

30 Quantum–tropical fluid realizations: operators, control and substrates 243
30.1 What a realization must specify . . . 243
30.2 A finite operator algebra with conservative quantum transport . . . 244
30.3 Toric features, tropical routing and the information they discard . . . 245
30.4 Nonlinear dynamics on an operator state space . . . 246
30.5 Control, observability and robustness . . . 247
30.6 An exact nonlinear quantum-state feedback example . . . 249
30.7 Candidate engineered and natural substrates . . . 250
30.8 De novo substrate design with generative models . . . 252
30.9 Connections to the wider research program . . . 253

31 Coverings, orders and learned quantum renormalization 254
31.1 Voltage covers preserve the specified local presentation . . . 254
31.2 Integral relations can obstruct a special-fiber cover . . . 255
31.3 Neural pullback, descent and the cost of arbitrary identifiers . . . 255
31.4 A learned hierarchy of physical quantum channels . . . 256
31.5 What the learned renormalization map actually learns . . . 256
31.6 An explicit infinite-width theorem . . . 257
31.7 Cover symmetry and the dependence structure of wide quantum networks . . . 258
31.8 Testing whether gate learning improves the representation . . . 259

VII Conformal geometry, information and language 260

32 Conformal aeons, elliptic integrals and identifiable cosmic information 261
32.1 Four questions that require different evidence . . . 261
32.2 Conformal geometry with the singular boundary kept explicit . . . 262
32.3 A genuine elliptic-curve bridge in the Friedmann equation . . . 263
32.4 A spinor contraction that gives a conditional conserved charge . . . 265
32.5 From a possible carrier to an identifiable message . . . 266
32.6 Quantum constraints on a proposed communication channel . . . 268
32.7 Ring statistics, matched filters and the cost of searching . . . 269
32.8 What a neural search and a new cosmic observable can establish . . . 271
32.9 A reproducible synthetic search, with its actual result . . . 272
32.10 A testable program for learned tensor and graph models . . . 273

33 CCC geometry and a real-map search with frozen learned statistics 275
33.1 The original mathematical interfaces are retained . . . 275
33.2 Actual datasets and resolution-aware tests . . . 275
33.3 Learned statistics selected before observation scoring . . . 276
33.4 The null is part of the result . . . 276
33.5 Observed results and frequency sensitivity . . . 277
33.6 Interactive replay without moving the goalposts . . . 278
33.7 A joint phase action that retains cross-product dependence . . . 278

34 Neural computation, formal languages and a direct reply to Chomsky 279
34.1 The claim being challenged . . . 279
34.2 A constructive recursive recognizer . . . 280
34.3 Transformer theorems require their hypotheses . . . 281
34.4 A neural circuit can compute a counterfactual . . . 281
34.5 Empirical linguistic evidence and its strongest objection . . . 284
34.6 Learning a rule selector without surrendering guarantees . . . 284
34.7 Finite-family generalization with an explicit selection cost . . . 285
34.8 What follows for the proposed cosmological grammar . . . 286

VIII Certified surface computation and its measured realizations 288

35 A calculus of certified realizations 289
35.1 The organizing principle . . . 289

35.2 From a relation to topology, energy and quantum checks . . . 289
35.3 An explicit order whose special fiber is gentle . . . 290
35.4 A finite spectral gap, proved and executed . . . 291
35.5 The continuum transition as a precise theorem . . . 292
35.6 Probability diffusion as an open quantum evolution . . . 292
35.7 Controller-to-circuit robustness . . . 293
35.8 The arithmetic spectral target . . . 293
35.9 One program, distinct routes to evidence . . . 293

36 Nonlinear fluid coordinates and unitary half-density transport 296
36.1 A complete flow from the finite energy identity . . . 296
36.2 A nonlinear triad with an exact Jacobian law . . . 297
36.3 A unitary discretization and its actual quantum operation . . . 298
36.4 Operator learning, cyclic structure and the continuum problem . . . 298

37 Source-led graph universality: preserved proofs and constructive repairs 300
37.1 What was missing from the preceding synthesis . . . 300
37.2 A direct uniform graph-to-graph construction . . . 300
37.3 Finite interpolation, circuits and unbounded computation . . . 302
37.4 Dependency repairs and explicit counterexamples . . . 303
37.5 How these results connect to the physical program . . . 304
37.6 Additional primary references . . . 304
37.7 A complete finite graph-function theorem in Lean . . . 304

38 Graph languages and computational mechanisms 306
38.1 From function approximation to graph computation: the precise claim . . . 306
38.2 Graph grammars, uniform computation and the language hierarchy . . . 310
38.3 Circuit interpretation as a testable form of semantic competence . . . 312
38.4 ConvexTok, graph incidence and the preservation of information . . . 314
38.5 GFlowNets over derivations, programs and complete forests . . . 316
38.6 Dynamical expressivity, protected memory and error accumulation . . . 318
38.7 Quantum and hybrid generalizations . . . 319
38.8 Quiver mutation, cluster structures and the semantics of rewriting . . . 319
38.9 Complexity: new interfaces do not make resource costs disappear . . . 322
38.10 An extended response to anti-neural and restrictive function-learning claims . . . 322
38.11 Cosmological graph grammars and falsifiable inference . . . 324
38.12 Evaluation and implementation program . . . 324
38.13 Analogical transfer is a proposal; deduction is a checked consequence . . . 325

39 Physical data, coding and the measured study 327
39.1 What a discovery test must discriminate . . . 327
39.2 Coding definitions with an honest denominator . . . 327
39.3 Executed bounded experiments . . . 328
39.4 Additional realistic tests, not executed results . . . 329
39.5 Recorded results from the executed pilot . . . 330

40 Measured signals, communicative controls and complete code lengths 331
40.1 Information available versus computation made accessible . . . 331
40.2 Sequential evidence and the cost of searching a grammar . . . 331
40.3 Two additional measured controls . . . 332
40.4 An emitted lossless code, not a likelihood-only surrogate . . . 333
40.5 A physical control that rejects an attractive inference . . . 333
40.6 Reproduction and source rights . . . 333

41 BGG, category $\mathcal{O}$ and structural synthesis 335
41.1 Three BGG constructions, with their domains kept separate . . . 335
41.2 An exact highest-weight example and its truncation failure . . . 336
41.3 Blocks, Ext data and quivers with relations . . . 337
41.4 Exterior BGG certificates, exactness and integral obstructions . . . 338
41.5 How to use the material in the neural research program . . . 339
41.6 Tropical, toric, hypertoric and surface connections . . . 339
41.7 Quantum information, orders and the GRH/Hilbert–Pólya boundary . . . 340
41.8 An executable agenda and the limits of the current evidence . . . 341

42 Reading the linked surface experiment 342

IX Completed computational steps and the unified research program 344

43 Completed next steps: uniform arithmetic and cover-compatible operators 345
43.1 One continuing manuscript and one project entry point . . . 345
43.2 Uniform reversible field primitives instead of lookup synthesis . . . 345
43.3 A trained neural operator that respects covers and quantum channels . . . 347
43.4 Higher-resolution CCC calibration with explicit simulation uncertainty . . . 348
43.5 Completion ledger and remaining research obligations . . . 348

44 A worked bridge: orders, periodic complexes and spectra 350
44.1 The order and its two complementary elements . . . 350
44.2 Tensor products detect information that ranks omit . . . 351
44.3 Generic splitting and singular specialization . . . 352
44.4 The finite Hodge operator of the same complex . . . 352
44.5 How the cryptographic comparison enters . . . 353

45 A quantitative synthesis and a research program 354
45.1 The information that each theory actually supplies . . . 354
45.2 A general comparison estimate . . . 355
45.3 A diagram with explicit logical boundaries . . . 355
45.4 Three concrete research directions . . . 356
45.5 What this monograph establishes . . . 357

A Reproducible finite experiments and interactive geometry 358
A.1 What the computations verify . . . 358
A.2 The 19-element elliptic-curve example . . . 358
A.3 The Fourier annihilator experiment . . . 359
A.4 A genuine four-coordinate variety . . . 359

A.5 Index, kernel and gap are separate readouts . . . 363
A.6 The torus and plaquette displays . . . 363
A.7 Reproduction and interpretation . . . 364
A.8 Exact algebraic companion checks . . . 365

B The Lean 4 companion: verified statements and explicit limits 366
B.1 Current seven-module verification . . . 366
B.2 Original module verification result . . . 366
B.3 Why this is selected formalization . . . 367
B.4 Linear complexes and rank . . . 367
B.5 The cone and its tangent calculation . . . 368
B.6 Hodge positivity and the characteristic polynomial . . . 368
B.7 Integral shears, cyclic kernels and noisy arithmetic . . . 369
B.8 Reading an axiom audit . . . 369
B.9 Rebuilding and extending the companion . . . 369
B.10 Seven additional quantum and binary-chain declarations . . . 370
B.11 Sixteen additional diagram and transport declarations . . . 370
B.12 Complete verified companion sources . . . 371

C Fluid verification, reproducibility and proof dependencies 383
C.1 Current verification boundary . . . 383
C.2 What was checked and what each check means . . . 383
C.3 The separate upstream Navier–Stokes formalization . . . 384
C.4 Independent finite computational checks . . . 385
C.5 Reproduction commands . . . 385
C.6 The complete selected fluid formalization . . . 386

D Source concordance, review coverage, and mathematical provenance 394
D.1 What the review establishes . . . 394
D.2 The fixed source snapshot . . . 394
D.3 The classical books: extraction and passage coverage . . . 395
D.4 Three decisive anchors in Volume II . . . 395
D.5 What survives scrutiny in the upstream research program . . . 396
D.6 From sources to the new theorems in this monograph . . . 397
D.7 A finite-field correction in the source matrix reductions . . . 397
D.8 Formal proofs, computations, and remaining coverage . . . 398

E Primary sources and source-pinned implementation register 399

F Arithmetic and quantum-renormalization source register 402

G Source-bound verification of the earlier computational continuation 403

Part I

Arithmetic structure and the inherited research program

Chapter 1

A common language with explicit interfaces

1.1 The mathematical objects that must be kept distinct

The word “lattice” appears in three different settings in this project. An arithmetic lattice is a finitely generated torsion-free module over an integral base, often with an action of an order. A Euclidean lattice is a discrete subgroup of a real inner-product space, of full rank in its real span. A lattice gauge theory is a field theory on a discrete cell structure. The first can produce the second after choosing an embedding and a metric. The third can use modules and chain complexes as part of its discretization. Neither passage is an equality of definitions.

Similarly, an algebraic complex, a differential complex and an elliptic curve are unrelated uses of two familiar words until additional maps have been specified. A complex of modules satisfies $d_{i+1}d_i = 0$. A differential complex is a sequence of differential operators with the same composition property. Its ellipticity is a condition on the principal symbols away from the zero section of the cotangent bundle. An elliptic curve is a smooth projective genus-one curve with a chosen point. The term “elliptic” in elliptic-curve cryptography does not make the group operation an elliptic partial differential operator.

The meaningful connections are more precise. The endomorphism ring of an elliptic curve acts on its integral ℓ-adic Tate modules for ℓ different from the field characteristic, and can be a quadratic or quaternion order. A presentation of an order-module can be studied through a complex of projective modules. Families of such complexes define algebraic varieties with rank strata. Geometric differential complexes can have these algebraic patterns as their pointwise symbols. Index theory assigns a stable integer to the resulting Fredholm problem. Quantum theory imposes further analytic conditions on correlation functions and on a reconstructed Hilbert space. This series of interfaces gives a coherent research program without treating every object as an instance of the same theorem.

Definition 1.1 (Interface data)**.** An interface between two mathematical settings consists of a specified object map, a specified morphism map, hypotheses under which the map is defined, and the properties it preserves. If the application is computational, it also includes representations, algorithms and a cost model. If the application is spectral, it also includes inner products, operator domains and convergence topology.

This definition is methodological rather than a new category-theoretic construction. It prevents a common failure: a functor preserves an algebraic invariant, and one then informally assumes that

it preserves a norm, a sampling law or a computational complexity class. The latter conclusions require their own theorems.

1.2 A small order that already contains the key issues

Let k be a field, let $R = k[[u]]$, and consider

$$\Lambda = R[X]/(X^2 - uX).$$

Write x for the image of X and $y = u - x$. Then $xy = 0$, and Λ is the simplest completed nodal order in the inherited surface-order program. Its geometric presentation is $k[[x, y]]/(xy)$, with the central parameter $u = x + y$. The following calculation makes the order structure transparent without assuming a general classification theorem. The broader surface construction and its conventions are discussed in the original research sources [137–139].

Proposition 1.2 (Normalization and conductor of the nodal order)**.** *The map*

$$\Phi : \Lambda \longrightarrow R \times R, \qquad \Phi(a + bx) = (a + bu, a)$$

is an injective R-algebra homomorphism with image

$$\{(f, g) \in R \times R : f \equiv g \pmod{u}\}.$$

Consequently Λ is free of rank two over R, has generic algebra $k((u)) \times k((u))$, has normalization $R \times R$, and has conductor $uR \times uR$ in that normalization.

Proof. Every class in $R[X]/(X^2 - uX)$ has a unique representative $a + bX$, because the relation is monic of degree two. Thus $1, x$ form an R-basis. The proposed image of x is $(u, 0)$, whose square is $u(u, 0)$, so the map respects the defining relation. If $\Phi(a + bx) = 0$, the second coordinate gives $a = 0$, and the first gives $bu = 0$. Since R is a domain, $b = 0$. A pair is in the image exactly when $f - g$ is divisible by u, proving the image statement.

After inverting u, any pair (f, g) can be written as $\Phi(g + (f - g)x/u)$, so the generic algebra is the product of two fields. The ring $R \times R$ is generated as a Λ-algebra by the integral idempotent $(1, 0)$, so it is finite and integral over Λ. Conversely, an element of $k((u)) \times k((u))$ integral over Λ is integral over R by transitivity, since Λ is finite over R. Both coordinates therefore belong to the integrally closed ring R. This proves the normalization assertion. An element (f, g) of $R \times R$ belongs to the conductor precisely when its product with every (r, s) lies in the image of Φ. Multiplication by $(1, 0)$ and $(0, 1)$ shows that both f and g must be divisible by u. That condition is sufficient because both resulting coordinates are zero modulo u. Hence the conductor is exactly $uR \times uR$. □

The example exhibits three levels at once. Over the fraction field there are two separate branches. Integrally their constant terms are glued. At the special fibre $u = 0$, the relation becomes $x^2 = 0$, so the fibre is the algebra of dual numbers. The generic algebra, the integral order and the special fibre are different objects with different module categories. A proof using one level cannot be transferred to another merely because the same symbol x appears.

Example 1.3 (A parameter is not an eigenvalue)**.** In the basis $(1, x)$, multiplication by x is

$$M_x = \begin{pmatrix} 0 & 0 \\ 1 & u \end{pmatrix}, \qquad M_y = \begin{pmatrix} u & 0 \\ -1 & 0 \end{pmatrix}.$$

They satisfy $M_xM_y = M_yM_x = 0$ and $M_x + M_y = uI$. These identities hold in $M_2(R)$. To ask about Hermitian positivity, one must first choose a specialization of u, an embedding into a field with involution and a positive inner product. The displayed integral matrices are not automatically self-adjoint in the standard complex inner product.

This is directly relevant to the proposed relationship with elliptic complexes. The matrix relation gives a complex, but it gives no symbol exactness away from zero. It also gives no elliptic estimate and no continuum Hamiltonian. These distinctions are productive: the algebraic model provides equations to formalize, while the analytic model identifies the additional estimates that would be required.

1.3 What survives from the GRH program

The pinned GRH repository develops surface orders from clean dessins, projective resolutions indexed by forbidden paths, determinant identities, representation-theoretic components and several arithmetic estimates. Its current book explicitly records that the original unconditional GRH claims do not follow from the constructions. It identifies an unresolved arithmetic spectral comparison, and distinguishes finite computational checks from analytic and categorical proofs [137].

The extension here uses the surviving architecture: specify a category, prove exact algebraic identities, introduce a metric only when it is needed, and isolate the arithmetic comparison as a separate statement. The construction of a positive matrix can be valuable even when it does not yet compute an L-function. An exact projective resolution can be valuable even when its determinant has cycle factors rather than prime factors. The absence of the final comparison does not invalidate those intermediate mathematical results.

For orientation, a primitive nonprincipal Dirichlet character χ of conductor q has a completed L-function

$$\Lambda(s,\chi) = \left(\frac{q}{\pi}\right)^{(s+a)/2} \Gamma\left(\frac{s+a}{2}\right) L(s,\chi), \qquad \chi(-1) = (-1)^a.$$

This normalization centers its functional equation at $1/2$. For a modular elliptic-curve L-function in its usual arithmetic normalization, the center is 1. Changing to unitary normalization shifts the variable. A sentence about a "critical line" has no mathematical content until the normalization and the L-function family are fixed. Standard analytic definitions and the Dirichlet normalization are recorded in the NIST reference [107].

Proposition 1.4 (Positive pencils have real generalized eigenvalues)**.** *Let $A, B \in M_n(\mathbb{C})$ be Hermitian, with B positive definite. If*

$$\det(A - zB) = 0,$$

then $z \in \mathbb{R}$. In particular, the polynomial

$$P(s) = \det\big(A + i(s - \tfrac{1}{2})B\big)$$

has every zero on the line $\Re s = 1/2$.

Proof. Singularity gives a nonzero v such that $Av = zBv$. Multiplying by the conjugate transpose v^* gives $v^*Av = zv^*Bv$. Both quadratic forms are real because the matrices are Hermitian; the second is strictly positive. Hence $z = (v^*Av)/(v^*Bv)$ is real. For the second statement set $z = -i(s - 1/2)$. Its imaginary part is $-(\Re s - 1/2)$, which vanishes exactly on the stated line. □

The proof is short, exact and powerful. Its limitation is equally exact: it says nothing about whether $P(s)$ equals an arithmetic function, or whether a sequence of these polynomials converges to one after multiplication by a nonvanishing factor. The required arithmetic comparison is an independent theorem.

Proposition 1.5 (A precise conditional zero-location transfer)**.** *Let $U \subset \mathbb{C}$ be connected and open, let $\ell = \{s : \Re s = 1/2\}$, and let F_N be holomorphic on U. Suppose that every F_N is zero-free on $U \setminus \ell$, that F_N converges uniformly on compact subsets of U to a holomorphic function F, and that F is not identically zero. Then F is zero-free on $U \setminus \ell$.*

Proof. Assume $F(s_0) = 0$ for some $s_0 \in U \setminus \ell$. Since U is connected and F is not identically zero, its zeros are isolated. Choose a closed disk D about s_0 contained in $U \setminus \ell$, with no zero of F on its boundary. The continuous function $|F|$ has a strictly positive minimum m on that boundary. Uniform convergence on the boundary gives $|F_N - F| < m$ for all sufficiently large N. Rouché's theorem then gives the same positive number of zeros for F_N and F inside D, counted with multiplicity, contradicting the zero-free hypothesis. This argument uses the standard complex-analytic Rouché theorem; it does not establish the hypotheses for any L-function. □

For an arithmetic application one must supply F_N, prove the exact normalization, establish compact-uniform convergence on a domain large enough for the conclusion, and rule out the zero function. Pointwise agreement at a finite list of sample values is insufficient. So is convergence only on the critical line. The proof uses a two-dimensional complex neighborhood around every alleged off-line zero.

Example 1.6 (Why a finite Euler product is not the desired approximant)**.** For $p > 1$ and $\Re s > 0$, one has $|p^{-s}| = p^{-\Re s} < 1$. Thus $1 - \chi(p)p^{-s} \neq 0$ when $|\chi(p)| \leq 1$. Every finite reciprocal Euler product is therefore zero-free wherever it is defined in that half-plane. This does not give compact-uniform convergence to a completed L-function throughout its critical strip. The continuation from the absolutely convergent region is part of the missing analytic argument; replacing it with a finite-product identity would assume the desired conclusion.

1.4 Determinants, cohomology and spectra

A determinant appears in at least four roles in this subject. An ordinary matrix determinant detects singularity. A determinant line of a perfect complex records alternating exterior powers. A regularized determinant encodes an infinite spectrum and depends on a regularization procedure. An Euler product describes arithmetic local factors in a region of convergence before analytic continuation. They can interact through theorems, but their names do not identify them.

The finite positive-pencil proposition concerns the first role. The algebraic theory of complexes naturally supplies the second. Elliptic operators lead to the third. The GRH program seeks a carefully normalized comparison with the fourth. A chain of identifications would have to respect multiplicity, normalization, convergence and the entire analytic domain. The size of the missing work is not measured by how similar the formulas look.

The same discipline matters for Yang–Mills. The index of an elliptic deformation complex is a stable alternating dimension. A mass gap is a strictly positive lower bound on the spectrum of a Hamiltonian above the vacuum. The first survives many continuous deformations that can drive nonzero eigenvalues arbitrarily close to zero. Later chapters give an explicit family proving that the first cannot imply the second without additional estimates.

1.5 A map of rigorous and prospective connections

Starting structure	Established connection	Additional obligation for the requested application
Integral order	Acts on lattices; when realized in $\mathrm{End}(E)$, on the appropriate Tate modules	A metric, efficient presentation and error law for cryptographic use
Elliptic curve	Endomorphism ring is an arithmetic order in the relevant cases	Efficiently known endomorphisms and a proved circuit cost improvement
Algebraic complex	Rank, exactness, homology and determinantal strata	A specified differential-operator realization and exact symbols
Elliptic complex	Fredholm cohomology and topological index	Quantitative coercivity to bound nonzero spectra
Wilson lattice gauge theory	Finite regularized correlation functions	A controlled continuum and infinite-volume construction
Wightman theory	Hilbert-space field theory with spectral axioms	A positive spectral threshold above the vacuum
Positive spectral pencil	Zeros constrained to a line	Exact comparison and analytic convergence to an arithmetic target

Every row will be expanded with formulas. The table also identifies where new research could be meaningful. The objective is a family of explicit bridges with testable claims, rather than a claim that one broad analogy has solved all of the underlying problems.

Chapter 2

Orders as integral arithmetic structures

2.1 The role of the two volumes

The first attached volume is *Lattices over Orders I*, by Klaus W. Roggenkamp and Verena Huber-Dyson, published in 1970 as volume 115 of *Lecture Notes in Mathematics*. Its five chapters develop modules, homological algebra, Morita equivalence and separability, maximal orders, and the Higman ideal with extensions of lattices. The second volume is by Roggenkamp. The first volume is particularly useful here because it explains precisely which integral information disappears upon extending scalars to a field. This is the information that must be retained when one passes between elliptic curves, ideals, finite modules, and geometric realizations of arithmetic lattices. Our notation uses ordinary left-to-right function application, whereas the book often places homomorphisms opposite the scalars. This change of convention must be remembered when comparing endomorphism rings and opposite rings [135].

The corrections printed at the end of Volume II are part of the source, not optional editorial detail. They add a restriction excluding totally definite quaternion matrix components to a cancellation assertion in the introduction, and replace an invalid torsion-decomposition step in the proof at chapter-page IV.13 by an argument using completion and evaluation maps. We use those corrections when interpreting the structural results; the elementary proofs below do not depend on the faulty decomposition [133, Corrections Vol. I].

There is a genuine common language, but not a common theorem that automatically solves all of the problems in this monograph. An order is an integral algebra. A lattice over it is a module satisfying a finiteness and torsion condition. A Euclidean lattice carries additional metric data. A lattice in a lattice gauge theory is a discretization of space or spacetime. These objects can interact through explicit constructions, but the shared word “lattice” supplies no such construction by itself. The present chapter builds the arithmetic objects and specifies the extra structure required at each subsequent step.

The scope of the proofs is deliberate. Elementary statements used in the later bridges are proved below. Structural theorems whose proofs require a substantial additional theory are identified as external results. A proof over an arbitrary field, a theorem over a Dedekind domain, and a theorem about a positive Hilbert-space operator have different hypotheses; replacing one of those settings by another is itself a mathematical obligation.

2.2 Full lattices and orders

Throughout this chapter, let R be a Dedekind domain with fraction field K, and let A be a finite-dimensional separable K-algebra. An R-algebra structure on A is understood through the central embedding $R \subseteq K \subseteq A$. The number-field example is $R = \mathcal{O}_K$, but many elementary arguments need only that R be a noetherian domain. Unless explicitly stated otherwise, modules are unital and finitely generated.

Definition 2.1 (Arithmetic lattice and order)**.** An R-lattice is a finitely generated torsion-free R-module. If V is a finite-dimensional K-vector space, a full lattice in V is an R-submodule $L \subseteq V$ satisfying $KL = V$ and the preceding finiteness condition. An R-order in A is a subring $\Lambda \subseteq A$ containing 1_A, which is a full R-lattice in A. A left Λ-lattice is a left Λ-module whose underlying R-module is an R-lattice.

The word "full" refers to the specified ambient vector space. It does not mean that the lattice is free over Λ. Even its underlying R-module need not be free when R is not a principal ideal domain. The basic Dedekind-domain theorem says that a finitely generated torsion-free R-module is projective. Projectivity over R does not imply projectivity over an arbitrary order Λ. This distinction is central to the first volume's treatment of lattices and extensions [135, Chs. I and IV].

Proposition 2.2 (Common denominators)**.** *Let L and M be full R-lattices in the same finite-dimensional K-space V. There exist nonzero $r, s \in R$ such that $rL \subseteq M$ and $sM \subseteq L$. Consequently $L \cap M$ is full, and both $(L + M)/L$ and $(L + M)/M$ are finitely generated torsion R-modules.*

Proof. Choose finite generating sets $l_1, \dots, l_a$ of L and $m_1, \dots, m_b$ of M. Since $KM = V$, each l_i is a finite K-linear combination of the m_j. A common nonzero denominator for these finitely many coefficients gives $r \in R \setminus \{0\}$ with $rl_i \in M$ for every i, hence $rL \subseteq M$. Interchanging the two lattices gives s. Thus rL lies in $L \cap M$ and spans V, while $L \cap M$ is finitely generated because it is a submodule of the noetherian module L. The quotient $(L + M)/L$ is finitely generated and killed by s; the other quotient is killed by r. □

Over the ring of integers of a number field, a finitely generated torsion module is finite as a set. Over an arbitrary Dedekind domain, that last conclusion requires finite residue fields and must not be silently inserted. Even in the number-field case, finiteness gives no useful complexity bound until the denominators and their bit lengths are controlled.

Proposition 2.3 (A stable integral model)**.** *If V is a finite-dimensional left A-module and Λ is an R-order in A, then V contains a full left Λ-lattice.*

Proof. Choose a K-basis $v_1, \dots, v_n$ of V and put $L = \sum_{i=1}^{n} \Lambda v_i$. Multiplication by Λ preserves L. If $a_1, \dots, a_t$ generate Λ over R, then the finitely many vectors $a_j v_i$ generate L over R. As a submodule of a K-vector space, L is torsion-free. Since $1 \in \Lambda$, the vectors v_i lie in L, so $KL = V$. □

This simple construction explains why rational representations admit integral models, while also explaining their nonuniqueness. Choosing a basis is not the same as classifying all stable lattices. The scalar extension functor $K \otimes_R -$ can send nonisomorphic lattices to isomorphic A-modules.

Example 2.4 (An order that retains a congruence)**.** For a nonzero ideal $\mathfrak{a} \subseteq R$, consider

$$\Lambda_{\mathfrak{a}} = \left\{ \begin{pmatrix} a & b \\ c & d \end{pmatrix} : a, b, d \in R,\ c \in \mathfrak{a} \right\} \subseteq M_2(K).$$

The lower-left entry of a product is $ca' + dc'$, which again belongs to $\mathfrak{a}$. Thus this is an order. Tensoring with K removes the condition $c \in \mathfrak{a}$, giving all of $M_2(K)$. The congruence data remains relevant to integral module theory even though the ambient rational algebra is unchanged.

2.3 Endomorphism orders and faithful models

Proposition 2.5 (Endomorphisms of a full lattice)**.** *Let L be a full R-lattice in V. Each R-linear endomorphism of L extends uniquely to a K-linear endomorphism of V, and*

$$\mathrm{End}_R(L) = \{T \in \mathrm{End}_K(V) : T(L) \subseteq L\}$$

is an R-order in $\mathrm{End}_K(V)$.

Proof. Torsion-freeness identifies L with its image in $K \otimes_R L = V$. Tensoring an endomorphism with K supplies the extension; uniqueness follows because L spans V. The displayed set is closed under addition and composition and contains the identity. Since L is a finite projective R-module, it is a summand of R^N for some N. Its endomorphism module is therefore a summand of, in particular a finite submodule of, $\mathrm{Hom}_R(R^N, R^N)$, so it is finitely generated and torsion-free. Finally, if $T \in \mathrm{End}_K(V)$, common denominators for the images of finitely many generators of L give nonzero $r \in R$ with $rT(L) \subseteq L$. Hence every T belongs to $K\,\mathrm{End}_R(L)$. □

For a full lattice $I \subseteq A$, its two multiplier orders are

$$O_L(I) = \{a \in A : aI \subseteq I\}, \qquad O_R(I) = \{a \in A : Ia \subseteq I\}.$$

The distinction matters in a quaternion algebra: a left ideal of O usually has a different right order. A computation that changes the side without changing the convention can reverse composition or identify the wrong endomorphism ring.

Proposition 2.6 (Multiplier orders)**.** *Both $O_L(I)$ and $O_R(I)$ are R-orders in A.*

Proof. Consider $O_L(I)$. Left multiplication embeds A into $\mathrm{End}_K(A)$, since an element acting as zero sends 1 to zero. The inverse image of $\mathrm{End}_R(I)$ is therefore a torsion-free R-module, and it is finite because it is isomorphic to a submodule of the finite module $\mathrm{End}_R(I)$. It is a ring containing 1. To prove fullness, choose any full lattice $L \subseteq A$. For finite generators of L and I, clear the finitely many denominators needed to put each product into I. This gives nonzero $r \in R$ such that $(rL)I \subseteq I$, so $rL \subseteq O_L(I)$. The same argument using right multiplication proves the other assertion. □

An arithmetic representation can now be studied through three related objects: its module L, the order acting on the left, and the commuting endomorphisms acting on the right. The first volume systematizes this viewpoint through generators and Morita theory [135, Ch. III]. One must nevertheless keep the integral module, not just its rational endomorphism algebra. For example, knowing that the latter is a quadratic field does not determine the conductor of the integral order.

2.4 Localization, support, and what local data detects

For a maximal ideal $\mathfrak{p}$ of R, write $R_{\mathfrak{p}}$ for the localization and $M_{\mathfrak{p}} = R_{\mathfrak{p}} \otimes_R M$. Write $\widehat{R}_{\mathfrak{p}}$ for the completion. These are distinct constructions. A statement proved for complete discrete valuation rings does not automatically hold for an uncompleted local ring. The first volume treats localization and completion separately [135, Ch. I, §§6–9].

Lemma 2.7 (A local membership test)**.** *Let M be an R-module and $x \in M$. If the image of x in $M_{\mathfrak{p}}$ is zero for every maximal ideal $\mathfrak{p}$, then $x = 0$.*

Proof. Suppose $x \neq 0$. Its annihilator $J = \{r \in R : rx = 0\}$ is a proper ideal, hence is contained in a maximal ideal $\mathfrak{p}$. If $x/1 = 0$ in $M_{\mathfrak{p}}$, the definition of localization supplies $s \notin \mathfrak{p}$ with $sx = 0$. This puts $s \in J \subseteq \mathfrak{p}$, a contradiction. □

Proposition 2.8 (Equality from localized inclusions)**.** *If $L, M \subseteq V$ are lattices, then $L \subseteq M$ if and only if $L_{\mathfrak{p}} \subseteq M_{\mathfrak{p}}$ for every maximal ideal $\mathfrak{p}$. In particular, equality of the embedded local lattices implies $L = M$.*

Proof. One implication follows from functoriality. For the converse, apply Lemma 2.7 to the image of each $l \in L$ in $(L + M)/M$. Localization is exact, so all its localized images vanish. □

Equality of embedded lattices is stronger than abstract local isomorphism. To see why, let R be a Dedekind domain with a nonprincipal invertible ideal I. For each $\mathfrak{p}$, the local ring $R_{\mathfrak{p}}$ is a discrete valuation ring and $I_{\mathfrak{p}}$ is a free module of rank one. Thus $I_{\mathfrak{p}} \simeq R_{\mathfrak{p}}$ for every prime. But $I \not\simeq R$ as an R-module: an isomorphism would send 1 to a generator of I, making it principal. The missing data is the ideal class, an obstruction to making the local trivializations compatible globally.

Definition 2.9 (Genus)**.** Two Λ-lattices are in the same genus if their completions are isomorphic as $\widehat{\Lambda}_{\mathfrak{p}}$-modules for every nonzero prime $\mathfrak{p}$. When a local-isomorphism version is used, its equivalence to this definition must be established in the applicable setting.

Roggenkamp's 1969 work on irreducible lattices studies precisely the interaction between maximal overorders, commuting endomorphism orders, and classes within a genus. Its opening proposition also supplies a useful cancellation-of-denominators observation: if two lattices already carry the action of an overorder, a homomorphism linear over the smaller order is linear over the larger one [131, Prop. 1]. We prove that observation in our conventions because it prevents several side-changing errors.

Lemma 2.10 (Linearity over an overorder)**.** *Let $\Lambda \subseteq \Gamma$ be R-orders in A, and let M, N be Γ-lattices. Every Λ-linear map $f : M \to N$ is Γ-linear.*

Proof. By Proposition 2.2, there is nonzero $r \in R$ with $r\Gamma \subseteq \Lambda$. For $\gamma \in \Gamma$ and $m \in M$, R-linearity and Λ-linearity give

$$rf(\gamma m) = f((r\gamma)m) = (r\gamma)f(m) = r\gamma f(m).$$

Torsion-freeness of N permits cancellation of r, giving $f(\gamma m) = \gamma f(m)$. □

The hypothesis that both modules carry the overorder action is essential. It does not assert that every lattice over Λ extends to a lattice over every maximal order containing Λ. Extending the action is itself a structural problem.

2.5 An explicit Morita equivalence

An equivalence of module categories is one of the strongest algebraic bridges available in order theory. It preserves module-theoretic information, but it does not preserve a metric that was never part of the category. The following elementary case makes both facts visible without invoking the general Morita theorem.

Theorem 2.11 (Matrix Morita equivalence)**.** *Let S be a unital ring, $B = M_n(S)$ with $n \geq 1$, and $e = E_{11}$. The functors*

$$F(M) = eM, \qquad G(N) = S^n \otimes_S N$$

give inverse equivalences between left B-modules and left S-modules. Here $eBe \simeq S$, and S^n is the usual (B, S)-bimodule of columns.

Proof. For a left B-module M, define

$$\alpha_M : S^n \otimes_S eM \longrightarrow M, \qquad e_i \otimes m \longmapsto E_{i1}m.$$

The balancing relation is respected because multiplication by an entry of S in a column agrees with the action of the corresponding corner matrix on eM. An inverse is

$$\beta_M(m) = \sum_{i=1}^{n} e_i \otimes E_{1i}m.$$

Each $E_{1i}m$ lies in eM. The identities $\sum_i E_{i1}E_{1i} = \sum_i E_{ii} = 1$ give $\alpha_M\beta_M = 1_M$. On a generator $e_j \otimes m$, $\beta_M\alpha_M$ equals $\sum_i e_i \otimes E_{1i}E_{j1}m = e_j \otimes m$, since $E_{11}m = m$. The maps commute with module homomorphisms and α_M is B-linear, as is checked on matrix units. Finally, $e(S^n \otimes_S N)$ consists of the first-coordinate copy of N, giving a natural isomorphism $FG(N) \simeq N$. These natural isomorphisms prove the equivalence. □

If S is an R-order, these functors restrict to the appropriate finite R-lattices: taking a direct summand preserves finite generation and torsion-freeness, and G is a finite direct sum on underlying R-modules. More general Morita equivalences replace S^n by a finitely generated projective generator, or *progenerator*; their treatment is a central component of Roggenkamp and Huber-Dyson's first volume [135, Ch. III, §§1–2].

To connect this with complexes, apply F degree by degree. Since F is exact, kernels and images are preserved up to canonical isomorphism, and therefore

$$H^i(F(C^\bullet)) \simeq F(H^i(C^\bullet)).$$

This is a categorical statement about homology. It provides no identity between spectra of analytic Laplacians placed on the two complexes after choosing unrelated inner products.

Example 2.12 (The missing metric)**.** The abelian groups $\mathbb{Z}^2$ with quadratic forms $Q_1(x,y) = x^2+y^2$ and $Q_T(x,y) = T^2x^2 + T^{-2}y^2$, $T > 0$, are identical as $\mathbb{Z}$-modules. Their covolumes in the corresponding Euclidean realizations are both one, but their shortest nonzero lengths are 1 and $\min(T, T^{-1})$. No theorem of module equivalence can force those numbers to agree. Any proposed cryptographic or spectral transfer must additionally control the chosen form and the distortion of the transporting maps.

2.6 Trace pairings, dual lattices, and discriminants

Let $t : A \to K$ be a K-linear functional such that $\langle x, y\rangle = t(xy)$ is nondegenerate and $t(xy) = t(yx)$. Important examples are the field trace in a separable field extension and the reduced trace in a central simple algebra. We explicitly assume $t(\Lambda) \subseteq R$ whenever integrality of a pairing on Λ is used. The distinction between an ordinary matrix trace and a reduced trace is relevant in positive characteristic.

Definition 2.13 (Trace dual)**.** For a full lattice $L \subseteq A$, define

$$L^{\#} = \{x \in A : t(xL) \subseteq R\}.$$

When $R = \mathbb{Z}$, $A = K$ is a number field, and $L = \mathcal{O}_K$, the trace dual is the inverse different, also called the codifferent.

Proposition 2.14 (Trace duality)**.** *The lattice $L^{\#}$ is naturally isomorphic to* $\mathrm{Hom}_R(L,R)$ *under* $x \mapsto (y \mapsto t(xy))$*. It is full and* $(L^{\#})^{\#} = L$*.*

Proof. Nondegeneracy identifies A with $\mathrm{Hom}_K(A,K)$. Every R-linear functional $L \to R$ extends uniquely to a K-linear functional on $KL = A$, so the displayed map is a bijection onto $\mathrm{Hom}_R(L,R)$. The latter is a finite projective module of the same rank as L, proving fullness. Since finite projective modules are canonically isomorphic to their double duals, applying the identification twice gives the final equality inside A. The elementary double-dual statement follows first for R^n by coordinates and then for its direct summands. □

There is a bimodule structure compatible with this identification: if $\phi \in \mathrm{Hom}_R(\Lambda, R)$, put $(a\phi b)(x) = \phi(bxa)$. Cyclicity of the trace converts this into left and right multiplication on $\Lambda^{\#}$. The trace dual need not be a ring. Treating it as an overorder without checking multiplicative closure would confuse a dual module with an algebra.

When $R = \mathbb{Z}$ and L has basis $b_1, \dots, b_d$, its trace Gram matrix is $G = (t(b_i b_j))_{ij}$. In these coordinates, $L^{\#}$ has coordinate module $G^{-1}\mathbb{Z}^d$; here symmetry of G follows from the assumption on t. If the trace pairing is integral on L, then

$$[L^{\#} : L] = |\det G|.$$

This equality follows by multiplying coordinates by G: the quotient is isomorphic to $\mathbb{Z}^d/G\mathbb{Z}^d$, whose order is the absolute determinant. The determinant formula is understood for a nondegenerate pairing; a singular Gram matrix does not define a finite dual quotient.

Proposition 2.15 (Change of lattice and discriminant)**.** *Let $L \subseteq M$ be full $\mathbb{Z}$-lattices in A of finite index h. Then*

$$\mathrm{disc}_t(L) = h^2 \,\mathrm{disc}_t(M).$$

Proof. If $C \in M_d(\mathbb{Z})$ expresses a basis of L in a basis of M, their Gram matrices satisfy $G_L = C^{\mathsf{T}} G_M C$. Taking determinants gives $\det G_L = (\det C)^2 \det G_M$. The index is $|\det C|$, yielding the formula. □

This is one of the cleanest quantitative bridges in the subject. The underlying change of integral structure is measured by an index; the bilinear invariant changes by its square. Yet the discriminant alone does not determine the shortest vector or the complexity of a search problem. Different positive forms can share a determinant and have very different minima.

2.7 From arithmetic lattices to Euclidean geometry

For a number field K with r_1 real embeddings and r_2 pairs of complex embeddings, the Minkowski embedding sends x to

$$\sigma(x) = (\sigma_1(x), \dots, \sigma_{r_1}(x), \sqrt{2}\,\mathrm{Re}\,\tau_1(x), \sqrt{2}\,\mathrm{Im}\,\tau_1(x), \dots) \in \mathbb{R}^{r_1+2r_2}.$$

The factor $\sqrt{2}$ is a convention that makes the squared norm $\sum_i |\sigma_i(x)|^2 + 2\sum_j |\tau_j(x)|^2$. The real-vector space isomorphism $K \otimes_{\mathbb{Q}} \mathbb{R} \simeq \mathbb{R}^{r_1} \times \mathbb{C}^{r_2}$ makes the image of any full $\mathbb{Z}$-lattice discrete and cocompact. A different coordinate basis describes the same lattice, whereas a different positive form changes the geometry.

The field trace $\mathrm{Tr}_{K/\mathbb{Q}}(xy)$ is nondegenerate, but it need not be positive definite. In a totally imaginary quadratic field, $\mathrm{Tr}(x\bar{x})$ supplies a positive form, with a specific use of complex conjugation.

In a general number field, the embedding-based Euclidean form is the safe uniform construction. Writing a trace expression without specifying the involution can inadvertently replace a positive form by an indefinite one.

Proposition 2.16 (When a transport controls minima)**.** *Let $T : V \to W$ be a linear isomorphism of finite-dimensional real inner-product spaces with $\dim V > 0$, and suppose $a\|v\| \le \|Tv\| \le b\|v\|$ for all v, with $0 < a \le b$. For a full Euclidean lattice $L \subseteq V$,*

$$a\lambda_1(L) \le \lambda_1(TL) \le b\lambda_1(L),$$

where λ_1 is the length of a shortest nonzero vector.

Proof. The lower bound follows by applying the lower inequality to every nonzero lattice vector and taking the minimum. The upper bound follows by applying the upper inequality to a vector realizing $\lambda_1(L)$. Such a vector exists because a discrete lattice has only finitely many points in any bounded ball. □

This proof also indicates what a useful algorithmic reduction would need: a computable map, polynomially bounded bit complexity, and a controlled ratio b/a. An abstract isomorphism of modules provides none of these quantitative guarantees. They must be proved for the specific encoding and norm used by a cryptographic problem.

2.8 Quaternion orders as four-dimensional arithmetic

Over a field K of characteristic different from two, for $a, b \in K^\times$ a quaternion algebra $(a, b)_K$ has basis $1, i, j, ij$ with $i^2 = a$, $j^2 = b$, and $ji = -ij$. The standard involution sends

$$x = x_0 + x_1 i + x_2 j + x_3 ij \quad \text{to} \quad \bar{x} = x_0 - x_1 i - x_2 j - x_3 ij.$$

Its reduced trace and norm are

$$\operatorname{trd}(x) = x + \bar{x} = 2x_0, \qquad \operatorname{nrd}(x) = x\bar{x} = x_0^2 - ax_1^2 - bx_2^2 + abx_3^2.$$

The presentation is a convenient coordinate system, not a guarantee that the displayed $\mathbb{Z}$-span is maximal when $K = \mathbb{Q}$.

Proposition 2.17 (Multiplicativity of the quaternion norm)**.** *For x, y in a quaternion algebra,* $\operatorname{nrd}(xy) = \operatorname{nrd}(x)\operatorname{nrd}(y)$.

Proof. The defining relations give $\overline{xy} = \bar{y}\bar{x}$. As $y\bar{y}$ is a scalar,

$$(xy)\overline{xy} = xy\bar{y}\bar{x} = \operatorname{nrd}(y)x\bar{x} = \operatorname{nrd}(y)\operatorname{nrd}(x).$$

□

For the real Hamilton algebra $\mathbb{H} = (-1, -1)_{\mathbb{R}}$, the reduced norm is the ordinary sum of four squares. A definite rational quaternion algebra becomes $\mathbb{H}$ after tensoring with $\mathbb{R}$, so a full order is a positive four-dimensional Euclidean lattice after choosing coordinates. The three-dimensional trace-zero space is useful for visualizing ternary norm forms. A three-dimensional projection of the full four-dimensional lattice can hide distinct points and shorten projected vectors; it must not be used as a shortest-vector certificate.

Left and right orders of a quaternion ideal describe how its two ends act. Multiplication of ideals is compatible only when those ends match. If I is an (O_1, O_2)-bimodule and J is an (O_2, O_3)-bimodule, then $I \otimes_{O_2} J$ is an (O_1, O_3)-bimodule. The multiplication map from this tensor product into the product lattice IJ requires its own hypotheses to be an isomorphism; it is an isomorphism in the invertible-ideal setting used for Morita transport.

2.9 Maximality, separability, and their limits

Definition 2.18. An order is maximal if it is not properly contained in any other R-order in the same K-algebra. An R-algebra Λ is separable if it is projective as a module over $\Lambda^e = \Lambda \otimes_R \Lambda^{\mathrm{op}}$, with the action $(a \otimes b^{\mathrm{op}})x = axb$.

The separability of A over K does not imply the separability of an integral order over R. Ramification is an obstruction. Roggenkamp's short paper *A Remark on Separable Orders* characterizes separable orders over rings of integers using maximality together with unramified conditions on the centers and the simple algebra components. This result is a useful warning: "maximal" and "separable over the base ring" cannot be used interchangeably [132].

For the main arithmetic setting, the first volume develops the existence and local structure of maximal orders in separable algebras, and their relationship with progenerators [135, Ch. IV]. Applying those results to an algorithm requires a representation of the order, arithmetic in its coefficient ring, and a procedure for recognizing membership. The abstract existence theorem does not itself construct a basis with useful complexity guarantees.

There is also a conceptual boundary worth stating precisely. An order can support a category of modules, a Grothendieck group of projectives, and finite complexes. Those are algebraic invariants. An elliptic operator requires a smooth manifold, vector bundles, and a principal symbol satisfying an exactness or invertibility condition away from the zero cotangent vector. A quantum field theory requires further analytic and probabilistic data. The intervening structures are the subject of subsequent chapters; they cannot be inferred merely from the existence of a maximal order.

2.10 A finite certificate for maximal overorders

There is a useful elementary special case of the existence theory. Assume now $R = \mathbb{Z}$ and that A carries the nondegenerate trace pairing used above, with the additional property that $t(\Gamma) \subseteq \mathbb{Z}$ for every order Γ in A. This holds for the field trace on number fields and for reduced trace on rational quaternion algebras. The trace-integrality fact uses that an element of an order is integral and that the relevant trace of an integral element is an integer; it is part of the separable-algebra machinery in the first volume [135, III, §6; IV, Theorem 1.4].

Proposition 2.19 (Bounding all overorders)**.** *If Λ is an order and Γ is an order containing it, then $\Lambda \subseteq \Gamma \subseteq \Lambda^{\#}$. In particular, there are only finitely many overorders of Λ, and at least one of them is maximal.*

Proof. For $\gamma \in \Gamma$ and $\lambda \in \Lambda$, their product belongs to Γ, so $t(\gamma\lambda) \in \mathbb{Z}$. This is exactly the condition $\gamma \in \Lambda^{\#}$. The quotient $\Lambda^{\#}/\Lambda$ is a finite abelian group. An intermediate lattice is determined by its subgroup in this quotient, and a finite group has finitely many subgroups. Overorders are a subset of those intermediate lattices; hence they form a finite nonempty partially ordered set. A maximal element Γ of this set cannot have a proper overorder, because any such overorder would still contain Λ. □

This argument gives a finite verification strategy for small examples. Compute a trace-dual basis, enumerate the relevant finite quotient's subgroups, lift each subgroup to a lattice, and test that the lifted lattice contains the identity and is closed under multiplication. A multiplication test requires only the products of basis elements: bilinearity then covers every pair of elements. Each product can be expressed exactly in the candidate basis, and membership is equivalent to integral coordinates.

The proposition is an existence and certification statement, not a polynomial-time maximal-order algorithm. The number of subgroups can be large, and the quotient size depends on the discriminant. A complexity theorem would have to control those quantities in terms of the bit length of the input. The general existence theorem in the book is IV.4.6; local detection of maximality is IV.4.8, and the hereditary property is IV.4.19 [135]. These are separate theorems, not interchangeable definitions.

2.11 Averaging and integral extension obstructions

The first volume's final chapter studies projective homomorphisms and ideals that annihilate extensions. A finite-group example shows the mechanism directly and also explains why a denominator that is harmless over a field can matter over an integral base [135, Ch. V].

Proposition 2.20 (An integral averaging lift)**.** *Let G be a finite group, let R be a commutative ring, and let*

$$0 \longrightarrow N \longrightarrow M \xrightarrow{q} L \longrightarrow 0$$

be an exact sequence of left $R[G]$-modules that splits as a sequence of R-modules. Then there is an $R[G]$-linear map $T : L \to M$ with $qT = |G|\, 1_L$.

Proof. Choose an R-linear section $s : L \to M$ of q, and define

$$T(l) = \sum_{g \in G} g\, s(g^{-1} l).$$

The map is R-linear. For $h \in G$, replacing g by hu in the sum gives

$$T(hl) = \sum_{u \in G} hu\, s(u^{-1} l) = hT(l),$$

so T is $R[G]$-linear. Since q is G-equivariant and $qs = 1_L$, each summand satisfies $q(g\, s(g^{-1} l)) = g(g^{-1} l) = l$. Summing gives $qT(l) = |G| l$. □

If $|G|$ is a unit in R, dividing T by $|G|$ produces an equivariant section, so the original exact sequence splits as $R[G]$-modules. This is the averaging mechanism behind Maschke's theorem in the field case. If $|G|$ is not invertible, the proof still supplies a lift of a multiple of the identity, but not necessarily a section. This residual multiple is precisely the kind of information measured by annihilators of extension classes.

For completeness, one can state the extension consequence without concealing the role of the lift. Pull the sequence back along the map $|G| 1_L : L \to L$. The pullback module is

$$P = \{(m, l) \in M \oplus L : q(m) = |G| l\}.$$

The map $l \mapsto (T(l), l)$ is an $R[G]$-linear section of $P \to L$. Thus this pullback extension is split. Under the standard identification of extension classes with $\mathrm{Ext}^1_{R[G]}(L, N)$, pullback by scalar multiplication is multiplication by that scalar, so $|G|$ annihilates the extension class. The identification itself is developed in the book's Chapter II, §5.

The condition that the sequence split over R holds in the arithmetic-lattice setting when L is projective as an R-module. It need not hold for arbitrary R-modules. The argument therefore records its exact input rather than silently extending an assertion about lattices to all modules.

Example 2.21 (A denominator that cannot be removed)**.** Let k have characteristic p, and let $G = C_p$ have generator g. The two-dimensional representation with $g \mapsto \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ is well-defined because $(1+N)^p = 1$ when $N^2 = 0$. Its first-coordinate line is a trivial subrepresentation, and the quotient is also trivial. There is no invariant complement: a vector $(a, 1)$ spanning a candidate complement is sent to $(a+1, 1)$, which is not in its span. Thus the extension is nonsplit. Here $|G| = p = 0$ in k, so the averaging proposition does not supply a nonzero multiple of the identity that can be inverted.

This example concerns an exact sequence and a representation. It does not assert a physical mass gap, nor a cryptographic hardness gap. Its contribution to the broader program is a precise obstruction: a rational or generic-fiber splitting can fail integrally at primes dividing a denominator. Such failure must be retained in any proposed bridge that uses integral complexes or reductions modulo primes.

Chapter 3

Elliptic curves through their endomorphism lattices

3.1 Two arithmetic questions attached to the same curve

An elliptic curve has a group of points and a ring of endomorphisms. These are different mathematical objects, and cryptographic problems can use either one. In conventional elliptic-curve cryptography, one usually fixes a large cyclic subgroup $\langle P \rangle$ of $E(\mathbb{F}_q)$ and studies the recovery of d from P and $[d]P$. In isogeny constructions, one studies maps between curves, endomorphism rings, or actions of ideal class groups on suitable families. The theory of orders speaks especially directly to the second collection of objects. It also clarifies arithmetic used in point counting and efficient scalar multiplication, but a description of an endomorphism ring is not in itself a discrete-logarithm algorithm.

To keep these roles separate, write $\mathrm{End}_k(E)$ for maps defined over k, $\mathrm{End}_{\bar{k}}(E)$ for geometric endomorphisms, and $\mathrm{End}_k^0(E) = \mathrm{End}_k(E) \otimes_{\mathbb{Z}} \mathbb{Q}$. Over finite fields, confusing the first two rings can turn a quadratic order into a claimed quaternion order with the wrong field of definition. Sutherland explicitly emphasizes this distinction in his treatment of ordinary and supersingular curves [161].

3.2 Complex tori and scalar endomorphisms

Let $L = \mathbb{Z}\omega_1 + \mathbb{Z}\omega_2 \subseteq \mathbb{C}$ with ω_1, ω_2 linearly independent over $\mathbb{R}$. The quotient $\mathbb{C}/L$ is a compact one-dimensional complex Lie group. The analytic uniformization theorem identifies such groups with complex elliptic curves; that theorem is an external input here. The following description of their homomorphisms has a short direct proof once the elementary lifting property of the covering map $\mathbb{C} \to \mathbb{C}/L$ is available.

Proposition 3.1 (Homomorphisms of complex tori)**.** *Every holomorphic group homomorphism $f : \mathbb{C}/L \to \mathbb{C}/M$ is induced by a unique scalar $\alpha \in \mathbb{C}$ satisfying $\alpha L \subseteq M$. It is an isomorphism if and only if $\alpha L = M$.*

Proof. Lift f to a holomorphic map $F : \mathbb{C} \to \mathbb{C}$ with $F(0) = 0$. The group-homomorphism property implies $F(z + w) - F(z) - F(w) \in M$. This is a continuous function of (z, w) taking values in the discrete set M, so it is constant; at $(0, 0)$ it is zero. Thus F is additive. Differentiating $F(z + w) = F(z) + F(w)$ with respect to w at zero gives $F'(z) = F'(0) = \alpha$. Therefore $F(z) = \alpha z$. For $\lambda \in L$, the points z and $z + \lambda$ have the same image in the source torus, so $\alpha\lambda \in M$. Conversely,

that containment makes $z + L \mapsto \alpha z + M$ well-defined and holomorphic. A nonzero scalar has an inverse on the quotient exactly when $\alpha L = M$. The lift fixing zero is unique, so α is unique. □

In particular,

$$\mathrm{End}(\mathbb{C}/L) = O(L) := \{\alpha \in \mathbb{C} : \alpha L \subseteq L\}.$$

This is the multiplier-order construction from the previous chapter, now restricted to scalar multiplication on the plane. The connection is literal: a holomorphic map of the torus is recorded by an arithmetic operation preserving its period lattice.

Theorem 3.2 (The elementary complex-multiplication dichotomy)**.** *For a full lattice $L \subseteq \mathbb{C}$, either $O(L) = \mathbb{Z}$ or $O(L)$ is an order in an imaginary quadratic field.*

Proof. Multiplication by $\alpha \in O(L)$ is represented in a $\mathbb{Z}$-basis of L by a matrix $A_\alpha \in M_2(\mathbb{Z})$. Its real-linear eigenvalues, after complexification, are α and $\bar{\alpha}$. Consequently α satisfies $X^2 - \mathrm{tr}(A_\alpha)X + \det A_\alpha = 0$. If α is real, multiplication on the underlying real plane is α times the identity in every real basis. Its matrix entries are integral, so $\alpha \in \mathbb{Z}$.

If all elements are real, we have obtained $O(L) = \mathbb{Z}$. Otherwise choose nonreal $\alpha \in O(L)$. Its quadratic polynomial is irreducible over $\mathbb{Q}$, since a real quadratic polynomial with a rational root has both roots real. Put $K = \mathbb{Q}(\alpha)$, an imaginary quadratic field. Every $\beta \in O(L)$ commutes with α, so A_β lies in the centralizer of A_α in $M_2(\mathbb{Q})$. This centralizer is $\mathbb{Q}[A_\alpha]$: choose any nonzero $v \in \mathbb{Q}^2$; $v, A_\alpha v$ form a basis, and a commuting operator is uniquely determined by its value on v. Hence the centralizer has dimension two, and it already contains the two-dimensional algebra $\mathbb{Q}[A_\alpha]$.

It follows that every β belongs to K. The additive group $O(L)$ embeds in $M_2(\mathbb{Z})$ and is therefore finitely generated. Its rational span is K, since it contains 1 and α. Thus it is a full $\mathbb{Z}$-lattice and a ring containing 1, which is exactly an order in K. □

This argument proves an algebraic classification of scalar endomorphisms. It does not prove that every finite-field elliptic curve can be lifted with all of its geometric endomorphisms to characteristic zero. In particular, a quaternion algebra cannot embed into the scalar algebra $\mathbb{C}$. The supersingular case below therefore needs a different geometric setting.

3.3 Quadratic orders and the conductor

Let K be a quadratic number field with ring of integers $\mathcal{O}_K = \mathbb{Z} + \mathbb{Z}\omega$. A subring finite over $\mathbb{Z}$ consists of algebraic integers, so every order $O \subseteq K$ is contained in $\mathcal{O}_K$. The following description gives the integral datum that is lost on passing from O to $O \otimes \mathbb{Q}$.

Theorem 3.3 (Quadratic conductor classification)**.** *Every order in K has the form*

$$O_f = \mathbb{Z} + f\mathcal{O}_K = \mathbb{Z} + \mathbb{Z}(f\omega)$$

for a unique integer $f \geq 1$. Its conductor $\{x \in \mathcal{O}_K : x\mathcal{O}_K \subseteq O_f\}$ equals $f\mathcal{O}_K$, and $\mathrm{disc}(O_f) = f^2 \mathrm{disc}(\mathcal{O}_K)$.

Proof. Since $\mathbb{Z} \subseteq O$, the subgroup $O/\mathbb{Z}$ of $\mathcal{O}_K/\mathbb{Z} \simeq \mathbb{Z}$ is $f\mathbb{Z}$ for a unique $f \geq 1$; it is nonzero because O has rank two. Therefore $O = \mathbb{Z} + \mathbb{Z}(f\omega)$. Conversely, if $\omega^2 = t\omega - n$ with $t, n \in \mathbb{Z}$, then $(f\omega)^2 = ft(f\omega) - f^2 n$, proving closure under multiplication for every f.

Certainly $f\mathcal{O}_K$ lies in the conductor. Write an element of the conductor as $x = a + b\omega$. Since $x = x \cdot 1 \in O_f$, we have $f \mid b$. Since $x\omega = -bn + (a + bt)\omega$ also lies in O_f, we have $f \mid a + bt$, hence $f \mid a$. Thus $x \in f\mathcal{O}_K$, giving the reverse containment. The change from the basis $(1, \omega)$ to $(1, f\omega)$

has determinant f, so the discriminant statement follows from the change-of-lattice formula proved in the previous chapter. □

Example 3.4 (A noninvertible ideal in a nonmaximal order)**.** Let $O = \mathbb{Z}[2i] = \mathbb{Z}+2\mathbb{Z}i$ and $I = 2\mathbb{Z}[i]$. Then I is an O-ideal and $O(I) = \mathbb{Z}[i]$, which strictly contains O. If I were invertible as an O-ideal with inverse J, the containment $xI \subseteq I$ would imply $xO = xIJ \subseteq IJ = O$, hence $x \in O$. The element $i \in O(I) \setminus O$ contradicts this. Thus the ideal is not invertible over O, despite being a valid full arithmetic lattice.

This example pinpoints a necessary qualification in class-group arguments. One cannot form a group from all nonzero ideals of an arbitrary order merely by declaring multiplication. The invertible fractional ideals do form a group; the noninvertible ideals are additional module-theoretic data.

3.4 Ideal classes as a torsor of elliptic curves

Let O be an imaginary quadratic order. Denote by $\mathrm{Pic}(O)$ the group of invertible fractional O-ideals modulo principal fractional ideals. For an invertible ideal I, its multiplier ring is O by the argument in the preceding example, so $\mathbb{C}/I$ has endomorphism ring O. The usual complex-multiplication correspondence identifies the proper ideal classes with the corresponding complex elliptic curves; the general correspondence is an external theorem, while the following action on invertible ideal lattices is elementary [157].

Proposition 3.5 (The ideal-class torsor)**.** *On homothety classes of invertible fractional O-ideals, the rule*

$$[\mathfrak{a}] \cdot [I] = [\mathfrak{a}^{-1} I]$$

defines a free and transitive action of $\mathrm{Pic}(O)$*.*

Proof. Changing $\mathfrak{a}$ or I by a nonzero scalar changes $\mathfrak{a}^{-1}I$ by a nonzero scalar, so the rule is well-defined. Multiplication of ideals is associative and commutative, so it gives a group action. If $[\mathfrak{a}^{-1}I] = [I]$, multiply the equality of ideal classes by $[I]^{-1}$ to get $[\mathfrak{a}] = [O]$; thus the action is free. Given $[I]$ and $[J]$, the class $[\mathfrak{a}] = [I][J]^{-1}$ satisfies $[\mathfrak{a}^{-1}I] = [J]$, proving transitivity. □

A torsor is a set with a free transitive group action, not a group with a specified identity. Choosing a base curve identifies its isomorphism classes with a class group abstractly. That choice does not automatically make discrete logarithms, ideal recovery, or group multiplication on curve encodings computationally easy. In algorithmic discussions, this distinction separates an efficient group action from an efficiently represented group.

For an integral invertible ideal $\mathfrak{a} \subseteq O$, $I \subseteq \mathfrak{a}^{-1}I$, and the quotient map induces an isogeny of complex tori

$$\phi_{\mathfrak{a}} : \mathbb{C}/I \longrightarrow \mathbb{C}/(\mathfrak{a}^{-1}I).$$

Its kernel is exactly the finite group $\mathfrak{a}^{-1}I/I$. This follows by testing which cosets $z + I$ become zero in the target. For a principal ideal (α), the kernel size is $[I : \alpha I]$, equal to the absolute determinant of multiplication by α and hence to $|N_{K/\mathbb{Q}}(\alpha)|$. The standard general formula is $\deg \phi_{\mathfrak{a}} = [O : \mathfrak{a}]$ for invertible integral ideals; its extension from principal to invertible ideals uses local freeness and multiplication of ideal norms [160].

Example 3.6 (A visible endomorphism of the square torus)**.** For $I = \mathbb{Z}[i]$, multiplication by $1 + i$ has matrix

$$\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$$

in the basis $(1, i)$. Its determinant is two. It therefore gives an isogeny of degree two and a kernel of order two. The matrix also shows that multiplication rotates by $\pi/4$ and scales length by $\sqrt{2}$. This is an example where the integral operation, the geometric norm, and the degree can all be compared directly. The comparison works because the metric and embedding have been explicitly supplied.

3.5 Finite-field endomorphisms and Frobenius

The geometric classification is a foundational external theorem: in characteristic zero an elliptic curve has geometric endomorphism ring $\mathbb{Z}$ or an imaginary quadratic order. Over an algebraic closure of a finite field, an ordinary elliptic curve has an imaginary quadratic endomorphism order, while a supersingular curve has a maximal order in the rational quaternion algebra ramified at p and infinity. The degree is the reduced norm in the quaternion case. These statements require the theory of isogenies and cannot be deduced from the complex-torus argument alone [169, Theorem 42.1.9].

If E is defined over $\mathbb{F}_q$, its q-power Frobenius π satisfies

$$\pi^2 - [t]\pi + [q] = 0, \qquad t = q + 1 - \#E(\mathbb{F}_q), \qquad |t| \le 2\sqrt{q}.$$

The characteristic equation and Hasse's bound are external inputs from elliptic-curve theory [158, 159]. Once they are available, several useful arithmetic conclusions are immediate, and it is useful to isolate them from those inputs.

Proposition 3.7 (The Frobenius order)**.** *Suppose $E/\mathbb{F}_q$ is ordinary, and put $K = \mathbb{Q}(\pi)$. Then*

$$\mathbb{Z}[\pi] \subseteq \mathrm{End}_{\mathbb{F}_q}(E) \subseteq \mathrm{End}_{\overline{\mathbb{F}}_q}(E) \subseteq \mathcal{O}_K.$$

Moreover, if an intervening order O has index $h = [O : \mathbb{Z}[\pi]]$, then $t^2 - 4q = h^2 \operatorname{disc}(O)$.

Proof. Frobenius and integer multiplications are defined over $\mathbb{F}_q$, giving the first inclusion. Extending the field of definition gives the second. The ordinary classification identifies the geometric endomorphism algebra with an imaginary quadratic field; it contains the nonrational Frobenius and is therefore K. Endomorphisms are integral, giving the final inclusion. The trace and norm of π are t and q. The trace Gram matrix in the basis $(1, \pi)$ is

$$\begin{pmatrix} 2 & t \\ t & t^2 - 2q \end{pmatrix},$$

whose determinant is $t^2 - 4q$. The index-squared discriminant formula now proves the assertion. □

This proposition provides an explicit bridge to the conductor classification. It constrains which orders can contain the Frobenius order. It does not determine which one is the full endomorphism ring without further information. Nor does knowing $\#E(\mathbb{F}_q)$ produce discrete logarithms in a large prime-order subgroup. Point counting and discrete logarithms ask different questions even though both use the same curve.

3.6 Tate modules as arithmetic lattices

Let $\ell \neq p$ be a prime. Over an algebraic closure, $E[\ell^n]$ is abstractly isomorphic to $(\mathbb{Z}/\ell^n\mathbb{Z})^2$. The inverse limit under multiplication by ℓ is the Tate module

$$T_\ell(E) = \varprojlim_n E[\ell^n],$$

a free $\mathbb{Z}_\ell$-module of rank two. Endomorphisms act on every torsion level and therefore on this inverse limit. This is a direct meeting of the completed-lattice theory in the first volume with elliptic-curve arithmetic.

Lemma 3.8 (Faithfulness on prime-to-characteristic torsion)**.** *If an endomorphism α annihilates $E[\ell^n]$ for every $n \geq 1$, then $\alpha = 0$.*

Proof. A nonzero endomorphism is a finite isogeny, so its kernel has at most $\deg \alpha$ geometric points. But $E[\ell^n]$ has ℓ^{2n} geometric points for every n. If all these groups lie in the kernel, their cardinalities eventually exceed $\deg \alpha$, a contradiction. □

Thus endomorphisms can be represented by compatible integral matrices on Tate modules. Obtaining a full description of the endomorphism ring from Galois-equivariant Tate-module maps is a deeper theorem of Tate, not merely the faithfulness lemma above [164]. The distinction mirrors the difference between an injective map into a matrix algebra and a theorem identifying its image.

For supersingular E with all geometric endomorphisms defined, the maximal order has $\mathbb{Z}$-rank four. Away from p, its completed algebra is split, and its action on prime-power torsion relates it to $M_2(\mathbb{Z}_\ell)$. At p, the algebra is a division quaternion algebra. The exceptional prime is not a minor implementation detail: the geometry of inseparable isogenies and the local ramification require a different description. This explains why replacing all local factors by the same 2×2 matrix model loses arithmetic information.

3.7 Quaternion ideals and the Deuring correspondence

Fix a supersingular curve $E_0/\overline{\mathbb{F}}_p$ and put $O = \mathrm{End}(E_0)$, $B = O \otimes_{\mathbb{Z}} \mathbb{Q}$. The Deuring correspondence relates supersingular elliptic curves to ideal lattices over O. Its module convention is crucial. The group $\mathrm{Hom}(E_0, E)$ is a *right* O-module by precomposition:

$$f \cdot \alpha = f \circ \alpha.$$

Postcomposition by an endomorphism of E acts on the left. With suitable identifications in the common rational quaternion space, the Hom group becomes a full right O-ideal. In this convention the other endomorphism ring is its left multiplier order. The categorical correspondence is an external theorem; Voight's treatment develops its ideals and morphisms explicitly [169, §42.3].

The next elementary lemma makes the multiplier identification transparent once the correspondence has supplied the ideal.

Lemma 3.9 (Endomorphisms of a right ideal lattice)**.** *Let O be an order in a quaternion algebra B, and let $I \subseteq B$ be a full right O-ideal. Then*

$$\mathrm{End}_O(I) \simeq O_L(I).$$

The map sends $a \in O_L(I)$ to left multiplication by a.

Proof. Left multiplication preserves I exactly for elements of $O_L(I)$ and commutes with right multiplication by O, so the map is well-defined. It is injective because I spans B. For surjectivity, extend $f \in \mathrm{End}_O(I)$ to a $\mathbb{Q}$-linear map $f_{\mathbb{Q}} : B \to B$. Since O spans B, the extension is right B-linear. A right B-linear endomorphism of B is left multiplication by $a = f_{\mathbb{Q}}(1)$, because $f_{\mathbb{Q}}(x) = f_{\mathbb{Q}}(1 \cdot x) = ax$. Its preservation of I says $a \in O_L(I)$. Composition of left multiplications agrees with multiplication in $O_L(I)$, so this is a ring isomorphism. □

One must also distinguish ideal classes from conjugacy classes of maximal orders. Several ideal classes may have isomorphic multiplier orders. Recording only an order up to conjugacy can forget which ideal class, and therefore which object in the correspondence, was present. This is an arithmetic instance of a general principle: a derived invariant can classify less information than the object from which it was computed.

3.8 The quaternion norm and the meaning of degree

For an elliptic-curve endomorphism α, the dual isogeny $\widehat{\alpha}$ satisfies $\widehat{\alpha}\alpha = \alpha\widehat{\alpha} = [\deg \alpha]$. These identities make the degree behave like a norm. In the quaternion endomorphism algebra, the dual isogeny is the standard involution, and degree agrees with the reduced norm. For a nonzero quasi-endomorphism α/n, the norm is $\deg(\alpha)/n^2$. Thus passage from integral maps to rational quasi-isogenies changes the integrality, while retaining the norm law.

Proposition 3.10 (Degree multiplicativity from duality)**.** *Assume the dual-isogeny identities and* $\widehat{\beta\alpha} = \widehat{\alpha}\widehat{\beta}$*. Then* $\deg(\beta\alpha) = \deg\beta \deg\alpha$*.*

Proof. Compute

$$\widehat{\beta\alpha}(\beta\alpha) = \widehat{\alpha}\widehat{\beta}\beta\alpha = \widehat{\alpha}[\deg\beta]\alpha = [\deg\beta \deg\alpha].$$

On the other hand, the dual-isogeny identity for $\beta\alpha$ makes the left side $[\deg(\beta\alpha)]$. The map $\mathbb{Z} \to \mathrm{End}(E)$, $n \mapsto [n]$, is injective, so the integers are equal. □

This proof illustrates the appropriate form of a conditional bridge: its assumptions are already strong theorems of elliptic-curve geometry, and the conclusion is a short algebraic consequence. It would be misleading to present the short calculation as a replacement for the construction of the dual isogeny or the proof of its properties.

For numerical visualization, a positive definite quaternion norm can display shells $\mathrm{nrd}(x) = c$ in four real coordinates. A sequence of three-dimensional slices, or a projection controlled by a fourth-coordinate slider, can reveal how an ideal lattice meets those shells. Such a picture helps explain why a norm equation is a Diophantine constraint. A rendered shell is not evidence that an unshown integer point exists, and a point drawn with floating-point error is not a verified ideal element.

3.9 What these bridges establish for cryptography

The arithmetic theory now supplies several precise translations. Complex multiplication translates scalar endomorphisms into imaginary quadratic orders. Ideal multiplication translates into actions and isogenies when the necessary invertibility hypotheses hold. Supersingular endomorphisms translate into maximal quaternion orders. Hom groups translate into integral bimodules. Prime-to-characteristic torsion translates part of the theory into compatible finite matrix modules.

Each translation has a different computational cost. An abstract isomorphism, an efficiently computable isomorphism, and a reduction preserving a promised distribution are three separate claims. To apply the translations to cryptographic security, one must specify which of them is actually available. In particular, an algorithm given an explicit endomorphism basis need not solve the problem of obtaining that basis from a curve equation. A class-group action given an ideal factorization need not reveal the factorization from the resulting curve. A statement about all ideals is not a statement about the small ideals selected by a protocol's distribution.

The most useful way to state a proposed new algorithm is therefore as a tuple of mathematical data. The input type specifies a curve, an order, a basis, a subgroup, or an ideal. The promise specifies ordinary or supersingular, maximal or nonmaximal, and any supplied orientation. The output specifies a scalar, a morphism, a basis, or a witness to equivalence. The cost model counts bit operations or fault-tolerant quantum resources. The proof then separates correctness from cost and from any average-case assumption.

These requirements expose several invalid shortcuts. The fact that an order is a lattice does not reduce arbitrary lattice cryptography to an elliptic-curve problem. A four-dimensional quaternion lattice is not the same family as a growing-dimensional module lattice. A positive norm on an endomorphism algebra is not a spectral gap for a quantum field theory. An exact functor preserving module homology does not by itself supply a reversible quantum circuit. Later chapters can use the constructions above, but each claimed transfer must carry its own missing hypotheses.

3.10 A worked comparison of three rank-two objects

It is instructive to hold the rank fixed while changing the structure. First consider $L_1 = \mathbb{Z}^2$ as an abelian group. Its automorphisms are $\mathrm{GL}_2(\mathbb{Z})$, and it has no preferred notion of length. Next equip the same module with the standard dot product. Its isometries are the much smaller group $\mathrm{GL}_2(\mathbb{Z}) \cap O(2)$, consisting of signed permutation matrices. Finally identify it with $\mathbb{Z}[i] \subseteq \mathbb{C}$ and remember multiplication by i. Its module endomorphisms commuting with i form the ring $\mathbb{Z}[i]$.

Proposition 3.11 (Centralizer of the quarter-turn)**.** *Let $J = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$. An integral matrix commutes with J if and only if it has the form $\left(\begin{smallmatrix} a & -b \\ b & a \end{smallmatrix}\right)$ with $a, b \in \mathbb{Z}$. Its determinant is $a^2 + b^2$.*

Proof. Write the matrix as $T = \left(\begin{smallmatrix} u & v \\ w & z \end{smallmatrix}\right)$. Equating entries in $TJ = JT$ gives $z = u$ and $v = -w$. Set $a = u$, $b = w$. The determinant formula follows by direct expansion. Conversely, any displayed matrix commutes with J by multiplication. □

The determinant is the norm of $a + bi$, and the matrix is a unit precisely when $a^2 + b^2 = 1$. Hence the scalar automorphisms of the square complex torus are $\{1, -1, i, -i\}$. The same underlying rank-two group has therefore supported three different automorphism questions. This small example is a useful check whenever the monograph passes between algebraic, metric, and complex-analytic structures.

3.11 An exact arithmetic certificate for a quadratic order

The order $O_f = \mathbb{Z} + f\mathbb{Z}i$ provides a complete small certificate that can be checked without a numerical approximation. In the basis $b_1 = 1$, $b_2 = fi$, its multiplication table is

$$b_1 b_1 = b_1, \quad b_1 b_2 = b_2 b_1 = b_2, \quad b_2 b_2 = -f^2 b_1.$$

Every structure constant is an integer, so the table certifies closure under multiplication. The field trace sends $a + bi$ to $2a$, and the trace matrix is

$$G_f = \begin{pmatrix} 2 & 0 \\ 0 & -2f^2 \end{pmatrix}, \qquad \det G_f = -4f^2.$$

Its negative determinant is not an error: the trace pairing $\mathrm{Tr}(xy)$ is indefinite on this imaginary quadratic field. The positive pairing $\mathrm{Tr}(x\bar{y})$ has matrix $\mathrm{diag}(2, 2f^2)$ instead. These two matrices encode different bilinear structures on the same order.

Proposition 3.12 (Trace dual of the Gaussian suborder)**.** *For every integer $f \geq 1$,*

$$O_f^{\#} = \tfrac{1}{2}\mathbb{Z} + \tfrac{i}{2f}\mathbb{Z}, \qquad [O_f^{\#} : O_f] = 4f^2.$$

Proof. Write $x = a + bi$ with $a, b \in \mathbb{Q}$. The conditions $\mathrm{Tr}(x \cdot 1) \in \mathbb{Z}$ and $\mathrm{Tr}(x \cdot fi) \in \mathbb{Z}$ are $2a \in \mathbb{Z}$ and $-2fb \in \mathbb{Z}$. They are sufficient because $1, fi$ generate O_f additively. This proves the displayed description. In the first coordinate, the quotient $\frac{1}{2}\mathbb{Z}/\mathbb{Z}$ has size two. In the imaginary coordinate, $\frac{1}{2f}\mathbb{Z}/f\mathbb{Z}$ has size $2f^2$. Multiplying their sizes gives $4f^2$. □

For $f = 2$, the order has discriminant -16, its maximal overorder $\mathbb{Z}[i]$ has discriminant -4, and the index is two. The discriminant formula reads $-16 = 2^2(-4)$. The conductor is $2\mathbb{Z}[i]$, which also supplied the noninvertible ideal above. Its multiplier ring is larger than O_2, so no inverse ideal over O_2 can exist. Thus multiplication, discriminant, conductor, trace dual, and failure of invertibility are mutually consistent in one exact example.

An implementation should output these integer matrices and containment witnesses, not merely a decimal determinant. A diagram may plot the corresponding point sets, but the matrix identities are the certificates. The same separation between exact witnesses and geometric display is useful in larger quaternion calculations and in the interactive companions to this monograph.

3.12 A research program that is actually implied

The preceding material supports a concrete, testable program. One can represent orders and their ideals exactly by integral bases, verify multiplication tables, compute trace matrices, and check discriminant-index identities. One can formalize scalar actions and centralizers in finite dimensions before attempting the full geometric classification. One can compare localizations of a specified lattice, keeping the difference between local isomorphism and global equivalence visible. One can design interactive diagrams whose point coordinates are generated from verified arithmetic rather than artistic approximations.

The next mathematical advance would require a specific theorem that the existing translations do not yet provide. Examples are a distortion-controlled reduction between explicit arithmetic norm problems, a sharper resource analysis for an exact quantum arithmetic subroutine, or a new effective bound with fully stated analytic assumptions. Each is a well-posed research target. None follows simply by adjoining the names of order theory, elliptic complexes, quantum algorithms, and Yang–Mills theory.

The didactic value of Roggenkamp and Huber-Dyson's framework is that it forces the integral and categorical data to be stated before any such claim is made. In that sense the books provide a rigorous foundation for investigation: they clarify the objects and identify obstructions that a successful new connection would have to overcome.

Part II

Complexes, topology and spectral information

Chapter 4

From integral modules to complexes and stable invariants

4.1 What the second volume contributes

The second supplied volume is Klaus W. Roggenkamp's *Lattices over Orders II*, published in 1970 as volume 142 of the Lecture Notes in Mathematics. Its organizing questions concern integral modules over orders: how local isomorphisms differ from global isomorphisms, how genera are classified, when cancellation is valid, and when only finitely many indecomposable lattices occur. Chapters VI and VII develop modules, ideals and genera; Chapter VIII constructs Grothendieck and Whitehead groups; Chapters IX and X study hereditary and Bass orders and questions of finite lattice type. These subjects supply a precise language for passing from individual arithmetic objects to stable invariants. They do not identify cryptographic hardness with a topological index or an analytic spectral gap.[133]

There is a particularly direct historical connection to the analytic side of this monograph. The remark following VIII.2.4, printed page 114, discusses Atiyah–Hirzebruch topological K-theory and the correspondence between vector bundles and projective modules. This is an actual connection in the supplied author's text, rather than an analogy inferred from a shared word. Its modern use requires distinguishing several groups that have similar notation. In this chapter these distinctions are part of the mathematical statement, not qualifications added after a conclusion.[133, Chapter VIII, §§1–2]

All results below that have a displayed proof are derived within the stated framework. Deep classification and index results are explicitly identified as external theorems. The purpose is to expose the exact points at which an arithmetic construction can be transported to a complex, a complex can be transported to a stable class, and a stable class can be evaluated by an index. Each transport needs data and hypotheses. A chain of related subjects is not itself a theorem connecting their computational difficulties.

4.2 The exact category of lattices

Let R be a Dedekind domain, let K be its fraction field, and let A be a finite-dimensional separable K-algebra. An R-order $\Lambda \subseteq A$ is a unital subring, finitely generated as an R-module, such that $K\Lambda = A$. A left Λ-lattice is a left Λ-module M that is finitely generated and torsion-free over R. Because R is Dedekind, this is equivalent to M being finitely generated and projective over R.

Projectivity over R must not be confused with projectivity over Λ. The distinction is responsible for much of the homological content of the theory.[133]

Proposition 4.1 (Extension closure of lattices)**.** *Suppose*

$$0 \longrightarrow M' \xrightarrow{u} M \xrightarrow{v} M'' \longrightarrow 0$$

is an exact sequence of finitely generated Λ-modules. If M' and M'' are Λ-lattices, then M is a Λ-lattice. The sequence splits as a sequence of R-modules, although it need not split as a sequence of Λ-modules.

Proof. Finite generation over R follows by taking generators of M', lifting generators of M'', and using exactness. If $r \in R \setminus \{0\}$ and $rm = 0$, then $rv(m) = 0$ in the torsion-free module M'', so $v(m) = 0$. Write $m = u(m')$. Injectivity of u gives $rm' = 0$, hence $m' = 0$ and $m = 0$. Thus M is torsion-free. Finally, projectivity of M'' over R supplies an R-linear section of v. Nothing in this argument produces a Λ-linear section. □

The lattice category therefore has a natural exact structure: its admissible exact sequences are those short sequences that are exact as module sequences and whose three terms are lattices. This permits extension classes without requiring the category itself to contain every kernel and cokernel of every morphism. For instance, multiplication by a nonunit on a nonzero lattice can have a torsion cokernel, which is not a lattice.

Example 4.2 (Why quotient conditions matter)**.** The map $\mathbb{Z} \xrightarrow{p} \mathbb{Z}$, for a prime p, has cokernel $\mathbb{Z}/p\mathbb{Z}$. Both its source and target are $\mathbb{Z}$-lattices, while its cokernel is not. It is therefore incorrect to regard all maps between lattices as producing exact sequences inside the lattice category. The ambient module category and the exact lattice category encode different information.

An extension of M'' by M' is an exact sequence as above, considered up to an isomorphism that is the identity on the two end terms. Its equivalence class lies in $\mathrm{Ext}^1_\Lambda(M'', M')$. When both endpoints are lattices, the preceding proposition shows that every such module extension has a lattice as middle term. Thus ordinary module Ext^1 really does govern extensions within this exact category. The zero class is precisely the class of split extensions. This gives an initial, concrete bridge between the arithmetic category in Roggenkamp's books and the deformation calculations developed later.

4.3 Three different Grothendieck groups

For an exact category $\mathcal{E}$, write $G_0(\mathcal{E})$ for the abelian group generated by isomorphism classes $[M]$ of objects, subject to

$$[M] = [M'] + [M'']$$

whenever $0 \to M' \to M \to M'' \to 0$ is admissibly exact. For the same additive category one can instead impose only direct-sum relations. Denote the resulting group by $K_0^{\mathrm{split}}(\mathcal{E})$. There is a surjection

$$K_0^{\mathrm{split}}(\mathcal{E}) \longrightarrow G_0(\mathcal{E}),$$

because split exact sequences are exact sequences. In general it is not an isomorphism. Roggenkamp makes this direct-sum versus short-exact-sequence distinction at the beginning of Chapter VIII.[133, Definitions VIII.1.2 and VIII.1.6]

For a unital ring Λ, the usual algebraic $K_0(\Lambda)$ is formed from finitely generated projective Λ-modules under direct sums. Since every short exact sequence of projective modules splits when its quotient is projective, the split and exact constructions agree on this category. In contrast,

$$G_0^{\mathrm{lat}}(\Lambda) := G_0(\mathrm{Lat}(\Lambda))$$

uses all lattices and their exact sequences. The natural map $K_0(\Lambda) \to G_0^{\mathrm{lat}}(\Lambda)$ must not be declared an isomorphism without additional hypotheses.

Proposition 4.3 (Universal property)**.** *Let B be an abelian group and let a assign an element $a(M) \in B$ to every object of $\mathcal{E}$. Suppose a is invariant under isomorphism and additive on admissible short exact sequences. There is a unique group homomorphism*

$$\overline{a} : G_0(\mathcal{E}) \longrightarrow B, \qquad \overline{a}([M]) = a(M).$$

Proof. Map the free abelian group on isomorphism classes to B by extending a linearly. Every defining relation $[M] - [M'] - [M'']$ maps to zero by additivity, so the map descends to the quotient. Its values on the generators prescribe it uniquely. □

This simple property is the proper explanation of the usefulness of Grothendieck groups. They collect additive invariants. They do not retain every invariant of an object, and they are not automatic certificates of isomorphism.

Proposition 4.4 (Equality in a split group)**.** *For objects M, N in an additive category, equality of their classes in K_0^{split} holds if and only if there exists an object T such that $M \oplus T \cong N \oplus T$.*

Proof. The implication from a stable isomorphism follows directly from the relation for direct sums. For the reverse implication, form the commutative monoid of isomorphism classes under direct sum. Its group completion consists of pairs (x, y), where (x, y) and (x', y') are equivalent when some z satisfies $x + y' + z = x' + y + z$. Equality of $(M, 0)$ and $(N, 0)$ therefore says that $M + T = N + T$ for some T. This is exactly the asserted stable isomorphism. □

Cancellation would strengthen the conclusion to $M \cong N$. Roggenkamp's local cancellation results and Jacobinski's global cancellation theorem address when such a strengthening is valid. Their hypotheses are substantial. In particular, restrictions involving totally definite quaternion components are relevant precisely where quaternion orders also arise in supersingular elliptic-curve arithmetic. Removing these restrictions would erase part of the arithmetic content.[133, Chapters VI–VIII]

Example 4.5 (An invariant may deliberately forget torsion)**.** Let $G_0(\mathbb{Z})$ use all finitely generated abelian groups. The exact sequence

$$0 \longrightarrow \mathbb{Z} \xrightarrow{n} \mathbb{Z} \longrightarrow \mathbb{Z}/n\mathbb{Z} \longrightarrow 0$$

shows $[\mathbb{Z}/n\mathbb{Z}] = 0$. Every finitely generated abelian group is a sum of a free group and a finite group, so rank induces $G_0(\mathbb{Z}) \cong \mathbb{Z}$. By contrast, the Grothendieck group of the category of finite abelian groups has one independent generator for each prime, represented by $\mathbb{Z}/p\mathbb{Z}$. The category and the permitted exact sequences determine what the invariant remembers. This example is also among the exercises accompanying VIII.1 in the supplied volume.[133]

4.4 Euler classes of complexes

Let $C^\bullet$ be a bounded cochain complex in an abelian category:

$$\cdots \longrightarrow C^{i-1} \xrightarrow{d^{i-1}} C^i \xrightarrow{d^i} C^{i+1} \longrightarrow \cdots, \qquad d^i d^{i-1} = 0.$$

Write $Z^i = \ker d^i$, $B^i = \operatorname{im} d^{i-1}$, and $H^i = Z^i/B^i$. In a Grothendieck group containing these objects define

$$\chi(C^\bullet) = \sum_i (-1)^i [C^i].$$

Theorem 4.6 (Algebraic Euler identity)**.** *Suppose all terms, cycles, boundaries and cohomology objects belong to the exact category in question, and both sequences* $0 \to Z^i \to C^i \to B^{i+1} \to 0$ *and* $0 \to B^i \to Z^i \to H^i \to 0$ *are admissibly exact in its specified exact structure. Then*

$$\sum_i (-1)^i [C^i] = \sum_i (-1)^i [H^i(C^\bullet)].$$

Proof. The two exact sequences

$$0 \to Z^i \to C^i \to B^{i+1} \to 0, \qquad 0 \to B^i \to Z^i \to H^i \to 0$$

give $[C^i] = [B^i] + [H^i] + [B^{i+1}]$. Multiply by $(-1)^i$ and sum. Each boundary term occurs twice with opposite signs. Boundedness makes the sum finite and eliminates endpoint terms after inserting zero objects outside the given range. □

The hypothesis about the category is essential. A complex of lattices can have torsion cohomology, so its cohomology terms need not lie in G_0^{lat}. The identity always makes sense in the Grothendieck group of all finitely generated modules when the ring is noetherian, but this is a different group. Similarly, a complex of projectives need not have projective cohomology, so one cannot insert its cohomology classes directly into $K_0(\Lambda)$ without a suitable resolution construction.

Proposition 4.7 (Bounded acyclic projective complexes)**.** *A bounded acyclic complex of projective modules is contractible. Its Euler class in* $K_0(\Lambda)$ *is zero.*

Proof. Let n be the largest degree with a nonzero term. Acyclicity makes $d^{n-1} : C^{n-1} \to C^n$ surjective. Since C^n is projective, choose a section. Consequently $C^{n-1} \cong Z^{n-1} \oplus C^n$, and the last two terms contain a direct summand equal to the two-term identity complex $C^n \xrightarrow{1} C^n$. Its complement is a shorter acyclic complex of projectives: Z^{n-1} is projective because it is a direct summand of C^{n-1}. Induction decomposes the entire complex into identity complexes. Each has an explicit contracting homotopy and contributes $[P] - [P] = 0$ to the Euler class. □

For a bounded complex of projectives, this proves that the Euler class is invariant under quasi-isomorphism: the mapping cone of a quasi-isomorphism is a bounded acyclic projective complex, and its Euler class is the difference of the two Euler classes. This is the elementary core of the language of perfect complexes. It is enough for several useful arithmetic comparisons without claiming that every lattice has a finite projective resolution.

Example 4.8 (A rationally exact integral complex)**.** For a prime p, consider

$$0 \longrightarrow \mathbb{Z} \xrightarrow{A} \mathbb{Z}^2 \xrightarrow{B} \mathbb{Z} \longrightarrow 0, \quad A(t) = (pt, 0), \quad B(u, v) = v.$$

Then $BA = 0$, $\ker A = 0$, $\operatorname{coker} B = 0$, and the middle cohomology is $\mathbb{Z}/p\mathbb{Z}$. Tensoring with $\mathbb{Q}$ makes the complex exact. Thus a generic-rank calculation over the fraction field detects none of its torsion defect. Modulo p, the first map becomes zero and the two nonzero cohomology groups each have dimension one. The alternating dimension remains zero throughout. Exactness, torsion and Euler characteristic are three different measurements of this example.

4.5 Schanuel's lemma as a comparison mechanism

Theorem 4.9 (Schanuel's lemma)**.** *Suppose*

$$0 \to K \to P \xrightarrow{f} M \to 0, \qquad 0 \to K' \to P' \xrightarrow{f'} M \to 0$$

are exact and P, P' are projective. Then

$$K \oplus P' \cong K' \oplus P.$$

Proof. Form the pullback

$$Q = \{(p, p') \in P \oplus P' : f(p) = f'(p')\}.$$

The projection $Q \to P'$ is surjective because f is surjective, and its kernel is K. Since P' is projective, $Q \cong K \oplus P'$. Interchanging the two presentations gives $Q \cong K' \oplus P$. Composing these isomorphisms proves the assertion. □

This lemma explains why alternating presentation classes are well defined even though the presentation itself is not unique. Roggenkamp uses it in VIII.3.1 to compare the Grothendieck group of lattices with that of all finitely generated modules. The point is not that a torsion module becomes a lattice; it is represented by a difference of two lattice classes.[133, Theorem VIII.3.1]

Proposition 4.10 (A one-step lattice presentation)**.** *Every finitely generated left Λ-module T admits an exact sequence*

$$0 \to M \to P \to T \to 0$$

with P finite free over Λ and M a Λ-lattice.

Proof. Choose finitely many generators of T over Λ to obtain a surjection $P = \Lambda^n \to T$. The ring Λ is noetherian on either side because it is module-finite over the central noetherian ring R, so its kernel M is finitely generated over R. It is an R-submodule of the torsion-free module P, hence is torsion-free. Therefore M is a lattice. □

In particular $[P] - [M]$ is independent of the chosen one-step presentation by Schanuel's lemma. Additivity on exact sequences follows by the standard pullback construction of compatible presentations, yielding the comparison used in the book. This is a genuine arithmetic use of complexes and stable classes. It does not supply a short algorithm for finding module isomorphisms, since the existence of a presentation comparison is separate from the complexity of constructing it.

4.6 What hereditary and Bass hypotheses change

A left hereditary ring is one in which every left ideal is projective; equivalently every submodule of a projective left module is projective. For orders, this sharply simplifies lattice resolutions. A lattice can be embedded into a finite direct sum of copies of the ambient semisimple algebra and, after clearing denominators, into a finite free Λ-module. Under a hereditary hypothesis it is therefore projective. Here the embedding must be justified from the semisimplicity of the generic fiber; torsion-freeness alone over an arbitrary ring would not suffice. Roggenkamp's Chapter IX treats the structure of hereditary orders and explains why their local and global classification still requires care.[133, Chapter IX, §2]

Proposition 4.11 (Embedding a lattice into a free module)**.** *Let Λ be an order in a finite-dimensional semisimple K-algebra A. Every Λ-lattice M embeds into Λ^n for some n.*

Proof. The finite-dimensional A-module $K \otimes_R M$ is a direct summand of A^n for a sufficiently large n, because A is semisimple and its regular representation contains every simple module. Thus there is an injective A-linear map $j : K \otimes_R M \to A^n$. Choose finitely many R-generators $m_1, \dots, m_\ell$ of M. Since $K\Lambda^n = A^n$, each $j(m_i)$ becomes an element of Λ^n after multiplication by a nonzero denominator in R. Taking the product of these denominators gives a common nonzero $r \in R$ for which $rj(M) \subseteq \Lambda^n$. Multiplication by r does not destroy injectivity in the K-vector space A^n. Restricting rj to M gives the desired embedding. □

Corollary 4.12. *If Λ is hereditary, every Λ-lattice is projective. Consequently every short exact sequence of lattices splits, and $K_0(\Lambda) \cong G_0^{\mathrm{lat}}(\Lambda)$.*

Proof. The embedding proposition and the submodule characterization of hereditary rings make each lattice projective. A short exact sequence splits because its quotient is projective. Hence the two categories have the same objects and the same additive relations. □

Bass orders occupy a broader setting. In the standard order-theoretic terminology, an order is Gorenstein when its R-dual is projective on the appropriate opposite side, and it is Bass when every overorder is Gorenstein. The duality, finiteness and local hypotheses must be specified when using this terminology outside orders. Roggenkamp's 1971 paper computes Grothendieck groups of Bass orders over complete discrete valuation rings with finite residue field and distinguishes an exceptional local type. The paper also explains why a local surjectivity statement from maximal overorders does not simply persist globally.[134]

The useful conclusion here is methodological. A structural class can make a module problem more tractable and can expose stable invariants. One still has to determine whether the particular cryptographic order belongs to that class, which invariants survive a reduction, and what information an algorithm must recover. A theorem about finite representation type is not a theorem asserting polynomial-time recovery of a cryptographic secret.

4.7 A Gram-matrix bridge already present in the book

There is a second direct connection to the finite Hodge calculations below. In the proof of VIII.3.5, Roggenkamp compares a decomposition matrix and a Cartan matrix by means of a positively weighted Gram matrix. On printed page 139, the argument explicitly uses the equality between the rank of a real matrix and the rank of its Gram matrix. This is a concrete positivity argument inside the supplied arithmetic source. It is not the analytic index theorem, but it is the same elementary mechanism that identifies harmonic vectors with simultaneous differential kernels.[133, Theorem VIII.3.5 and Exercise VIII.3.5]

Proposition 4.13 (Weighted Gram kernel and rank)**.** *Let X be an $m \times n$ real matrix, and let $W = \operatorname{diag}(w_1, \dots, w_m)$ with every $w_i > 0$. Then*

$$\ker(X^T W X) = \ker X, \qquad \operatorname{rank}(X^T W X) = \operatorname{rank} X.$$

Moreover $X^T W X$ is positive semidefinite.

Proof. For every column vector $v \in \mathbb{R}^n$,

$$v^T X^T W X v = \sum_{i=1}^{m} w_i (Xv)_i^2 \geq 0.$$

Symmetry is immediate from the diagonal form of W, proving positive semidefiniteness. If $Xv = 0$, then $X^T W X v = 0$. Conversely, if $X^T W X v = 0$, the displayed sum is zero. Each summand is nonnegative and each weight is strictly positive, so every $(Xv)_i$ is zero. This proves the equality of kernels. Both matrices define linear maps with the same n-dimensional domain, so rank–nullity gives equality of ranks. □

The positive weights matter. If W is merely positive semidefinite, nonzero coordinates of Xv in a zero-weight direction can disappear. If W is indefinite, cancellation can occur between positive and negative contributions. For example, $X = (1, 1)^T$ and $W = \mathrm{diag}(1, -1)$ give $X^T W X = 0$ although X has rank one. Similarly, over a finite field a nonzero vector can have zero sum of squares. Thus the real positivity step in the book is not a formal identity that can be transferred unchanged to arbitrary coefficient fields.

For a real two-step complex with chosen inner products, place A^* and B in a single vertically stacked matrix X. Its unweighted Gram matrix is $AA^* + B^*B$, the middle Hodge Laplacian. The proposition then identifies its kernel with $\ker A^* \cap \ker B$. This gives a precise arithmetic-to-Hodge connection at the level of a matrix lemma. Its proof requires no claim that the arithmetic category and the differential-geometric category are equivalent. It also supplies no scale-independent estimate: equality of kernels and ranks does not determine the smallest positive singular value.

4.8 From projective modules to topological bundles

For a compact Hausdorff space X, let $C(X, \mathbb{C})$ be the ring of continuous complex-valued functions. Swan's theorem identifies finite-rank complex vector bundles on X with finitely generated projective modules over $C(X, \mathbb{C})$ through the section functor. It induces

$$K_0(C(X, \mathbb{C})) \cong K^0(X).$$

This is a precise comparison theorem, not an equality between arbitrary algebraic and topological constructions.[162]

Its idempotent mechanism is worth deriving. A finitely generated projective module P is the image of an idempotent matrix $e \in M_n(C(X, \mathbb{C}))$. At each $x \in X$, the matrix $e(x)$ satisfies $e(x)^2 = e(x)$. Its image is a vector subspace of $\mathbb{C}^n$. The rank is locally constant: the eigenvalues of an idempotent are zero and one, and its trace is an integer-valued continuous function. Around a point where an appropriate minor is nonzero, its image admits a continuous frame. These image spaces therefore form a vector bundle. Conversely, a vector bundle on a compact space admits a complementary bundle making it a summand of a trivial bundle; the resulting projection gives an idempotent. This describes how the two kinds of object meet.

For an arithmetic order Λ, one cannot silently replace Λ by $C(X, \mathbb{C})$. One needs an actual representation or scalar-extension construction. For example, a unital homomorphism

$$\rho : \Lambda^{\mathrm{op}} \longrightarrow M_n(C(X, \mathbb{C}))$$

specifies a right Λ-action on the trivial rank-n bundle. Tensoring it with a finitely generated projective left Λ-module then produces a vector bundle, because the latter module is an idempotent summand of a finite free module. This yields a homomorphism on K_0. Its definition depends on ρ; no canonical map is supplied merely by knowing that both domains involve K-theory.

Remark 4.14 (A left–right convention). A homomorphism into endomorphisms gives a left action by default. To use it on the right, one must provide a homomorphism from Λ^{op}, or equivalently specify

a right action directly. Throughout this scalar-extension construction the data are a bundle E with a right Λ-action and a left projective Λ-module P, and the output is $E \otimes_{\Lambda} P$. This convention avoids reversing the multiplication in a noncommutative order.

There is also an important degree-one distinction. Algebraic $K_1(\Lambda)$ is formed from the stable general linear group modulo elementary matrices. Topological K_1 uses a homotopy construction. For the one-point space,

$$K_1^{\mathrm{alg}}(\mathbb{C}) \cong \mathbb{C}^{\times}, \qquad K_1^{\mathrm{top}}(\mathbb{C}) = 0.$$

The first equality follows from the determinant and elementary elimination; the second follows because complex invertible matrices are path connected. The agreement in degree zero under Swan's theorem therefore does not justify identifying all higher algebraic and topological K-groups. Likewise the compact space is recovered as the character or maximal-ideal space of its continuous-function algebra with the appropriate topology; the full algebraic prime spectrum is a different object. These distinctions refine the historical discussion on printed page 114 of the supplied volume.[133, 162]

Chapter 5

Varieties of complexes, exactness, and deformation

5.1 A parameter space with an equation

Fix finite-dimensional vector spaces V_0, V_1, V_2 over a field k, with dimensions a, b, c. A pair of maps

$$V_0 \xrightarrow{A} V_1 \xrightarrow{B} V_2$$

is a complex precisely when $BA = 0$. After choosing bases, its parameter scheme is

$$\mathcal{C}_{a,b,c} = \operatorname{Spec} k[a_{ij}, b_{\ell i}]/\big((BA)_{\ell j} : 1 \le \ell \le c, 1 \le j \le a\big).$$

The equations are quadratic. They express compatibility of successive maps, but they do not assert exactness. Exactness is an additional condition on images and kernels.

Buchsbaum and Eisenbud developed structure and exactness results for finite free complexes. De Concini and Strickland studied the algebraic geometry of varieties of complexes; the later *Hodge algebras* exposition by De Concini, Eisenbud and Procesi gives a direct account in Section 16. It is therefore preferable to identify the precise result being used instead of attributing the entire geometric theory to one name. The calculations below are direct derivations for two successive maps and require neither the full straightening-law machinery nor a classification of all singularities.[27, 40, 41]

Proposition 5.1 (Rank constraint and cohomology)**.** *If $BA = 0$, $r = \operatorname{rank} A$, and $s = \operatorname{rank} B$, then*

$$r + s \le b.$$

The cohomology dimensions of the three-term complex are

$$h^0 = a - r, \qquad h^1 = b - r - s, \qquad h^2 = c - s.$$

In particular the complex is exact at V_1 if and only if $r + s = b$.

Proof. The equation $BA = 0$ implies $\operatorname{im} A \subseteq \ker B$. Rank–nullity gives $\dim \ker B = b - s$, so $r \le b - s$. The formula for h^1 is the dimension of the quotient $\ker B / \operatorname{im} A$. The endpoint formulas are rank–nullity for A and the dimension of the cokernel of B. Equality of the middle image and kernel follows from inclusion and equality of their dimensions. □

Corollary 5.2 (Euler characteristic)**.** *For every pair in $\mathcal{C}_{a,b,c}$,*

$$h^0 - h^1 + h^2 = a - b + c.$$

Proof. Substitution cancels r and s. □

This invariant is constant even when cohomology jumps. It is the numerical version of the Euler identity in the preceding chapter. Its constancy is useful, but it is exactly why it cannot by itself detect a jump or an estimate on small singular values.

5.2 Classification under changes of bases

The group

$$G = \mathrm{GL}(V_0) \times \mathrm{GL}(V_1) \times \mathrm{GL}(V_2)$$

acts by

$$(g_0, g_1, g_2) \cdot (A, B) = (g_1 A g_0^{-1},\ g_2 B g_1^{-1}).$$

This action preserves $BA = 0$ because the composite transforms to $g_2 B A g_0^{-1}$. It represents isomorphism of complexes with the graded spaces fixed up to independent bases.

Theorem 5.3 (Normal form over a field)**.** *Two points of $\mathcal{C}_{a,b,c}$ lie in the same G-orbit if and only if they have the same pair of ranks (r, s).*

Proof. Ranks are unchanged by the action. Conversely, choose a decomposition

$$V_0 = H_0 \oplus U, \qquad V_1 = I \oplus H_1 \oplus W, \qquad V_2 = J \oplus H_2,$$

where $H_0 = \ker A$, $I = \operatorname{im} A$, $\ker B = I \oplus H_1$, and $J = \operatorname{im} B$. Choose U complementary to H_0 and W complementary to $\ker B$. The restrictions $A : U \to I$ and $B : W \to J$ are isomorphisms. Choose bases of U and W, transport them to I and J, and choose arbitrary bases of H_0, H_1, H_2. In these bases the only nonzero blocks of A and B are identity matrices of sizes r and s. The dimensions of all blocks are determined by (a, b, c, r, s), so any two such pairs have the same normal form. □

This theorem depends strongly on the coefficient field. It fails over an order or over $\mathbb{Z}$ if ranks are the only invariants retained. The factor p in the integral example above is invisible over $\mathbb{Q}$ and essential over $\mathbb{Z}$. A visual rank classification over $\mathbb{R}$ or $\mathbb{C}$ must therefore be presented as a model of the linear-algebraic portion of the theory, not as a complete classification of integral lattice complexes.

5.3 Dimension of a rank stratum

Let $\mathcal{C}^{r,s}_{a,b,c}$ be the locally closed set where A has rank exactly r and B has rank exactly s. It is nonempty exactly when

$$0 \le r \le \min(a, b), \quad 0 \le s \le \min(b, c), \quad r + s \le b.$$

The normal form proves sufficiency. The rank inequalities prove necessity.

Theorem 5.4 (Dimension formula)**.** *Over an algebraically closed field, every nonempty rank stratum is smooth and irreducible, and*

$$\dim \mathcal{C}^{r,s}_{a,b,c} = r(a + b - r) + s(b - r + c - s).$$

Proof. First choose $I = \operatorname{im} A \subseteq V_1$, a point of the Grassmannian $\mathrm{Gr}(r, V_1)$ of dimension $r(b-r)$. For this I, the choices of a surjection $V_0 \to I$ form an open subset of $\mathrm{Hom}(V_0, I)$, of dimension ar. Thus the rank-r locus for A has dimension $r(a+b-r)$ and is smooth and irreducible.

For fixed A, the equation $BA = 0$ says precisely that B factors through V_1/I. A rank-s map from this $(b-r)$-dimensional quotient to V_2 is specified by its image $J \in \mathrm{Gr}(s, V_2)$ and a surjection $V_1/I \to J$. The choices have dimension $s(c-s) + s(b-r)$. Locally on Grassmannian charts, the quotient spaces form a trivial vector bundle and these choices form a product with the rank-s matrix locus. This gives smoothness and the displayed dimension. The Grassmannians and the open surjection loci are irreducible, and the bundle description preserves irreducibility. □

The formula is a count of algebraic parameters, not a complexity bound. A search space of large dimension can still be easy to parameterize, while a finite set can encode a hard search problem. For cryptographic use one must also fix the field size or integer height, the input encoding, and which maps or invariants are visible to an adversary.

Example 5.5 (The first nontrivial quadric)**.** Take $(a, b, c) = (1, 2, 1)$. With $A = (x_1, x_2)^T$ and $B = (y_1, y_2)$, the sole equation is

$$q = x_1 y_1 + x_2 y_2 = 0.$$

The $(1,1)$-stratum has dimension $1(1+2-1) + 1(2-1+1-1) = 3$. It is the locus where neither x nor y is zero. The $(1,0)$ and $(0,1)$ strata each have dimension two, and $(0,0)$ is the origin. Every point in the $(1,1)$-stratum is an exact three-term complex.

5.4 First-order deformations and tangent equations

Let $k[\varepsilon]/(\varepsilon^2)$ be the ring of dual numbers. A first-order perturbation of (A, B) has the form

$$A_\varepsilon = A + \varepsilon U, \qquad B_\varepsilon = B + \varepsilon V.$$

Multiplication gives

$$B_\varepsilon A_\varepsilon = BA + \varepsilon(BU + VA).$$

Thus the Zariski tangent space at a point with $BA = 0$ is

$$T_{(A,B)}\mathcal{C}_{a,b,c} = \{(U, V) : BU + VA = 0\}.$$

This calculation is valid over any field. It is the derivative of the polynomial equation defining the scheme, so it applies to the scheme tangent space whether or not the point is smooth.

Proposition 5.6 (Tangent-space dimension)**.** *At a point of rank (r, s),*

$$\dim T_{(A,B)}\mathcal{C}_{a,b,c} = ab + bc - as - cr + rs.$$

Proof. Consider the linear map

$$L : \mathrm{Hom}(V_0, V_1) \oplus \mathrm{Hom}(V_1, V_2) \longrightarrow \mathrm{Hom}(V_0, V_2), \quad L(U, V) = BU + VA.$$

The maps BU are exactly the maps $V_0 \to \operatorname{im} B$. Their space has dimension as: every such map lifts through the surjection $B : V_1 \to \operatorname{im} B$, since vector-space sequences split. The maps VA are exactly the maps $V_0 \to V_2$ that vanish on $\ker A$. Their space has dimension cr, since A identifies $V_0/\ker A$ with $\operatorname{im} A$ and a map on that image extends to V_1.

The intersection consists of maps from $V_0/\ker A$ to $\operatorname{im} B$, so has dimension rs. Therefore $\operatorname{rank} L = as + cr - rs$. The tangent space is $\ker L$, and rank–nullity gives the result. □

Corollary 5.7 (Transverse infinitesimal directions)**.** *The difference between tangent-space dimension and rank-stratum dimension is*

$$(b-r-s)(a+c-r-s) = h^1(h^0+h^2).$$

Proof. Subtract the preceding dimension formulas and factor; substitute the cohomology dimensions already proved. □

This difference measures directions that leave the fixed-rank stratum to first order. Its positivity alone does not prove that the ambient variety is singular: a lower-dimensional stratum can lie inside a smooth ambient locus. For the quadric $q = 0$ with $(a, b, c) = (1, 2, 1)$, a point with $A \neq 0$ and $B = 0$ lies in a two-dimensional stratum but is smooth on the three-dimensional quadric.

Proposition 5.8 (The singularity of the four-coordinate model)**.** *Assume* $\operatorname{char} k \neq 2$*. The hypersurface*

$$x_1y_1 + x_2y_2 = 0 \subseteq \mathbb{A}^4_k$$

is smooth away from the origin and singular at the origin.

Proof. Its gradient is (y_1, y_2, x_1, x_2). This vector is zero exactly at the origin. Away from the origin the hypersurface Jacobian has rank one, yielding a three-dimensional tangent space and smoothness. At the origin the tangent space is the entire four-dimensional ambient space, while the hypersurface has dimension three. Thus the origin is singular. □

The proof in fact extends to this particular polynomial in characteristic two using the same derivatives and the hypersurface criterion. The characteristic restriction is retained here because the orthogonal coordinate change used in the real visualization below is part of the model being illustrated.

5.5 The endomorphism complex and the meaning of gauge

There is a uniform way to package these first-order calculations. For a bounded graded vector space $C^\bullet$, define

$$\operatorname{End}^m(C) = \bigoplus_i \operatorname{Hom}(C^i, C^{i+m}).$$

If d is a differential of degree one, put

$$\delta f = df - (-1)^m fd \quad \text{for } f \in \operatorname{End}^m(C).$$

Composition is the multiplication; the formula is a graded commutator.

Proposition 5.9 (Endomorphism differential)**.** *If* $d^2 = 0$*, then* $\delta^2 = 0$*. For homogeneous* f *of degree* m *and any homogeneous* g*,*

$$\delta(fg) = (\delta f)g + (-1)^m f(\delta g).$$

Proof. For f of degree m,

$$\begin{aligned}\delta^2 f &= d(df - (-1)^m fd) - (-1)^{m+1}(df - (-1)^m fd)d \\ &= d^2 f - (-1)^m dfd + (-1)^m dfd - fd^2 = 0.\end{aligned}$$

For the second identity expand both sides. The two terms containing fdg appear with opposite signs and cancel, leaving $dfg - (-1)^{m+\deg g} fgd$. □

A first-order change $d_\varepsilon = d + \varepsilon a$ satisfies $d_\varepsilon^2 = 0$ if and only if $\delta a = da + ad = 0$. A degree-zero change of basis $g_\varepsilon = 1 + \varepsilon u$ has inverse $1 - \varepsilon u$, and

$$g_\varepsilon d g_\varepsilon^{-1} = d + \varepsilon(ud - du) = d - \varepsilon \delta u.$$

Consequently first-order deformations modulo infinitesimal changes of basis are described by $H^1(\mathrm{End}^\bullet(C), \delta)$.

The word "gauge" in this calculation means a change of bases in a graded algebraic object. In a differential-geometric setting, a gauge transformation is a bundle automorphism and can vary over a manifold. The formulas have a common deformation-theoretic pattern, but the analytic function spaces, global topology and equations of motion are additional data. We will derive the geometric version explicitly rather than treating the terminology as a proof of equivalence.

Proposition 5.10 (Second-order obstruction)**.** *Let $a \in \mathrm{End}^1(C)$ satisfy $\delta a = 0$. A lift*

$$d_t = d + ta + t^2 b \pmod{t^3}$$

is a differential if and only if $\delta b = -a^2$. In particular the class of a^2 in $H^2(\mathrm{End}^\bullet(C), \delta)$ must vanish, and its vanishing is sufficient for such a second-order lift.

Proof. Expanding the square gives

$$d_t^2 = d^2 + t(da + ad) + t^2(db + bd + a^2) \pmod{t^3}.$$

The first two coefficients vanish by hypothesis. The third vanishes precisely when $\delta b = -a^2$. The graded Leibniz rule gives $\delta(a^2) = (\delta a)a - a(\delta a) = 0$, so a^2 is a cocycle. Existence of b is exactly the assertion that its cohomology class is zero. □

Over orders, extension and deformation groups can have torsion, and reduction modulo a prime can change their dimensions. Thus the field calculation should be regarded as one layer in a larger integral deformation problem. It suggests a concrete research workflow: construct an integral complex, compute its cohomology and endomorphism cohomology, compare their reductions, and only then investigate whether the deformation parameters have any relation to a cryptographic task.

5.6 Why ranks alone fail over commutative rings

For a matrix M with entries in a commutative ring S, let $I_t(M)$ be the ideal generated by its $t \times t$ minors, with $I_0(M) = S$. Minors behave well under base change: applying a ring homomorphism to the matrix applies the homomorphism to every minor. The condition $I_t(M) = S$ is much stronger than saying that one $t \times t$ minor is nonzero. It says that locally everywhere some such minor is invertible.

Lemma 5.11 (Splitting off a unit pivot)**.** *Let S be a commutative ring and let a differential in a complex of finite free S-modules have a unit matrix entry. After changes of bases the complex has a direct summand consisting of $S \xrightarrow{1} S$ in the two affected degrees.*

Proof. Permute rows and columns to place the unit in the upper-left entry and scale its row to make it one. Elementary operations clear the rest of its column and row. The differential then has block form $\mathrm{diag}(1, M')$. The identity $d^i d^{i-1} = 0$ forces the component of the preceding map into this pivot coordinate to be zero. Likewise $d^{i+1} d^i = 0$ forces the following map to vanish on its image coordinate. The two coordinates therefore split off as an identity complex. □

This lemma is the ring-theoretic remnant of Gaussian elimination. A nonzero scalar is invertible over a field, so one can keep eliminating. Over a ring, a nonzero scalar such as $p \in \mathbb{Z}$ need not be invertible, and its residual arithmetic becomes part of the homology.

For a noetherian local ring $(S, \mathfrak{m})$ and a proper ideal I, its grade is the maximum length of an S-regular sequence contained in I. A sequence $x_1, \dots, x_t$ is regular if each x_j is a nonzerodivisor on $S/(x_1, \dots, x_{j-1})$ and the final quotient is nonzero. We set $\operatorname{grade} S = +\infty$ for the unit ideal. Over a general noetherian ring the same concept can be expressed by localized depth or Ext; the local formulation is sufficient here.

Theorem 5.12 (Buchsbaum–Eisenbud criterion; external theorem)**.** *Let S be a commutative noetherian local ring and let*

$$0 \to F_n \xrightarrow{d_n} F_{n-1} \to \cdots \xrightarrow{d_2} F_1 \xrightarrow{d_1} F_0$$

be a complex of finite free modules. Put

$$r_i = \sum_{j=i}^{n} (-1)^{j-i} \operatorname{rank} F_j.$$

Using the usual determinantal rank and the convention for the unit ideal, the complex is acyclic in positive degrees if and only if

$$\operatorname{rank} d_i = r_i, \qquad \operatorname{grade} I_{r_i}(d_i) \geq i \quad (1 \leq i \leq n).$$

Here acyclic in positive degrees leaves $\operatorname{coker} d_1$ *unconstrained. Zero free summands can be removed, or handled with the standard zero-rank conventions.[27]*

The full theorem is an external result. The manuscript does not claim that its proof is a few lines of rank–nullity, and the accompanying formalization should not replace it by an unproved axiom while advertising a verified proof. Its two ingredients have different roles: rank conditions control generic compatibility; grade conditions prevent homology from surviving on loci of insufficient codimension or depth.

The integral example $A = (p, 0)^T$, $B = (0, 1)$ illustrates the second ingredient after localizing $\mathbb{Z}$ at (p). The ideal of rank-one minors of A is (p), with grade one. The leftmost map is in homological degree two, so the required grade is at least two. The criterion correctly rejects exactness. Over $\mathbb{Q}$ the same ideal becomes the unit ideal, and the obstruction disappears. The change is arithmetic, not a disagreement between two rank computations.

5.7 A complete exactness proof: the two-generator Koszul complex

Let S be a commutative ring and let $x, y \in S$. Define

$$d_2 : S \to S^2, \quad d_2(c) = (-yc, xc), \qquad d_1 : S^2 \to S, \quad d_1(a, b) = xa + yb.$$

Then $d_1 d_2 = 0$ by commutativity.

Theorem 5.13 (Regular two-generator Koszul complex)**.** *Suppose x is a nonzerodivisor in S and the class of y is a nonzerodivisor in $S/(x)$. Then*

$$0 \to S \xrightarrow{d_2} S^2 \xrightarrow{d_1} S \to S/(x, y) \to 0$$

is exact.

Proof. If $d_2(c) = 0$, then $xc = 0$, hence $c = 0$. Thus d_2 is injective. Its image is contained in $\ker d_1$ by the complex identity. Conversely, suppose $xa + yb = 0$. Modulo (x) this becomes $\overline{y}\,\overline{b} = 0$. The hypothesis gives $b = xc$ for some $c \in S$. Substitution yields $x(a + yc) = 0$, hence $a = -yc$. Therefore $(a, b) = d_2(c)$, proving exactness at S^2. Finally the image of d_1 is exactly the ideal (x, y), which is the kernel of the quotient map to $S/(x, y)$. □

The proof shows precisely where regularity is used. Merely having two nonzero elements is insufficient. If $S = k[t]/(t^2)$ and $x = t$, multiplication by x already has a nonzero kernel; no generic intuition can remove that kernel. If $S = k[x, y]$, the variables form a regular sequence and the preceding resolution is exact. This example connects the local algebra of free resolutions with the symbol calculations of elliptic complexes: both use exterior-algebra differentials, but regularity in a ring and nonzero covectors on a manifold are different hypotheses.

5.8 Base change and the appearance of new homology

Let $S \to T$ be a ring homomorphism. Tensoring a complex with T preserves the equation $d^2 = 0$, because composition and tensor product commute. It need not preserve exactness.

Proposition 5.14 (Flat base change)**.** *If T is flat as an S-module, tensoring an exact complex of S-modules with T preserves exactness. For any complex, the natural comparison identifies its cohomology after tensoring with $H^i(C) \otimes_S T$.*

Proof. Flatness preserves the two short exact sequences involving cycles, boundaries and cohomology used in the Euler identity. Their tensor products identify the new cycles and boundaries with the tensors of the old ones. Taking their quotient gives the cohomology comparison. When the original cohomology vanishes, the new one vanishes as well. □

Localization is flat, so passing from an integral domain to its fraction field preserves exactness. The converse fails because localization can kill torsion. Reduction modulo a prime is usually not flat. For example, tensor the exact sequence

$$0 \to \mathbb{Z} \xrightarrow{p} \mathbb{Z} \to \mathbb{Z}/p\mathbb{Z} \to 0$$

with $\mathbb{F}_p$. The multiplication map becomes zero and loses injectivity. This loss is measured by a Tor term, not by an error in the tensor operation.

For an arithmetic application, one should record at least three coefficient regimes: the order itself, its localizations or completions, and its residue fields. The same written matrix can define complexes with different homology in these regimes. A Lean theorem over a field verifies the field statement; it does not automatically establish the corresponding integral statement. A well-designed theorem interface makes the coefficient assumptions explicit enough that this distinction is enforced by type checking.

5.9 The real four-dimensional quadric as an exact visual model

Return to $q(x, y) = x \cdot y$ with $x, y \in \mathbb{R}^2$. Define

$$u = \frac{x + y}{\sqrt{2}}, \qquad v = \frac{x - y}{\sqrt{2}}.$$

Direct expansion gives

$$q(x,y) = \frac{\|u\|^2 - \|v\|^2}{2}, \qquad \|x\|^2 + \|y\|^2 = \|u\|^2 + \|v\|^2.$$

The variety $q = 0$ is therefore the null cone of a quadratic form of signature $(2, 2)$. Its intersection with a sphere of radius $\rho > 0$ is described by

$$\|u\| = \|v\| = \frac{\rho}{\sqrt{2}}.$$

This link is a product of two circles. An exact parameterization is

$$u = \frac{\rho}{\sqrt{2}}(\cos\alpha, \sin\alpha), \quad v = \frac{\rho}{\sqrt{2}}(\cos\beta, \sin\beta), \quad x = \frac{u+v}{\sqrt{2}}, \quad y = \frac{u-v}{\sqrt{2}}.$$

It supplies a genuine four-coordinate object whose three-dimensional projections can be rotated interactively. The displayed projection can conceal intersections or identify points that are distinct in four dimensions; the underlying coordinates and the residual q should therefore remain visible in the interface.

There are several useful controls with an exact mathematical interpretation. Changing ρ moves along the cone and reaches the singular point at $\rho = 0$. Changing $\alpha - \beta$ changes the relative sizes of x and y; some values make one vanish and move to an endpoint rank stratum. A four-dimensional orthogonal rotation changes only the viewing coordinates if the original coordinates are retained for evaluating q. An arbitrary ambient rotation does not necessarily preserve the fixed expression $x \cdot y$, so the software must distinguish a change of view from a change of the mathematical model.

This example does not illustrate every variety of complexes. It illustrates the smallest nontrivial three-term case in four real coordinates. Its value is that every visible feature has an exact equation, a rank interpretation, and a verifiable relation to cohomology. It can be used to test the visualization code with known cases before any higher-dimensional object is drawn.

Chapter 6

Elliptic complexes, Hodge theory, and gauge deformations

6.1 Ellipticity is a condition on symbols

Let X be a closed smooth manifold: throughout this chapter "closed" means compact and without boundary. Let $E_i \to X$ be smooth complex vector bundles, and let

$$D_i : \Gamma(E_i) \longrightarrow \Gamma(E_{i+1})$$

be first-order differential operators satisfying $D_{i+1}D_i = 0$. The collection is a differential complex. For a covector $\xi \in T^*_x X$, the principal symbol $\sigma(D_i)(x, \xi)$ records the highest-derivative part of D_i in direction ξ. One convention is

$$\sigma(D)(x, df_x)s_x = \sqrt{-1}(D(fs)_x - f(x)D(s)_x).$$

Omitting the factor of $\sqrt{-1}$ gives another common convention; exactness of the symbol sequence is unchanged. The displayed analytic convention makes the Laplacian symbol positive.

Definition 6.1 (Elliptic complex)**.** The differential complex is elliptic if, for every (x, ξ) with $\xi \neq 0$, the finite-dimensional sequence

$$0 \to E_{0,x} \xrightarrow{\sigma(D_0)(x,\xi)} E_{1,x} \to \cdots \to E_{m,x} \to 0$$

is exact.

Thus an elliptic complex produces, at every nonzero cotangent direction, a point in an exact rank locus of a variety of complexes. This is a direct connection to the preceding chapter. The converse needs additional geometric information: a family of exact matrices is not automatically the symbol of differential operators, and even an admissible symbol need not have a chosen operator realization satisfying $D_{i+1}D_i = 0$.

Ellipticity also differs from an elliptic curve. A smooth projective genus-one curve with a chosen point is an algebraic object. An elliptic differential operator is defined by the invertibility or exactness of its principal symbol away from the zero section. Elliptic curves carry useful elliptic complexes, such as their Dolbeault complex, but the two uses of the word do not identify their theories.

6.2 Exterior algebra gives a complete symbol calculation

Let V be a finite-dimensional real or complex vector space. For $\xi \in V^*$ let $\epsilon(\xi)$ be exterior multiplication by ξ on $\bigwedge^\bullet V^*$. For $v \in V$, let ι_v be contraction.

Lemma 6.2 (Wedge–contraction identity)**.** *For every $\alpha \in \bigwedge^\bullet V^*$,*

$$\iota_v(\xi \wedge \alpha) + \xi \wedge \iota_v \alpha = \xi(v)\alpha.$$

Proof. It suffices by linearity to take $\alpha = \eta_1 \wedge \cdots \wedge \eta_p$. The definition of contraction is

$$\iota_v \alpha = \sum_{j=1}^{p} (-1)^{j-1} \eta_j(v) \eta_1 \wedge \cdots \widehat{\eta_j} \cdots \wedge \eta_p.$$

Apply this definition to $\xi \wedge \alpha$. The contribution from the first factor is $\xi(v)\alpha$; every other contribution is the negative of the corresponding term in $\xi \wedge \iota_v \alpha$. Adding the two expressions cancels those remaining terms. □

Theorem 6.3 (Exactness of the nonzero-covector Koszul complex)**.** *If $\xi \neq 0$, the complex*

$$0 \to \bigwedge^0 V^* \xrightarrow{\epsilon(\xi)} \bigwedge^1 V^* \xrightarrow{\epsilon(\xi)} \cdots \xrightarrow{\epsilon(\xi)} \bigwedge^{\dim V} V^* \to 0$$

is contractible and hence exact.

Proof. Choose v with $\xi(v) = 1$. The wedge–contraction identity says that ι_v is a contracting homotopy. More explicitly, if $\xi \wedge \alpha = 0$, then $\alpha = \xi \wedge \iota_v \alpha$, so each cycle is a boundary. The endpoint statements are included in the same identity. □

For the de Rham differential d, the symbol is $\sqrt{-1}\epsilon(\xi)$; the nonzero scalar does not change exactness. The theorem proves ellipticity of the de Rham complex. It is a pointwise proof and does not require a global coordinate system on X. The construction shows exactly how the algebraic condition $d^2 = 0$ and the analytic symbol condition work together: the former makes cohomology meaningful, while the latter supplies local high-frequency control.

With a Riemannian metric, contraction by the metric dual $\xi^\sharp$ is the adjoint of wedge multiplication by ξ. The same identity becomes

$$\epsilon(\xi)\epsilon(\xi)^* + \epsilon(\xi)^*\epsilon(\xi) = \|\xi\|^2 \,\mathrm{id}\,.$$

This is the symbol of the Hodge Laplacian. Its positivity for $\xi \neq 0$ is an ellipticity statement. It is not yet a quantitative lower bound on the nonzero spectrum of the global Laplacian.

6.3 Finite-dimensional Hodge theory with a full proof

Before invoking analytic elliptic theory, one can prove its finite-dimensional model completely. Give a finite cochain complex $C^\bullet$ Hermitian inner products. Let $(d^i)^*$ denote the adjoint, and define

$$\Delta_i = (d^i)^* d^i + d^{i-1}(d^{i-1})^*.$$

A vector is harmonic if it belongs to $\ker \Delta_i$.

Theorem 6.4 (Finite-dimensional Hodge decomposition)**.** *For each degree i,*

$$C^i = \operatorname{im} d^{i-1} \ \perp \ \ker \Delta_i \ \perp \ \operatorname{im}(d^i)^*,$$

and the map from a harmonic vector to its cohomology class is an isomorphism

$$\ker \Delta_i \cong H^i(C^\bullet).$$

Proof. For $v \in C^i$,

$$\langle \Delta_i v, v\rangle = \|d^i v\|^2 + \|(d^{i-1})^* v\|^2.$$

Thus $\ker \Delta_i = \ker d^i \cap \ker(d^{i-1})^*$. The two image spaces are orthogonal because

$$\langle d^{i-1}u, (d^i)^* w\rangle = \langle d^i d^{i-1} u, w\rangle = 0.$$

Their common orthogonal complement is exactly $\ker(d^{i-1})^* \cap \ker d^i$, proving the orthogonal direct-sum decomposition.

If v is a cycle, write it in that decomposition as $v = d^{i-1}u + h + (d^i)^* w$. The last term is itself a cycle because the other two terms are. Taking its inner product with itself gives

$$\|(d^i)^* w\|^2 = \langle w, d^i (d^i)^* w\rangle = 0.$$

Therefore v differs from the harmonic vector h by a boundary. This proves surjectivity onto cohomology. A harmonic boundary has zero inner product with itself because harmonic vectors are orthogonal to boundaries, so it is zero. This proves injectivity. □

There is a spectral consequence with equally clear limits. On the orthogonal complement of its kernel, a fixed finite-dimensional Δ_i has a positive smallest eigenvalue whenever that complement is nonzero. Positivity follows from finite dimensionality and positive definiteness on that subspace. A family of such matrices need not have a uniform positive lower bound.

Example 6.5 (A fixed index and a variable spectral gap)**.** In the real four-coordinate model, suppose $x \cdot y = 0$. The three Hodge Laplacians are

$$\Delta_0 = \|x\|^2, \qquad \Delta_1 = xx^T + yy^T, \qquad \Delta_2 = \|y\|^2.$$

When x and y are both nonzero, they are orthogonal and span $\mathbb{R}^2$. The eigenvalues of Δ_1 are $\|x\|^2$ and $\|y\|^2$, because applying it to x gives $\|x\|^2 x$ and similarly for y. The complex is exact and its smallest eigenvalue on the total graded space is

$$\min\{\|x\|^2, \|y\|^2\}.$$

Taking $x = (t, 0)$ and $y = (0, 1)$ with $t > 0$ preserves exactness and Euler characteristic zero, while the gap tends to zero as $t \to 0$. At $t = 0$ cohomology appears. An index can remain fixed through this entire process.

6.4 The analytic Hodge theorem and its input

For an elliptic complex on a closed manifold, choose Hermitian bundle metrics and a positive volume form. These define formal adjoints D_i^*. The Laplacians

$$\Delta_i = D_i^* D_i + D_{i-1} D_{i-1}^*$$

are elliptic under the first-order elliptic-complex assumptions. Elliptic regularity, compactness of Sobolev inclusions and the spectral theory of elliptic operators supply finite-dimensional smooth harmonic spaces and closed ranges. They then yield the analytic Hodge decomposition and identify harmonic sections with cohomology. These are analytic theorems, not consequences of finite-dimensional linear algebra alone.[10, 12]

Once those analytic inputs are available, the formal proof of the decomposition follows the same orthogonality identities as in the finite-dimensional case. This division is useful for formalization. A Lean development can first prove the finite-dimensional algebraic theorem, then express the analytic theorem through explicit hypotheses about domains, adjoints, closed range and compactness. It must not suppress these hypotheses merely because the displayed identities look identical.

On noncompact manifolds the situation is more delicate. Continuous spectrum can occur, ranges may fail to be closed, and an operator may not be Fredholm. On manifolds with boundary, a boundary condition must be specified. The closed-manifold index statement below does not automatically extend to either case. Since the Yang–Mills mass-gap problem concerns a theory on four-dimensional Euclidean space and its reconstructed quantum dynamics, this distinction is central rather than peripheral.

6.5 The index as an additive invariant

An operator D between appropriate Hilbert spaces is Fredholm if its kernel and cokernel are finite dimensional and its range is closed. Its index is

$$\operatorname{ind} D = \dim \ker D - \dim \operatorname{coker} D.$$

For an elliptic complex, assemble $D + D^*$ from the even-degree to the odd-degree bundles. The analytic Hodge theorem identifies its index with

$$\chi(E_\bullet, D_\bullet) = \sum_i (-1)^i \dim H^i(E_\bullet, D_\bullet).$$

The symbol sequence defines a compactly supported topological K-class on T^*X, because it is exact away from the zero section.

Theorem 6.6 (Atiyah–Singer index theorem; external theorem)**.** *For a complex elliptic differential operator on a closed smooth manifold, the analytic index equals the topological index of its principal-symbol class:*

$$\operatorname{ind}_{\mathrm{an}}(D) = \operatorname{ind}_{\mathrm{top}}([\sigma(D)]), \qquad [\sigma(D)] \in K_c^0(T^*X).$$

The topological index is defined using the Thom construction, an embedding into Euclidean space and Bott periodicity, and is independent of the choices.[12]

The theorem turns analytic finite-dimensional defects into a computable topological invariant. It does not determine all eigenvalues, establish a uniform coercive estimate, or reconstruct an interacting quantum field theory. The class $[\sigma(D)]$ contains symbol information; an arithmetic class in $K_0(\Lambda)$ reaches it only after an explicitly constructed geometric map. This is the exact place where the projective-module discussion from Roggenkamp can enter an index-theoretic construction.

A useful specialization is the Dolbeault complex of a holomorphic vector bundle E on a compact complex manifold X. Its index is the holomorphic Euler characteristic. The cohomological index formula is

$$\chi(X, E) = \int_X \operatorname{ch}(E) \operatorname{Td}(T^{1,0}X).$$

This formula is stated here as a cited external theorem. It should not be confused with a proof that every arithmetic elliptic curve is canonically the same as a chosen complex torus.[13]

Example 6.7 (An elliptic curve over the complex numbers)**.** Let X be a complex elliptic curve. Its holomorphic tangent line bundle is trivial, so its Todd class is one. For a holomorphic line bundle L of degree d, the index formula gives

$$h^0(X, L) - h^1(X, L) = d.$$

For $L = \mathcal{O}_X$, the two dimensions are both one and the index is zero. This is an index-theoretic calculation on a genus-one curve. It says nothing by itself about the difficulty of a discrete logarithm in $E(\mathbb{F}_q)$, which is a finite-group problem with a different input model.

6.6 A complete counterexample to a uniform-gap inference

Consider the circle of length $L > 0$, with coordinate x modulo L. The functions

$$f_n(x) = \exp(2\pi\sqrt{-1}nx/L), \qquad n \in \mathbb{Z},$$

are eigenfunctions of $-d^2/dx^2$ with eigenvalues

$$\lambda_n = 4\pi^2 n^2/L^2.$$

The smallest positive eigenvalue is $4\pi^2/L^2$. Both the de Rham complex and its index are independent of L up to the relevant topological equivalence: the Euler characteristic of the circle is zero. Yet its positive spectral gap approaches zero as $L \to \infty$.

To verify that no smaller positive eigenvalue is missing, let $-f'' = \lambda f$ with periodic values and derivatives. Integration by parts gives $\lambda\|f\|^2 = \|f'\|^2$, so $\lambda \geq 0$. For $\lambda > 0$, the elementary ODE solution is a linear combination of $\cos(\sqrt{\lambda}x)$ and $\sin(\sqrt{\lambda}x)$. A nonzero periodic solution exists exactly when rotation through $\sqrt{\lambda}L$ fixes a nonzero vector, equivalently $\sqrt{\lambda}L \in 2\pi\mathbb{Z}$. This proves the stated least positive eigenvalue. The cohomology has dimension one in degrees zero and one: closed zero-forms are constant, and a one-form is exact precisely when its integral around the circle vanishes. The example gives a rigorous reason that an index theorem cannot alone supply the kind of scale-uniform lower bound needed in a continuum mass-gap argument. The missing ingredient is quantitative control under a limiting process, not another application of Euler-characteristic cancellation.

6.7 Connections and curvature: when a differential squares to zero

Let $P \to X$ be a principal bundle for a compact Lie group G, and let $\operatorname{ad} P$ be its adjoint bundle. A connection A defines the covariant exterior derivative

$$d_A : \Omega^i(X, \operatorname{ad} P) \longrightarrow \Omega^{i+1}(X, \operatorname{ad} P).$$

In a local trivialization one can write $d_A\alpha = d\alpha + [A \wedge \alpha]$. Its curvature is $F_A = dA + \frac{1}{2}[A \wedge A]$.

Proposition 6.8 (Curvature identity)**.** *For an adjoint-valued differential form α,*

$$d_A^2\alpha = [F_A \wedge \alpha].$$

Proof. Expand $d_A(d_A\alpha)$ using the graded Leibniz rule for exterior differentiation and the graded Lie bracket. The terms containing one ordinary derivative of α cancel. The remaining terms are $[dA \wedge \alpha] + \frac{1}{2}[[A \wedge A] \wedge \alpha]$, where the graded Jacobi identity supplies the factor $\frac{1}{2}$. This is precisely $[F_A \wedge \alpha]$. □

If A is flat, meaning $F_A = 0$, the covariant de Rham sequence is a complex. Its principal symbol is ordinary wedge multiplication by the covector tensored with the identity on $\operatorname{ad} P$. The previously proved Koszul exactness therefore makes it elliptic. Flatness is a sufficient condition for the complex identity; in the adjoint representation a central curvature can also act trivially. It would be inaccurate to say that $d_A^2 = 0$ always forces the entire curvature to vanish for an arbitrary structure group.

For a general Yang–Mills connection, the equation is $d_A^* F_A = 0$. This does not imply $F_A = 0$. Thus the full covariant de Rham sequence is not automatically a complex at a general Yang–Mills connection. Writing down d_A in successive degrees and calling it an elliptic complex without checking this composition is a mathematical error.

6.8 The anti-self-dual deformation complex

Assume now that X is an oriented Riemannian four-manifold. On two-forms the Hodge star satisfies $*^2 = 1$, so

$$\Omega^2 = \Omega^{2,+} \oplus \Omega^{2,-}, \qquad P_+ = \tfrac{1}{2}(1 + *).$$

A connection is anti-self-dual (ASD) when $F_A^+ = 0$. Set $d_A^+ = P_+ d_A$ on one-forms. At an ASD connection there is a complex

$$0 \longrightarrow \Omega^0(\operatorname{ad} P) \xrightarrow{d_A} \Omega^1(\operatorname{ad} P) \xrightarrow{d_A^+} \Omega^{2,+}(\operatorname{ad} P) \longrightarrow 0.$$

Proposition 6.9 (Complex identity at an ASD connection)**.** *If $F_A^+ = 0$, then $d_A^+ d_A = 0$ on adjoint-valued zero-forms.*

Proof. The curvature identity gives $d_A^+ d_A \phi = P_+[F_A, \phi]$. The Hodge projection acts only on the form component and commutes with the pointwise Lie bracket with the zero-form ϕ. Hence $P_+[F_A, \phi] = [F_A^+, \phi] = 0$. □

Proposition 6.10 (Ellipticity of the ASD complex)**.** *The ASD deformation complex is elliptic.*

Proof. It suffices to check a fiber and a nonzero covector ξ. Rescale and choose an oriented orthonormal basis with $\xi = e^1$. Suppress the adjoint factor, which is unchanged by the symbol maps. The first symbol sends a scalar to its multiple of e^1 and is injective. The second sends

$$\alpha = a_1 e^1 + a_2 e^2 + a_3 e^3 + a_4 e^4$$

to

$$\frac{1}{2}\left[a_2(e^{12} + e^{34}) + a_3(e^{13} - e^{24}) + a_4(e^{14} + e^{23})\right].$$

The three displayed two-forms are a basis of the self-dual two-forms. Thus the second symbol is surjective and its kernel is the span of e^1, exactly the image of the first. Tensoring with the adjoint fiber preserves exactness. □

These are complete proofs of the algebraic identities underlying the deformation complex. The subsequent global study of its moduli space requires elliptic analysis and nonlinear geometry. Atiyah, Hitchin and Singer's work supplies the classical four-dimensional framework and its index calculations.[16]

To see why this complex describes deformations, perturb the connection by a one-form a. The curvature expansion is

$$F^+_{A+a} = F^+_A + d^+_A a + \tfrac{1}{2}[a \wedge a]^+.$$

At an ASD background the linearized equation is $d^+_A a = 0$. An infinitesimal gauge transformation generated by a zero-form ϕ changes a by $d_A\phi$, up to the conventional overall sign. Thus infinitesimal deformations modulo infinitesimal gauge are

$$H^1_A = \ker d^+_A / \operatorname{im} d_A.$$

The group H^0_A measures infinitesimal stabilizers, and H^2_A is the obstruction space appearing after solving the nonlinear equation locally. The names describe their roles; the statement that a local moduli space has a given smooth dimension requires the stabilizer and obstruction hypotheses below.

6.9 Expected dimension and actual local dimension

Introduce the gauge-fixed operator

$$\mathcal{D}_A = d^*_A \oplus d^+_A : \Omega^1(\operatorname{ad} P) \to \Omega^0(\operatorname{ad} P) \oplus \Omega^{2,+}(\operatorname{ad} P).$$

It is elliptic. On a closed manifold its kernel is the harmonic representative space of H^1_A. Its cokernel is $H^0_A \oplus H^2_A$. The latter assertion follows from the adjoint formula and orthogonality: the images of d_A and $(d^+_A)^*$ are orthogonal because $d^+_A d_A = 0$. Therefore

$$\operatorname{ind} \mathcal{D}_A = \dim H^1_A - \dim H^0_A - \dim H^2_A.$$

For an $SU(2)$ bundle on a connected X with second Chern number k, the standard index calculation is

$$\operatorname{ind} \mathcal{D}_A = 8k - 3(1 - b_1(X) + b^+_2(X)).$$

This formula is cited as an external index-theoretic result. If the connection is irreducible and regular, so that the relevant infinitesimal stabilizer and obstruction spaces vanish, elliptic slice and implicit-function theorems identify the local moduli dimension with this index. Without those hypotheses, it is an expected dimension, not an unconditional dimension of a smooth manifold.[16]

This is a particularly clear instance of the common pattern across the manuscript. A differential identity gives a complex. Its cohomology controls an infinitesimal classification problem. An index evaluates an alternating sum. Nonlinear existence, global classification and quantitative estimates remain additional questions. Replacing the last three questions by the alternating sum would discard information rather than solve them.

6.10 Classical action bounds do not establish a quantum mass gap

Choose an invariant positive inner product on the compact Lie algebra. In four dimensions decompose the curvature as $F_A = F^+_A + F^-_A$. Orthogonality gives

$$\|F_A\|^2_{L^2} = \|F^+_A\|^2_{L^2} + \|F^-_A\|^2_{L^2}.$$

The quadratic topological expression obtained by pairing F_A with $*F_A$ is

$$Q(A) = \|F_A^+\|_{L^2}^2 - \|F_A^-\|_{L^2}^2.$$

Hence

$$\|F_A\|_{L^2}^2 \geq |Q(A)|,$$

with equality when one of the two curvature components vanishes. For the usual trace normalization on $SU(2)$, $|Q(A)| = 8\pi^2|k|$. The inequality itself follows only from nonnegativity of the two squared norms; the integrality and topological interpretation of k use Chern–Weil theory.

An ASD connection also solves the classical Yang–Mills equation: the Bianchi identity $d_A F_A = 0$ and $*F_A = -F_A$ imply $d_A * F_A = 0$, which is equivalent to the Yang–Mills equation. These are classical statements about a connection and its action. A quantum mass gap concerns a positive separation above the vacuum in the spectrum of a rigorously constructed quantum theory. A lower bound on classical action in a nontrivial topological sector does not construct that theory, identify its vacuum Hilbert space, or control its spectrum. The action bound is valuable but proves a different statement.

6.11 A lower-dimensional bridge through holomorphic bundles

Atiyah and Bott studied the Yang–Mills functional on a compact Riemann surface using gauge symmetry and Morse theory, connecting its critical behavior with moduli of holomorphic vector bundles. This gives a strong, established link between complex algebraic geometry, gauge theory and topology. It is a result in two real dimensions with its specified geometric setting; it is not a construction of four-dimensional quantum Yang–Mills theory.[11]

On a complex curve, a unitary connection determines a $(0,1)$ operator. Since $(0,2)$-forms vanish in complex dimension one, its integrability condition is automatic, so it defines a holomorphic structure. In higher complex dimensions the curvature component $F_A^{0,2}$ need not vanish, and the same step needs an additional condition. This elementary dimensional observation already explains why a surface result cannot simply be reused in four real dimensions.

For an elliptic curve viewed as a complex Riemann surface, the bundle and Dolbeault constructions fit this lower-dimensional bridge. For an elliptic curve over a finite field, Frobenius, endomorphism rings and finite groups supply arithmetic structures that are not encoded by this analytic surface argument. A comparison project can study both sides, but its claims should specify whether a curve, bundle, module, complex, cohomology group or algorithm is being transported. That discipline is what turns a collection of suggestive themes into a mathematically checkable program.

6.12 Interfaces suitable for formal verification

The preceding chapters identify a finite list of statements with a manageable algebraic core. The chain condition $BA = 0$ implies image containment in a kernel. Rank–nullity converts that containment into exactness when dimensions add correctly. The derivative of the chain condition is $BU + VA = 0$. A graded commutator squares to zero when the underlying differential does. The Euler identity follows from two exact sequences and cancellation. The wedge–contraction relation proves exactness of a symbol sequence. Each assertion can be represented without postulating the Yang–Mills conjecture or the full Atiyah–Singer theorem.

An honest formalization separates these checked statements from analytic imports. A proof of the finite-dimensional Hodge decomposition does not establish elliptic regularity. A polynomial

proof of the ASD symbol calculation does not establish the nonlinear slice theorem. A theorem conditional on a Fredholm hypothesis does not prove that a particular unbounded operator has a closed range. A definition of a topological index is not a formal proof of its equality with the analytic index. The accompanying claim manifest should retain these distinctions so that the word "verified" refers to an actual kernel-checked declaration and its visible assumptions.

For future work, the natural formalization layers are: finitely supported complexes and polynomial identities; vector-space dimensions and exactness; finite-dimensional adjoints and positivity; projective-module Euler classes; Sobolev and Fredholm interfaces; and finally geometric symbol classes and index comparisons. This order makes dependencies explicit. It also allows a genuinely useful verified core to be completed while leaving the deep analytic theorems marked as external mathematical inputs.

Chapter 7

Exterior algebra, tableaux, and Schur constructions

7.1 The algebra needed to describe a space of complexes

A matrix equation such as $BA = 0$ has several simultaneous meanings. It states that two linear maps compose to zero; it gives polynomial equations in their entries; and it makes every exterior power of their composition vanish. The first interpretation controls homology. The second defines an algebraic scheme. The third produces the determinant identities used to describe its coordinate ring. Standard monomial theory makes the third interpretation effective: it chooses products of minors that form a basis, and supplies relations that express all other products in that basis.

This requires more than remembering that determinants are alternating. One must keep track of integral coefficients, the direction of each linear map, the dual action on coordinate functions, and the distinction between a basis of a module and multiplication in an algebra. These chapters develop that bookkeeping explicitly. The elementary arguments are proved here. The general standard basis and normality theorems are identified as external results, with their hypotheses and their precise role stated.

Throughout this chapter R is a commutative ring with identity and V is a finite free R-module. Statements involving irreducible representations will instead specify an algebraically closed field of characteristic zero. Maintaining this distinction matters for orders: an argument using division by a hook length cannot silently be transported to an order over a residue field whose characteristic divides that hook length.

7.2 Exterior powers and determinant coordinates

Definition 7.1. The exterior algebra $\bigwedge V$ is the quotient of the tensor algebra of V by the two-sided ideal generated by $v \otimes v$, for all $v \in V$. Its degree p part is $\bigwedge^p V$. We use $\bigwedge^0 V = R$. For an ordered basis $e_1, \dots, e_n$, write

$$e_I = e_{i_1} \wedge \cdots \wedge e_{i_p}, \qquad I = \{i_1 < \cdots < i_p\} \subseteq \{1, \dots, n\}.$$

The relation $v \wedge v = 0$ is part of the definition. In characteristic two it does not follow merely from the formula $v \wedge w = -w \wedge v$. Expanding $(v + w) \wedge (v + w) = 0$ does, however, prove the latter formula over every R.

Proposition 7.2 (The exterior basis)**.** *The elements e_I, with $|I| = p$, form an R-basis of $\bigwedge^p V$. In particular its rank is $\binom{n}{p}$, interpreted as zero when $p > n$.*

Proof. Multilinearity expands any pure wedge in the basis vectors. A term with a repeated index vanishes; adjacent interchanges put the remaining indices in increasing order, so the displayed elements span. For every ordered p-element set J, define an alternating p-linear form on V^p by taking the determinant of the coordinates in the rows indexed by J. The universal property of the exterior quotient gives a functional $\ell_J : \bigwedge^p V \to R$. Its value on e_I is 1 for $I = J$, and zero otherwise. Applying these functionals to a proposed linear relation proves independence, without dividing by any integer. □

A linear map $T : V \to W$ therefore gives a linear map $\bigwedge^p T$ by applying T to each factor. If T has matrix M, then the coefficient of f_I in $(\bigwedge^p T)e_J$ is the minor $\det M_{I,J}$. This statement is a direct expansion of the determinant, so it works over every commutative coefficient ring.

Proposition 7.3 (Cauchy–Binet with specified indices)**.** *Let $A : R^a \to R^b$ and $B : R^b \to R^c$. If $I \subseteq [c]$ and $J \subseteq [a]$ have size p, then*

$$\det(BA)_{I,J} = \sum_{\substack{K \subseteq [b] \\ |K|=p}} \det B_{I,K} \det A_{K,J}.$$

Proof. Functoriality on pure wedges gives $\bigwedge^p(BA) = (\bigwedge^p B)(\bigwedge^p A)$. Taking its I, J matrix entry in the exterior bases gives the formula. This also proves it when $p > b$, since the sum is empty and the intermediate exterior module is zero. □

If $BA = 0$ and $p > 0$, these sums vanish. An individual summand need not vanish. Already for a row $B = (b_1, b_2)$ and column $A = (a_1, a_2)^{\mathsf{T}}$, the equation is $b_1a_1 + b_2a_2 = 0$, not the pair of equations $b_1a_1 = b_2a_2 = 0$. Confusing a determinant identity with termwise vanishing destroys the coordinate ring. The mixed straightening relations below retain the necessary sums and signs.

7.3 Exterior coproducts and shuffle signs

For disjoint increasing index sets I, J whose union is $\{1, \dots, p+q\}$, let $\epsilon(I, J)$ be the sign of the permutation that first lists I, then J.

Definition 7.4. The exterior coproduct of bidegree (p, q) is

$$\Delta_{p,q}(v_1 \wedge \cdots \wedge v_{p+q}) = \sum_{\substack{I \sqcup J = [p+q] \\ |I|=p}} \epsilon(I, J)\, v_I \otimes v_J.$$

Here the elements in each wedge are taken in their original increasing index order.

Proposition 7.5 (Shuffle identities over a ring)**.** *The map $\Delta_{p,q}$ is well defined and natural in V. Its iterates are coassociative in the graded sense. If $m_{p,q}(u \otimes v) = u \wedge v$, then*

$$m_{p,q}\Delta_{p,q} = \binom{p+q}{p} \mathrm{id}.$$

Proof. The expression is multilinear. Interchanging two adjacent input vectors pairs the summands whose subsets exchange those indices, or changes the sign inside the corresponding exterior factor. Thus it changes the overall sign. If the two inputs agree, the paired summands cancel and the other terms contain a repeated wedge. Consequently the expression is alternating and descends to the

exterior quotient. Applying a linear map to every input commutes with every summand, proving naturality.

For coassociativity, expand both ways of splitting a wedge into three factors of sizes p, q, r. Both expansions are indexed by ordered partitions of $[p+q+r]$ into three subsets of those sizes. In each term, the product of the two shuffle signs is the sign of the permutation that lists the three subsets in order. The terms therefore agree. Finally, each summand of $m\Delta$ acquires a second copy of its shuffle sign when its factors are reordered. It becomes the original wedge, and there are $\binom{p+q}{p}$ summands. □

The final identity is an integral identity, not an instruction to divide by the binomial coefficient. For example, dividing $\Delta_{1,1}$ by two is unavailable over $\mathbb{F}_2$. Integral Schur constructions use shuffle maps themselves rather than normalized averaging projectors.

The closely related contraction operation will connect this algebra to complexes. For $\ell \in V^* = \mathrm{Hom}_R(V, R)$, put

$$\iota_\ell(v_1 \wedge \cdots \wedge v_p) = \sum_{j=1}^{p} (-1)^{j-1} \ell(v_j) v_1 \wedge \cdots \widehat{v_j} \cdots \wedge v_p.$$

Lemma 7.6 (Contraction and a contracting homotopy)**.** *One has $\iota_\ell^2 = 0$. If $\ell(u) = 1$, the operation $h(\omega) = u \wedge \omega$ satisfies $\iota_\ell h + h\iota_\ell = \mathrm{id}$. Thus the contraction complex is exact.*

Proof. In ι_ℓ^2, the term deleting positions $i < j$ appears twice. Deleting i first changes the later position of j by one, so the two signs are opposite. They cancel over R. Expanding the definition on $u \wedge \omega$ gives $\iota_\ell(u \wedge \omega) = \ell(u)\omega - u \wedge \iota_\ell \omega$. This is the asserted homotopy identity. A cycle ω therefore equals $\iota_\ell h \omega$, so is a boundary. □

This is the finite algebraic mechanism behind several Koszul exactness arguments. In the de Rham symbol complex at a nonzero covector ξ, choose a vector u with $\xi(u) = 1$ and exchange the roles of wedge and contraction to obtain the same homotopy. The construction proves exactness of that symbol sequence. It contains no estimate uniform in a degenerating family of covectors or connections; such an estimate is a separate analytic issue.

7.4 Partitions and a characteristic-free Schur presentation

A partition $\lambda = (\lambda_1 \geq \cdots \geq \lambda_s > 0)$ is drawn with λ_i boxes in row i. Its transpose λ' records column heights. Write these heights as $c_1 \geq \cdots \geq c_{\lambda_1} > 0$, and set

$$E_\lambda(V) = \bigwedge^{c_1} V \otimes \cdots \otimes \bigwedge^{c_{\lambda_1}} V.$$

A tensor in the exterior basis gives a filling of the diagram whose columns are strictly increasing. A filling is *semistandard* here when its columns are strictly increasing and its rows are weakly increasing. Some sources transpose this convention and call the resulting tableaux standard; the convention is fixed here to prevent a hidden transpose in the Schur index.

For adjacent column heights $p \geq q$ and $1 \leq t \leq q$, consider

$$\theta_{p,q,t} : \bigwedge^{p+t} V \otimes \bigwedge^{q-t} V \longrightarrow \bigwedge^{p} V \otimes \bigwedge^{q} V,$$
$$\theta_{p,q,t} = (\mathrm{id} \otimes m_{t,q-t})(\Delta_{p,t} \otimes \mathrm{id}).$$

Tensor this map with the identity on the other columns.

Definition 7.7. The Schur module $S_\lambda(V)$, with the row-partition convention used here, is the quotient of $E_\lambda(V)$ by the sum of the images of these adjacent-column shuffle maps.

The following statement is the general basis theorem used in this chapter. The integral shuffle presentation and its standard basis are proved in Akin–Buchsbaum–Weyman, Section II.2, with transposed diagram conventions [1].

Theorem 7.8 (External general Schur basis theorem)**.** *For finite free V with ordered basis, the classes of semistandard tableaux of shape λ form an R-basis of $S_\lambda(V)$. This construction commutes with extension of the coefficient ring. In particular it is zero if the first column is longer than* $\operatorname{rank} V$.

There are two logically distinct parts to this theorem. Shuffle relations move a violating entry toward an earlier column and produce tableaux lower in a finite tableau order, which proves spanning. Independence uses a triangular comparison with monomials in a tensor product of symmetric powers; the comparison matrix has diagonal entries one. The latter fact is what makes the statement valid over $\mathbb{Z}$, rather than only over a field where a nonzero determinant could be inverted. The complete general triangularity argument is the cited result; the following cases give full proofs without invoking it.

Proposition 7.9 (Three elementary Schur cases)**.** *For finite free V,*

$$S_{(d)}(V) = \operatorname{Sym}^d(V), \qquad S_{(1^d)}(V) = \bigwedge^d V.$$

Moreover the sequence

$$0 \longrightarrow \bigwedge^3 V \xrightarrow{\Delta_{2,1}} \bigwedge^2 V \otimes V \longrightarrow S_{(2,1)}(V) \longrightarrow 0$$

is exact. If $\operatorname{rank} V = n$*, the last module has rank $n(n^2-1)/3$.*

Proof. For a one-row diagram, all exterior factors have degree one. The relations from adjacent columns are $v \otimes w - w \otimes v$, so the quotient is the symmetric power. For a one-column diagram there are no adjacent-column relations, and the module is the indicated exterior power.

For $(2,1)$, the defining relations are precisely the image of $\Delta_{2,1}$. If $i < j < k$, that image is

$$(e_i \wedge e_j) \otimes e_k - (e_i \wedge e_k) \otimes e_j + (e_j \wedge e_k) \otimes e_i.$$

Among these three basis tensors only the last fails the semistandard condition: the top entry j of the first column is larger than the second-column entry i. Every nonstandard basis tensor has this unique form for one triple $i < j < k$. The relation therefore eliminates that tensor with coefficient one, and no relation for another triple contains it. The relations are independent and the remaining standard tensors form a free basis of the quotient. This proves both injectivity on the left and the asserted basis directly. Counting gives $n\binom{n}{2} - \binom{n}{3} = n(n^2-1)/3$. □

For $n = 3$, this computation replaces

$$(e_2 \wedge e_3) \otimes e_1 = -(e_1 \wedge e_2) \otimes e_3 + (e_1 \wedge e_3) \otimes e_2 \quad \text{in } S_{(2,1)}(V).$$

The remaining eight tensors are a basis. This is a concrete straightening law, including its signs and its independence proof. It also shows why projecting onto a preferred representation by division by three is unnecessary.

An even smaller exact sequence is

$$0 \longrightarrow \bigwedge^2 V \xrightarrow{v\wedge w\mapsto v\otimes w-w\otimes v} V\otimes V \longrightarrow \operatorname{Sym}^2 V \longrightarrow 0.$$

Injectivity follows because each $e_i \wedge e_j$, $i < j$, has a distinct coefficient-one tensor $e_i \otimes e_j$. The quotient makes all tensor factors commute and has the usual ordered monomial basis. This proof remains valid in characteristic two, although splitting the sequence by symmetrization requires extra hypotheses.

7.5 Rational highest weights and duals

Let now k be algebraically closed of characteristic zero and $\dim V = n$. Polynomial Schur modules indexed by partitions of length at most n are irreducible rational representations of $\mathrm{GL}(V)$ [1]. To describe coordinate rings of complexes one also needs negative weights. For a dominant integer sequence $\eta_1 \geq \cdots \geq \eta_n$, define

$$S_\eta V = (\det V)^{\eta_n} \otimes S_{(\eta_1-\eta_n,\ldots,\eta_{n-1}-\eta_n,0)}V.$$

A negative tensor power of $\det V$ means the corresponding power of its one-dimensional dual.

If λ, μ are partitions with $\ell(\lambda) + \ell(\mu) \leq n$, the notation

$$S_{(\lambda,-\mu)}V = S_{(\lambda_1,\ldots,\lambda_a,0,\ldots,0,-\mu_b,\ldots,-\mu_1)}V$$

includes reversal of the negative block. Without that reversal the sequence generally would not be dominant. In particular, $(S_\lambda V)^*$ has highest weight $(0,\ldots,0,-\lambda_a,\ldots,-\lambda_1)$. The determinant twists are required even when all matrices in the original problem have ordinary polynomial entries.

For example, if $\dim V = 2$,

$$S_{(q,-p)}V = (\det V)^{-p} \otimes \operatorname{Sym}^{p+q} V, \qquad \dim S_{(q,-p)}V = p+q+1.$$

The dimension can be read from the monomial basis of a binary form. It will reappear as the dimension of a bihomogeneous piece of the simplest nontrivial complex variety.

Over a field of positive characteristic, the integral Schur module and its tableau basis still make sense. Its irreducibility and the splitting of representation filtrations are separate assertions and generally fail. Likewise, a characteristic-zero Young symmetrizer can be normalized to an idempotent only after its scalar hook factor becomes invertible. The integral presentation above avoids making this normalization an unspoken assumption.

7.6 Schur modules applied to complexes

A polynomial construction on modules is not automatically an exact functor. For the split short exact sequence $0 \to k \to k^2 \to k \to 0$, applying Sym^2 term by term gives dimensions $1, 3, 1$. Its middle homology contains the mixed monomial and has dimension one. Thus even splitting the original sequence does not make this termwise sequence exact. A Schur complex is a construction on the entire graded complex with its differential, not merely a list of Schur modules of its terms.

Here is a precise construction sufficient in characteristic zero. If $C^\bullet$ is a bounded complex of vector spaces, the differential on its d-fold tensor power is

$$d(c_1 \otimes \cdots \otimes c_d) = \sum_{j=1}^{d} (-1)^{|c_1|+\cdots+|c_{j-1}|} c_1 \otimes \cdots \otimes dc_j \otimes \cdots \otimes c_d.$$

The symmetric group permutes the homogeneous factors with the corresponding graded signs.

Proposition 7.10 (A characteristic-zero Schur complex)**.** *Let d be a positive integer. The displayed differential squares to zero and commutes with the graded permutation action. The image of any Young idempotent for a partition of d is a subcomplex. If $C^\bullet$ is acyclic, this image is acyclic.*

Proof. Terms in the square that differentiate the same factor vanish because the original differential squares to zero. Terms that differentiate distinct factors occur twice, with signs differing by one degree shift; they cancel. A transposition of adjacent homogeneous factors commutes with the differential by the same sign calculation. Since adjacent transpositions generate the symmetric group, every group algebra element commutes with the differential, including the idempotent. Its image and kernel are therefore complementary subcomplexes.

A bounded acyclic complex of vector spaces is contractible: choose a complement of the cycles in each degree and invert the differential from that complement onto the boundaries in the next degree. If h is the resulting contracting homotopy, $h \otimes \mathrm{id}$ contracts the tensor product of $C^\bullet$ with any other complex. The cross terms from the other differential cancel by the graded signs. Thus $(C^\bullet)^{\otimes d}$ is acyclic for $d > 0$. A direct summand of an acyclic complex has zero cohomology, because cohomology respects a direct sum decomposition. This proves the assertion for the idempotent image. □

The case $d = 0$ is deliberately excluded from the final argument: the zeroth tensor power is k, regardless of the input complex. Over a general coefficient ring, divided powers and the integral Schur-complex construction replace this idempotent argument [1]. In either setting, an assertion about exactness must identify the construction actually being used. Neither exterior powers nor Schur terminology alone supply a new exactness theorem for orders, a quantum algorithm, or an analytic spectral gap.

Chapter 8

Standard monomials on spaces of complexes

8.1 What a standard monomial theorem must establish

A useful straightening theorem does three jobs. It identifies a set of proposed basis monomials, proves that rewriting reaches their span, and proves that no nontrivial linear relation remains among them. Termination without independence is not a basis theorem. Independence without spanning is not a presentation of the coordinate ring. The examples in the next two sections separate these tasks and prove each one.

All algebras in these sections are commutative. The expression *standard monomial* depends on an explicitly chosen order; it is not an intrinsic property of an arbitrary product of determinants. The geometry is independent of that order, while the normal forms and the corresponding degeneration can depend strongly on it.

8.2 Rank-one matrices: the Segre example

Let $Z = (z_{ij})$ be an a-by-b matrix of independent variables over a field k. Denote by $I_2(Z)$ the ideal of its two-by-two minors. Consider the homomorphism

$$\varphi : k[z_{ij}] \longrightarrow k[u_1, \dots, u_a, v_1, \dots, v_b], \qquad z_{ij} \longmapsto u_i v_j.$$

Every two-by-two minor belongs to its kernel. The following proof shows that these are all the equations, and at the same time constructs the coordinate basis.

Theorem 8.1 (Rank-one straightening, with full proof)**.** *The kernel of φ is $I_2(Z)$. A basis of $k[Z]/I_2(Z)$ consists of* 1 *and the products*

$$z_{i_1 j_1} \cdots z_{i_d j_d}, \qquad i_1 \le \cdots \le i_d, \quad j_1 \le \cdots \le j_d.$$

Its degree d dimension is

$$\binom{a+d-1}{d}\binom{b+d-1}{d}.$$

Proof. Use commutativity to put the row indices of any product in nondecreasing order. If its column sequence has an adjacent inversion $j > l$, the two relevant factors are $z_{ij} z_{kl}$ with $i \le k$.

When $i = k$, commuting the two factors sorts their column indices without changing the row sequence. When $i < k$, the minor relation gives

$$z_{ij}z_{kl} = z_{il}z_{kj}.$$

The row sequence remains sorted, while the column inversion count decreases by one. Repetition terminates and proves spanning by the displayed monomials.

Their images under φ are $u_{i_1} \cdots u_{i_d} v_{j_1} \cdots v_{j_d}$. Distinct standard products have distinct sorted row or column multisets, hence distinct monomial images. Monomials in a polynomial ring are independent, so the proposed basis is independent. If a polynomial lies in $\ker \varphi$, reduce it modulo the minor relations to a standard linear combination. Independence forces all coefficients to be zero; this proves the kernel assertion. Finally, the two independently chosen weakly increasing sequences are counted by the two displayed binomial coefficients. □

The proof works without a division operation and therefore also gives a free basis over a commutative coefficient ring. Over a field it exhibits the coordinate ring as a subalgebra of a polynomial ring, so it is a domain. On the open set $z_{11} \neq 0$, all other entries satisfy $z_{ij} = z_{i1}z_{1j}/z_{11}$. The coordinates

$$z_{11}^{\pm 1}, \quad z_{21}, \ldots, z_{a1}, \quad z_{12}, \ldots, z_{1b}$$

are independent: arbitrary values with $z_{11} \neq 0$ produce a rank-one matrix by the displayed formula. The affine dimension is $a + b - 1$. Its projectivization is the Segre variety $\mathbb{P}^{a-1} \times \mathbb{P}^{b-1}$.

For later use, the exponent semigroup of the image is

$$S_{\mathrm{Segre}} = \{(\alpha, \beta) \in \mathbb{N}^a \times \mathbb{N}^b : |\alpha| = |\beta|\}.$$

The straightening proof explains why every such pair occurs: list the row and column indices with those multiplicities and pair them in sorted order. Thus its semigroup description is an exact description of the original ring, not merely a limiting approximation.

8.3 Two-row tableaux and the Plucker algebra

Let $x_1, \ldots, x_n, y_1, \ldots, y_n$ be independent variables, and define

$$p_{ij} = x_i y_j - x_j y_i, \qquad 1 \leq i < j \leq n.$$

These are the two-by-two minors of a two-by-n matrix. Write P_{ij} for formal generators, to distinguish a polynomial presentation from its image in $k[x, y]$.

Lemma 8.2 (The elementary Plucker shuffle)**.** *For $i < k < l < j$,*

$$p_{ij}p_{kl} = p_{il}p_{kj} - p_{ik}p_{lj}. \tag{P}$$

Proof. Expand the right side using the definitions. Terms containing $x_i x_j y_k y_l$ cancel, as do terms containing $x_k x_l y_i y_j$. The four remaining terms are precisely the expansion of $(x_i y_j - x_j y_i)(x_k y_l - x_l y_k)$. This proves the identity over $\mathbb{Z}$ and hence over any coefficient ring. □

Theorem 8.3 (The two-row standard monomial basis)**.** *The subalgebra $G_{2,n} = k[p_{ij}]$ has basis*

$$p_{i_1 j_1} \cdots p_{i_d j_d}, \qquad i_1 \leq \cdots \leq i_d, \quad j_1 \leq \cdots \leq j_d, \quad i_s < j_s. \tag{SM}$$

The relations (P), with p replaced by P, generate the kernel of $k[P_{ij}] \to G_{2,n}$.

Proof. Order the factors lexicographically by their column pairs (i, j). If the bottom row is not weakly increasing, some two columns are incomparable for the componentwise order. They have indices $i < k < l < j$, and (P) replaces their product by two products. In each replacement the lexicographically first changed column is either (i, l) or (i, k), strictly smaller than the old (i, j). After sorting all columns again, the whole column word has strictly decreased lexicographically. There are finitely many column words of the fixed degree, so straightening terminates. Its terminal products are exactly (SM).

For independence, order monomials lexicographically with $x_1 > \cdots > x_n > y_1 > \cdots > y_n$. The leading monomial of p_{ij} is $x_i y_j$. The leading monomial of a standard product is consequently

$$x_{i_1} \cdots x_{i_d} y_{j_1} \cdots y_{j_d}.$$

Such a monomial determines the two sorted rows, and those rows determine the standard product. Distinct standard products therefore have distinct leading monomials. In a nonzero linear combination, the largest of these leading monomials cannot cancel. This proves independence. Reducing a polynomial in the kernel to a standard linear combination now proves that the quadratic relations generate the entire kernel. □

The tableaux have rectangular shape (d, d): each column is strictly increasing, and both rows are weakly increasing. This gives the promised visible connection between minors and the tableau convention of the preceding chapter. For four indices the single relation takes the familiar form

$$p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23} = 0.$$

Only the final product is nonstandard in the componentwise column order.

The dimension can also be proved directly. For $n \geq 2$, localize at p_{12}. The equations imply

$$p_{ij} = \frac{p_{1i}p_{2j} - p_{1j}p_{2i}}{p_{12}}, \qquad 3 \leq i < j \leq n.$$

Conversely, give arbitrary values to $t \neq 0, u_j, v_j$, and take the matrix with columns

$$(1, 0)^{\mathsf{T}}, \quad (0, t)^{\mathsf{T}}, \quad (-v_j/t, u_j)^{\mathsf{T}} \quad (j \geq 3).$$

Its indicated minors are $p_{12} = t$, $p_{1j} = u_j$, $p_{2j} = v_j$. This proves that the localized ring is a Laurent polynomial ring in $2n - 3$ variables. The affine cone has dimension $2n - 3$, and the Grassmannian $\mathrm{Gr}(2, n)$ in Plucker coordinates has dimension $2(n - 2)$.

8.4 A complete complex example: dimensions $(1, 2, 1)$

Consider

$$k \xrightarrow{A} k^2 \xrightarrow{B} k, \qquad A = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, \quad B = \begin{pmatrix} b_1 & b_2 \end{pmatrix}.$$

Its coordinate ring is

$$C = k[a_1, a_2, b_1, b_2]/(f), \qquad f = a_1b_1 + a_2b_2. \tag{Q}$$

Proposition 8.4 (Geometry and basis of the basic complex)**.** *The residue classes of monomials not divisible by a_2b_2 form a basis of C. The ring is a domain of dimension three. Its affine scheme is smooth away from the origin, which is singular. Its Hilbert series in total degree is*

$$H_C(t) = \frac{1 - t^2}{(1 - t)^4} = \frac{1 + t}{(1 - t)^3}, \qquad \dim_k C_d = (d + 1)^2.$$

With $\deg a_i = (1, 0)$ *and* $\deg b_i = (0, 1)$, $\dim_k C_{p,q} = p + q + 1$.

Proof. The substitution

$$\begin{pmatrix} z_{11} & z_{12} \\ z_{21} & z_{22} \end{pmatrix} = \begin{pmatrix} a_1 & a_2 \\ -b_2 & b_1 \end{pmatrix}$$

takes the two-by-two determinant to f. Thus the ring is the rank-one Segre ring for $a = b = 2$, up to an invertible change of variables. The preceding theorem proves that it is a three-dimensional domain.

Alternatively, order the variables so that the leading monomial of f is a_2b_2. Every multiple of f has leading monomial divisible by a_2b_2, since leading monomials multiply in a polynomial ring over a field. Division by this one monic polynomial therefore gives unique remainders supported on the stated monomials. This proves that particular normal form.

The gradient is (b_1, b_2, a_1, a_2), so the hypersurface is smooth at every nonzero point by the hypersurface Jacobian criterion. At the origin its Zariski tangent space has dimension four whereas the local dimension is three, so it is singular. Multiplication by f is injective in the polynomial ring. The resulting graded exact sequence gives $H_C(t) = (1 - t^2)/(1 - t)^4$, whose degree d coefficient is $(d + 1)^2$. In bidegree (p, q) the same argument gives

$$(p + 1)(q + 1) - pq = p + q + 1$$

when $p, q > 0$; if either is zero, the identical answer comes directly from the polynomial monomial count. □

Here the straightening relation is simply $a_2b_2 = -a_1b_1$. It is both a shuffle-type relation and a determinant relation. The locus with $A \neq 0$ and $B \neq 0$ consists of complexes exact at their middle term: the inclusion $\operatorname{im} A \subseteq \ker B$ is then an inclusion of one-dimensional subspaces. The loci $A = 0$ and $B = 0$ are two planes in the threefold. Every nonzero point of either plane is smooth. Consequently a rank drop does not, by itself, identify a singular point.

There is an equally explicit representation-theoretic description. Let the input line be U, the middle space V, and the output line W. On points, change of bases acts by

$$(g_U, g_V, g_W) \cdot (A, B) = (g_V A g_U^{-1},\ g_W B g_V^{-1}).$$

On functions we use $(g \cdot h)(x) = h(g^{-1}x)$. The linear coordinate spaces are therefore $U \otimes V^*$ and $V \otimes W^*$.

Proposition 8.5 (The bihomogeneous Schur calculation)**.** *As a representation of this group,*

$$C_{p,q} \cong U^{\otimes p} \otimes (W^*)^{\otimes q} \otimes (\det V)^{-p} \otimes \operatorname{Sym}^{p+q} V. \tag{QS}$$

In characteristic zero, the middle factor is the irreducible representation $S_{(q,-p)}V$.

Proof. It suffices to check the middle $\mathrm{GL}(V)$ action; the powers of the outer lines follow from the bidegrees. The natural two-dimensional identity $V^* \cong (\det V)^{-1} \otimes V$, followed by symmetric multiplication, gives a surjection

$$\operatorname{Sym}^p V^* \otimes \operatorname{Sym}^q V \longrightarrow (\det V)^{-p} \otimes \operatorname{Sym}^{p+q} V.$$

In a basis with volume $e_1 \wedge e_2$, this identity sends e_1^* to $-e_2$ and e_2^* to e_1, with the determinant factor understood. The contraction relation $a_1b_1 + a_2b_2$ therefore maps to $-e_2e_1 + e_1e_2 = 0$. The map factors through $C_{p,q}$. It is surjective because each monomial of degree $p + q$ can be split into factors of degrees p and q. The source quotient and target both have dimension $p + q + 1$, so the induced map is an isomorphism. The last identification follows from the explicit highest-weight convention. □

This calculation checks all duals in a case where the entire answer can be verified without a general representation-theoretic decomposition theorem. Replacing $S_{(q,-p)}V$ by $\operatorname{Sym}^p V \otimes \operatorname{Sym}^q V$, for example, would give the wrong dimension and the wrong action.

8.5 The general complex and its standard products

Fix finite free modules $V_0, \dots, V_m$, of ranks $n_0, \dots, n_m$. We use the descending orientation

$$V_m \xrightarrow{D_m} V_{m-1} \longrightarrow \cdots \longrightarrow V_1 \xrightarrow{D_1} V_0. \tag{C}$$

Thus D_i has n_{i-1} rows and n_i columns, and the equations are $D_i D_{i+1} = 0$. At the endpoints we use $D_0 : V_0 \to 0$ and $D_{m+1} : 0 \to V_m$, both zero. Let A_R be the polynomial ring in all entries modulo all entries of these products.

For $I \subseteq [n_{i-1}]$ and $J \subseteq [n_i]$ with $|I| = |J| = p > 0$, write $[I \mid J]_i$ for the corresponding minor. Within a single matrix, order minors by

$$[I \mid J]_i \preceq [I' \mid J']_i \quad \Longleftrightarrow \quad p \geq q, \quad i_t \leq i'_t, \quad j_t \leq j'_t \quad (1 \leq t \leq q), \tag{M}$$

where $q = |I'|$. Minors of different matrices are declared incomparable in this partial order. This order by itself is not enough to describe standard products across different arrows.

Here is the additional rule, stated in the transpose convention of De Concini–Eisenbud–Procesi, Section 16 [41]. For a pair of minors, interchange the two factors if necessary so that their arrow indices satisfy $i \geq j$. The pair is standard if one of the following holds:

1. $i \geq j + 2$;
2. $i = j$ and the two minors are comparable under (M);
3. $j = i - 1$, and the following complement condition holds.

For the last condition write $x = [I \mid J]_i$, of size p, and $x' = [I' \mid J']_{i-1}$, of size q. The shared vector space has dimension $n = n_{i-1}$. List $[n] \setminus I = \{\ell_1 < \cdots < \ell_{n-p}\}$ and $J' = \{j'_1 < \cdots < j'_q\}$. The condition is

$$q \leq n - p, \qquad j'_t \leq \ell_{\,n-p-q+t} \quad (1 \leq t \leq q). \tag{MC}$$

A product of minors is standard when every pair of its factors is standard. For (Q), the rule says that $a_2 b_2$ is nonstandard while the other three mixed pairs are standard, exactly as the direct calculation found.

Theorem 8.6 (External standard basis theorem for complexes)**.** *The standard products just defined form an R-basis of A_R. For nonnegative integer bounds r_i, if all minors of D_i larger than r_i are set equal to zero, the surviving standard products form a basis of the quotient. These constructions commute with base change. Over a field, the quotients are reduced. If $r_0 = r_{m+1} = 0$ and*

$$r_i + r_{i+1} \leq n_i \qquad (0 \leq i \leq m), \tag{A}$$

the rank-bounded quotient is a normal Cohen–Macaulay domain.

The basis and reducedness assertions are De Concini–Eisenbud–Procesi, Theorem 16.1 and Corollary 16.2; the final assertion follows from their Theorem 16.3 and Proposition 16.4, which in turn refer to De Concini–Strickland [40, 41]. The general topological and commutative-algebraic

arguments establishing these results are not replaced here by the elementary examples. In particular, the existence of a resolution of a variety would not alone prove its normality.

Two features of the statement deserve attention. First, the basis includes products of minors of different sizes, whose polynomial degrees are those sizes. It is not a basis just in the degree-one matrix entries. Second, since nonadjacent arrows may form standard products despite their minors being incomparable in (M), this is a generalized Hodge algebra. Its standardness is not simply the chain condition in one ordinal partial order. That distinction will determine which simplicial complex appears in a monomial degeneration.

8.6 Rank strata and orbit closures without a source shortcut

For the following geometric argument let k be algebraically closed. Define C_r to be the closed locus in (C) with $\operatorname{rank} D_i \le r_i$, where r satisfies (A). Let C_r° denote the locus with all these ranks equal to r_i. The use of equality here is essential: strictly smaller ranks would not define the orbit indexed by r.

Proposition 8.7 (Orbits, closure, and dimension)**.** *The locus C_r° is a single orbit under $\prod_{i=0}^m \mathrm{GL}(V_i)$, and $C_r = \overline{C_r^\circ}$. It is irreducible for every admissible r. The irreducible components of the full space of complexes are precisely the C_r for coordinatewise maximal admissible r. Moreover*

$$\dim C_r = \sum_{i=1}^{m} r_i(n_i + n_{i-1} - r_i) - \sum_{i=1}^{m-1} r_i r_{i+1}. \tag{D}$$

Proof. At each vertex choose subspaces

$$V_i = B_i \oplus H_i \oplus T_i, \qquad B_i = \operatorname{im} D_{i+1}, \quad B_i \oplus H_i = \ker D_i.$$

The dimensions are r_{i+1}, $n_i - r_i - r_{i+1}$, and r_i, respectively. The only nonzero block of D_i is an isomorphism $T_i \to B_{i-1}$. Changing bases makes all these isomorphisms identity matrices. Thus the ranks determine a unique orbit.

Decrease one rank by one by replacing a single identity entry in the corresponding block by a parameter t. For $t \neq 0$ the old ranks remain; at $t = 0$ precisely that rank drops. The compositions stay zero because the direct-sum blocks have not changed. Iterating shows that every smaller rank orbit is in the closure. Conversely, the vanishing of the $(r_i + 1)$-minors is closed and prevents any larger rank from entering the closure. This proves the closure assertion. The group is irreducible, hence so is its orbit closure. There are finitely many admissible rank vectors; their containment order is coordinatewise. The maximal ones therefore give exactly the components.

For the dimension, first choose flags $B_i \subset Z_i \subset V_i$, with $\dim B_i = r_{i+1}$ and $\dim Z_i = n_i - r_i$. Choosing B_i contributes $r_{i+1}(n_i - r_{i+1})$ parameters. Choosing Z_i/B_i in V_i/B_i contributes $r_i(n_i - r_i - r_{i+1})$. Finally choose each isomorphism $V_i/Z_i \to B_{i-1}$, an open subset of a space of dimension r_i^2. Every complex in the stratum determines these flags and isomorphisms uniquely. Adding the dimensions and cancelling the repeated squares gives (D). □

This proof also corrects a possible ambiguity in phrases such as *the irreducible rank sequence*. Every admissible rank bound gives an irreducible orbit closure; maximality is needed for being an irreducible component of the full space. The equations imposing rank at most r_i use minors of size $r_i + 1$, not r_i. These distinctions are visible in the normal form itself and do not depend on terminology.

8.7 The Schur decomposition with all vertices and duals

Return to an algebraically closed field k of characteristic zero. On the descending complex, let

$$(g_0,\dots,g_m)\cdot D_i = g_{i-1}D_i g_i^{-1}, \qquad (g\cdot f)(D) = f(g^{-1}\cdot D). \tag{G}$$

The linear coordinates on the ith arrow are therefore $V_i \otimes V_{i-1}^*$. For a tuple of partitions $\lambda = (\lambda_1,\dots,\lambda_m)$, put $\lambda_0 = \lambda_{m+1} = \varnothing$. Require $\ell(\lambda_i) \le r_i$. At vertex j, define the dominant weight

$$\eta_j(\lambda) = \big(\lambda_j,\, 0^{\,n_j-\ell(\lambda_j)-\ell(\lambda_{j+1})},\, -\operatorname{rev}(\lambda_{j+1})\big). \tag{W}$$

Condition (A) ensures that the number of zero entries is nonnegative.

Theorem 8.8 (External Schur filtration and its characteristic-zero splitting)**.** *The coordinate ring $k[C_r]$ has a multiplicative tableau filtration whose successive modules are*

$$\bigotimes_{j=0}^{m} S_{\eta_j(\lambda)}V_j, \qquad \ell(\lambda_i)\le r_i. \tag{SF}$$

The summand indexed by λ has arrow multidegree $(|\lambda_1|,\dots,|\lambda_m|)$. In characteristic zero this gives a $\prod_j \mathrm{GL}(V_j)$-module isomorphism

$$k[C_r] \cong \bigoplus_{\substack{\lambda_1,\dots,\lambda_m\\ \ell(\lambda_i)\le r_i}} \bigotimes_{j=0}^{m} S_{\eta_j(\lambda)}V_j. \tag{SD}$$

The filtration is the Schur straightening construction of De Concini–Strickland, reviewed explicitly in Carroll–Weyman, Section 4 [32, 40]. Their tableaux concatenate an incoming dual diagram, expressed using a determinant twist, with the outgoing diagram. Formula (W) is that construction in the action convention (G). Rank bounds remove exactly the partitions whose column length exceeds the allowed rank. Each finite multidegree is a rational representation of a reductive group in characteristic zero, so its filtration splits as a representation. This last step does not assert a canonical splitting as an algebra.

There are $m+1$ factors in (SF), including both endpoints. Omitting the last vertex, or treating both appearances of an intermediate vector space as covariant, gives a wrong representation. For one arrow the formula reduces to the Cauchy decomposition

$$\mathrm{Sym}^d(V_1\otimes V_0^*) \cong \bigoplus_{|\lambda|=d} S_\lambda V_1 \otimes S_\lambda V_0^*,$$

with the partition length cut off by a rank bound when one is present. For (Q), take $\lambda_2 = (p)$ and $\lambda_1 = (q)$. The middle weight is exactly $(q,-p)$, and (SD) reproduces the independently proved formula (QS).

Multiplication contains further information. Multiplying two tableau expressions of shapes λ,μ lies in the filtration piece indexed by $\lambda+\mu$; straightening can add terms in lower pieces. The associated graded multiplication retains the top piece, often described in terms of highest-weight or Cartan multiplication. One cannot reconstruct every product in the original ring simply by writing a direct sum of representation spaces. An isomorphism of modules answers a linear question; an isomorphism of algebras must also answer this multiplicative question.

Chapter 9

Associated graded rings, semigroups, and toric limits

9.1 Three different constructions

There are three constructions in the preceding discussion that should be distinguished. A Schur filtration organizes the coordinate ring into representations. A straightening degeneration replaces nonstandard products by zero and produces a monomial quotient. A suitable subalgebra degeneration replaces determinants by their diagonal monomials and can produce an integral semigroup ring. They can start from the same ring and end in different rings. The difference is mathematical, not just a choice of notation.

For an ascending, exhaustive filtration $F_0A \subseteq F_1A \subseteq \cdots$ with $1 \in F_0A$ and $F_iA\,F_jA \subseteq F_{i+j}A$, the associated graded algebra is

$$\operatorname{gr}_F A = \bigoplus_{d\geq 0} F_dA/F_{d-1}A, \qquad F_{-1}A = 0.$$

Its product keeps only the highest filtration class of a product. A nonzero product in A can fall into a lower filtration piece and therefore become zero in the associated graded ring.

The Rees algebra records the lost information:

$$\mathcal{R}_F(A) = \bigoplus_{d\geq 0} F_dA\,t^d \subseteq A[t].$$

The quotient at $t = 1$ recovers A, while the quotient by t recovers $\operatorname{gr}_F A$. For the latter statement, the degree d coefficient of $t\mathcal{R}_F(A)$ is $F_{d-1}A$, which gives exactly the graded quotient above. The former follows by evaluation at one and finite telescoping of the coefficients. Over a field, choose a vector-space basis adapted to the filtration by successively extending a basis of $F_{d-1}A$ to one of F_dA. If a chosen basis vector b first appears at stage $\delta(b)$, the elements $b\,t^{\delta(b)}$ form a $k[t]$-basis of the Rees algebra: every later occurrence of b is a multiple by a power of t, and independence follows inside $A[t]$. Thus this algebra is free, hence flat, over $k[t]$. Finite generation as an algebra requires a suitable finite filtration; it is not automatic for every abstract filtration.

Flatness preserves the Hilbert function in the graded situations used below. It does not assert that every fiber is irreducible, smooth, or a domain. The following explicit construction makes that warning verifiable.

9.2 A monomial degeneration for a quadratic straightening algebra

Let P be a finite partially ordered set, and let A be a graded algebra generated by elements x_p, $p \in P$, all of degree one. Suppose that the chain products

$$x_{p_1} \cdots x_{p_d}, \qquad p_1 \le \cdots \le p_d,$$

form a basis. Suppose also that every incomparable pair p, q has a straightening relation

$$x_p x_q = \sum_{u \le v} c_{uv} x_u x_v, \qquad c_{uv} \neq 0 \Longrightarrow u < p,\ u < q. \tag{ASL}$$

These are the homogeneous quadratic hypotheses needed for the next argument. General Hodge algebras can have generators of different degrees, so they must not be substituted into this statement without checking those hypotheses.

Proposition 9.1 (An explicit straightening degeneration)**.** *Under these assumptions, A has a flat graded family over $k[t]$ whose fiber at one is A and whose fiber at zero is*

$$k[X_p : p \in P]/(X_p X_q : p, q \text{ incomparable}). \tag{SR}$$

The chain monomials form a free $k[t]$-basis of the family.

Proof. If P is empty, then $A = k$ and the family $k[t]$ proves the assertion. Assume from now on that $P \neq \varnothing$. Let $h(p)$ be the maximum number of strict steps in a chain starting at p and going upward. Put $D_p = 3^{h(p)}$, choose an integer $M > \max D_p$, and set $w_p = M - D_p > 0$. If a term $x_u x_v$ occurs on the right of (ASL), then

$$D_u \ge 3 \max(D_p, D_q) > D_p + D_q.$$

It follows that $w_u + w_v < w_p + w_q$. Thus straightening strictly decreases the positive integer weight of a monomial.

In a polynomial presentation, every element reduces by these relations to a combination of chain monomials. Since their images are independent in A, a relation reduces to zero. This proves both that the straightening relations generate the presentation ideal and that their initial ideal for this weight order is exactly the ideal of incomparable products. In particular no additional chain monomial is killed by taking initial terms.

Now replace each relation by

$$X_p X_q - \sum c_{uv} t^{\,w_p + w_q - w_u - w_v} X_u X_v.$$

All displayed exponents of t are strictly positive. The same decreasing weight argument gives spanning over $k[t]$ by chain monomials. After inverting t, the change $X_p = t^{w_p} x_p$ identifies this algebra with $A[t, t^{-1}]$. If a $k[t]$-linear combination of chain monomials were zero, this change would give a relation among the corresponding basis monomials of A, with Laurent polynomial coefficients. Each coefficient would be zero. The original polynomial coefficients are therefore zero as well. The family is free over $k[t]$, hence flat, and specializing its relations gives the two stated fibers. □

The quotient (SR) is the Stanley–Reisner ring of the order complex of P: a subset is a face exactly when it is a chain. Whenever there are incomparable p, q, the images of X_p and X_q are nonzero but their product is zero. Thus this ring has zero divisors.

For a two-by-two rank-one matrix, write its entries p, q, r, s, with relation $ps - qr = 0$. The only incomparable pair in the grid order is q, r. The monomial limit is

$$k[p, q, r, s]/(qr),$$

the union of the two coordinate hyperplanes $q = 0$ and $r = 0$. The original ring is a domain by the Segre proof. Calling this monomial limit a toric domain would therefore be an error.

For varieties of complexes, the general Hodge-algebra construction gives a stepwise Rees degeneration to the ring governed by its nonstandard products [41, Section 3]. Because the forbidden products are squarefree pairs, that discrete ring is the Stanley–Reisner ring of a flag simplicial complex: a set is a face exactly when all its pairs are standard. It need not be the order complex of the minor order (M), since standardness across different arrows uses the extra rule (MC).

This identifies precisely what the combinatorial basis supplies. It supplies a free module basis, a multiplication algorithm, and a controlled degeneration. Normality and the Cohen–Macaulay property of the original rank variety come from the stronger theorems cited in the preceding chapter. They cannot be inferred merely by observing that some associated graded ring has monomial equations.

9.3 Initial subalgebras and the diagonal semigroup

An initial subalgebra is a different construction from an initial quotient. If $G \subseteq k[z_1, \dots, z_N]$ is a subalgebra and a term order has been chosen, define

$$\mathrm{in}(G) = k[\mathrm{LM}(f) : 0 \neq f \in G],$$

where LM denotes the leading monomial. A set of generators is called a SAGBI basis if their leading monomials generate this entire initial subalgebra. This is the subalgebra counterpart of a Groebner basis; arbitrary algebra generators do not automatically have this property.

Proposition 9.2 (The two-row minors are a SAGBI basis)**.** *For the lexicographic order $x_1 > \cdots > x_n > y_1 > \cdots > y_n$, the minors p_{ij}, $i < j$, are a SAGBI basis of $G_{2,n}$. Their initial subalgebra is*

$$k[x_i y_j : 1 \leq i < j \leq n]. \tag{T}$$

Proof. Expand any nonzero element of $G_{2,n}$ in the standard monomial basis (SM). Those basis elements have distinct leading monomials, as proved earlier. The largest of their leading monomials cannot cancel. It is a product of the $x_i y_j$. Thus every leading monomial of an element of $G_{2,n}$ belongs to the subalgebra on the right of (T). The reverse inclusion holds because each $x_i y_j$ is the leading monomial of the corresponding minor. This proves the claim. □

Unlike (SR), the ring (T) is a subring of a polynomial ring and is therefore a domain. Its presentation uses binomial equalities between products having the same exponent vector. Those products do not become zero.

For clarity, an affine semigroup here means a finitely generated additive submonoid $S \subseteq \mathbb{N}^N$. Its semigroup ring has basis symbols z^s, $s \in S$, and multiplication $z^s z^{s'} = z^{s+s'}$. The embedding in a polynomial ring is part of this definition. For (T), the generators of S are the vectors (e_i, e_j), $i < j$, in $\mathbb{N}^n \times \mathbb{N}^n$.

Theorem 9.3 (The diagonal semigroup and its inequalities)**.** *This semigroup consists exactly of the pairs* $(\alpha, \beta) \in \mathbb{N}^n \times \mathbb{N}^n$ *satisfying*

$$|\alpha| = |\beta|, \qquad \sum_{j=1}^{k} \beta_j \le \sum_{i=1}^{k-1} \alpha_i \quad (1 \le k \le n). \tag{I}$$

Each such pair determines a unique semistandard two-row tableau by sorting its top and bottom rows. The semigroup is saturated in the lattice it generates.

Proof. An expression as a sum of generators is a matching of top indices i with bottom indices $j > i$, with multiplicities. Any bottom index at most k must be matched to a top index smaller than k. This proves all the prefix inequalities and the equality of total sizes.

Conversely, process the bottom indices in order $j = 1, \dots, n$. Before processing j, the available top-index capacity with $i < j$ is

$$\sum_{i<j} \alpha_i - \sum_{h<j} \beta_h.$$

All previous matches used indices in that set. Inequality (I) says that the available capacity is at least β_j. Choose that many units of capacity, making edges from the available $i < j$ to j. The process never fails. At the end, the equality of totals shows that all capacity has been used. Thus the pair is a sum of the specified generators.

Sort the top indices as $i_1 \le \cdots \le i_d$ and the bottom indices as $j_1 \le \cdots \le j_d$. If $j_s = k$, the number of bottom entries at most k is at least s. The prefix inequality ensures that at least s top entries are smaller than k. Hence $i_s < j_s$, so these rows form a semistandard tableau. Their multiplicities uniquely determine the sorted rows, proving uniqueness.

Finally let L be the subgroup generated by S, and suppose $qv \in S$ for an integer $q > 0$ and $v \in L$. Every coordinate of v is nonnegative because the same is true of qv. The equality and inequalities in (I) are homogeneous, so they also hold for v. The characterization just proved gives $v \in S$. □

The first and last prefix conditions force $\beta_1 = 0$ and $\alpha_n = 0$, as they should: there is no column with bottom entry one or top entry n. These boundary conditions are sometimes lost when a tableau semigroup is described only by equal total degrees.

9.4 Gelfand–Tsetlin patterns in the example

Put

$$A_r = \sum_{i \le r} \alpha_i, \qquad B_r = \sum_{j \le r} \beta_j, \qquad A_0 = B_0 = 0.$$

The conditions (I) become

$$A_r \ge A_{r-1} \ge B_r \ge B_{r-1}, \qquad A_n = B_n = d. \tag{GT}$$

The pair (A_r, B_r) is the shape of the subtableau formed by entries at most r. Its interlacing with the preceding pair is exactly the two-row Gelfand–Tsetlin condition. Taking successive differences recovers α, β, so (GT) is an equivalent integral description, not a second source of unverified basis vectors.

Addition of exponent vectors becomes componentwise addition of the patterns. This is the semigroup operation in the toric limit. In the original determinant algebra, multiplying the two

corresponding standard products may require straightening, with several terms. The pattern addition therefore records the leading product, not every coefficient in the original multiplication table.

The saturation in the preceding theorem also has a direct commutative-algebraic consequence.

Proposition 9.4 (Normality from these integral inequalities)**.** *The semigroup rings $k[S_{\mathrm{Segre}}]$ and $k[S]$ of (T) are normal domains.*

Proof. Both semigroups are the intersection of their generated lattice L with finitely many homogeneous integral linear half-spaces: coordinate nonnegativity and, for S, the prefix inequalities. The required total-degree equalities hold on the lattice L; the inequalities select its semigroup points. This follows from their explicit characterizations.

The group ring $k[L]$ is a Laurent polynomial ring after choosing a basis of the free abelian group L. It is a normal domain. Its fraction field is also the fraction field of $k[S]$, because every lattice element is a difference of two semigroup elements. For each defining integer linear functional ℓ, give a nonzero Laurent polynomial the value

$$v_\ell\left(\sum c_u z^u\right) = \min_{c_u \neq 0} \ell(u).$$

The lowest homogeneous parts of two Laurent polynomials have nonzero product, so these values add under multiplication. They extend by subtraction to a valuation of the fraction field.

Suppose f belongs to this common fraction field and is integral over $k[S]$. It is integral over $k[L]$, so belongs to that Laurent polynomial ring. It belongs to every valuation ring $v_\ell \geq 0$: otherwise an integral equation for f would have its monic highest power as the unique term of smallest valuation, which cannot sum to zero. Thus all exponents of f satisfy every defining inequality. They lie in S, and $f \in k[S]$. This proves integral closedness. The domain assertion follows from the monomial embeddings. □

This proof uses the actual inequality description; it does not infer normality from finite generation alone. A finitely generated semigroup can fail saturation, as $\langle 2, 3\rangle \subseteq \mathbb{N}$ shows: $2 \cdot 1$ belongs to it but 1 does not. Correspondingly, $k[t^2, t^3]$ is not normal, since t is integral over it and lies in its fraction field without belonging to the ring.

9.5 Two flat families from the same Grassmannian

The difference between the two degenerations can be made entirely explicit for $n = 4$. Let P denote its six Plucker variables and form the two $k[t]$-algebras

$$\begin{aligned} \mathcal{A}_{\mathrm{mon}} &= k[t, P]/(P_{14}P_{23} - tP_{13}P_{24} + tP_{12}P_{34}), \\ \mathcal{A}_{\mathrm{tor}} &= k[t, P]/(P_{14}P_{23} - P_{13}P_{24} + tP_{12}P_{34}). \end{aligned} \tag{F}$$

Both fibers at $t = 1$ are the affine Plucker cone over $\mathrm{Gr}(2, 4)$.

Proposition 9.5 (Flatness and different special fibers)**.** *Both algebras in (F) are free over $k[t]$, with basis the monomials not divisible by $P_{14}P_{23}$. Their special fibers are, respectively,*

$$k[P]/(P_{14}P_{23}), \qquad k[P]/(P_{14}P_{23} - P_{13}P_{24}).$$

The first has zero divisors and the second is a normal domain.

Proof. Choose a term order on the six P variables for which $P_{14}P_{23}$ is larger than the other two products. Each relation is monic in this leading monomial. Division over the coefficient ring $k[t]$ gives a remainder supported on the listed monomials. The leading term of a nonzero multiple of the relation is divisible by $P_{14}P_{23}$, so no nonzero remainder lies in the ideal. This proves uniqueness and freeness.

At zero, the first relation is the product of two distinct variables, which proves the zero-divisor assertion. The second relation is a two-by-two determinant in four of the variables. Its ring is the two-by-two Segre ring with two independent polynomial variables adjoined. The preceding normality proof, or the same semigroup argument with two free coordinates, proves it is a normal domain. □

For any nonzero value of t, rescaling P_{14} in the first family and P_{12} in the second identifies the fiber with the original quadric. Thus these are actual one-parameter degenerations, not merely two unrelated examples. In both cases the degree d dimension is

$$\binom{d+5}{5} - \binom{d+3}{5} = \frac{(d+1)(d+2)^2(d+3)}{12},$$

where the second binomial is zero when $d < 2$. Agreement of this Hilbert function does not make the two special fibers isomorphic. One is reducible and the other is integral.

The general study of toric limits of flag and Schubert varieties extends this pattern; it requires its own filtration or subalgebra argument [55]. The completely proved two-row case here is sufficient to show why an ordinary straightening degeneration cannot be renamed a toric degeneration without that additional argument.

9.6 How the algebra enters the larger research program

The construction gives a precise interface with invariant theory. At a vertex j of the descending complex, a scalar matrix c_jI acts on the Schur factor (W) by

$$c_j^{|\lambda_j|-|\lambda_{j+1}|}.$$

This follows either by summing the highest-weight entries or directly from the source and target coordinate actions in (G). It is therefore possible to test necessary central-character conditions for an invariant without expanding every polynomial. To obtain a semi-invariant one must also identify the appropriate one-dimensional determinant character at each vertex. Balancing scalar degrees alone does not prove that an entire Schur module is one-dimensional.

For gentle algebras whose colored paths describe complexes, the same vertex space can appear in several such factors. Taking invariants then requires decomposing their tensor product, or applying the precise semi-invariant criterion. Carroll–Weyman carry out that additional step under their quiver hypotheses [32]. The following chapters give the matching semigroup and determinantal examples. The coordinate ring of a single complex and the semi-invariant ring of a gentle algebra are distinct objects, even when both admit useful tableau descriptions.

There is also a useful integral interface with the theory of orders. The exterior and shuffle formulas are integral, and the standard bases are free over the base ring. One can therefore perform the same finite normal-form calculations before and after reduction modulo a prime. This controls the ranks of the indicated graded modules and their Hilbert functions. The characteristic-zero decomposition into irreducible representations does not survive merely because the basis does; semisimplicity is an additional property of the coefficient field.

For computation, the bases give exact finite representations of functions on the rank varieties. The straightening algorithm can multiply those functions and reduce equations to normal form. The leading semigroup can organize enumeration and Hilbert function checks. These are algebraic algorithms with costs that depend on degrees, dimensions, and coefficient sizes. A basis theorem supplies no automatic polynomial-time algorithm for elliptic-curve discrete logarithms and no error analysis for a lattice cryptosystem. Such claims would need their own reduction and complexity analysis.

The same separation applies to spectral questions. Schur characters, Hilbert functions, and K-theory classes record algebraic or representation-theoretic information. They do not specify a Hilbert-space inner product, a self-adjoint Hamiltonian, or a uniform lower bound on its positive spectrum. A flat family can preserve every graded dimension while changing its singularities. A family of finite exact complexes can likewise have positive eigenvalues tending to zero. The mass-gap transport problem therefore requires the analytic hypotheses stated elsewhere in this monograph.

Evidence status. The exterior basis, shuffle identities, contraction homotopy, elementary Schur sequences, characteristic-zero tensor-complex argument, Segre and two-row Plucker bases, basic-quadric calculations, rank-orbit dimension formula, quadratic straightening family, diagonal semigroup description, its normality, and the two explicit flat families are proved in these chapters. The general Schur basis theorem, general complex standard basis and normality theorems, and general Schur filtration are cited external theorems with stated hypotheses. The Lean files verify selected interfaces described in the formalization appendix; they do not formalize these general standard monomial theorems.

Chapter 10

Gentle algebras, matching semigroups and determinant coordinates

Evidence status. The elementary decompositions, determinant transformations and worked semigroup presentation below have proofs in this chapter. General results on gentle semi-invariants and rational quotients are attributed to their primary sources with hypotheses retained. These are additional mathematical proofs, not additional Lean declarations.

10.1 From a relation to a variety of complexes

Let k be an algebraically closed field of characteristic zero throughout this chapter, except where a result is explicitly stated over an arbitrary field. A bound quiver algebra $A = kQ/I$ is *gentle* when it is finite dimensional, every vertex has at most two incoming and two outgoing arrows, and the following restrictions hold. For any arrow, there is at most one possible successor whose composition is in I, at most one successor whose composition is not in I, and the analogous two restrictions hold for predecessors. The ideal I is generated by paths of length two. Thus two different kinds of combinatorics are present: paths that survive in A, and paths whose consecutive pairs are relations. A drawing of a quiver without its relations does not determine the algebra.

The basic example is

$$1 \xrightarrow{a} 2 \xrightarrow{b} 3, \qquad ba = 0.$$

The algebra has basis e_1, e_2, e_3, a, b and dimension five. A representation of dimension vector (p, q, r) is exactly a pair of matrices

$$k^p \xrightarrow{U} k^q \xrightarrow{V} k^r, \qquad VU = 0.$$

Consequently the representation scheme is the scheme of two consecutive differentials. Its defining equations, ranks, singular strata and coordinate ring are the objects studied in the preceding chapters. This is an equality of algebraic constructions, not an analogy between two diagrams.

Proposition 10.1 (Complete module decomposition for the two-arrow relation)**.** *Over any field, a representation $k^p \xrightarrow{U} k^q \xrightarrow{V} k^r$ with $VU = 0$, $\operatorname{rank} U = u$, and $\operatorname{rank} V = v$, decomposes as*

$$M_{12}^u \oplus M_{23}^v \oplus S_1^{p-u} \oplus S_2^{q-u-v} \oplus S_3^{r-v}. \tag{G1}$$

Here S_i is the simple module at vertex i, while M_{12} and M_{23} have a one-dimensional space at the indicated two vertices and the identity on the joining arrow. These five modules are precisely the indecomposable modules.

Proof. Choose a complement P to $\ker U$ in k^p. Then $U : P \to \operatorname{im} U$ is an isomorphism. Since $\operatorname{im} U \subseteq \ker V$, choose

$$k^q = \operatorname{im} U \oplus H \oplus T, \qquad \ker V = \operatorname{im} U \oplus H.$$

The restriction $V : T \to \operatorname{im} V$ is an isomorphism. Complete $\operatorname{im} V$ to a direct-sum decomposition of k^r. Bases transported along the two isomorphisms exhibit the two collections of identity arrows. The remaining summands have all arrow maps zero and split into the displayed simples. Their multiplicities are the dimensions of the chosen spaces, giving (G1). Each of the five listed modules has endomorphism ring k, so has no nontrivial idempotent and is indecomposable. The decomposition shows that no others occur. □

For dimension vector $(1, 2, 1)$, the four rank pairs give

(u, v)	module
$(0, 0)$	$S_1 \oplus S_2^2 \oplus S_3$
$(1, 0)$	$M_{12} \oplus S_2 \oplus S_3$
$(0, 1)$	$S_1 \oplus S_2 \oplus M_{23}$
$(1, 1)$	$M_{12} \oplus M_{23}$.

Writing $U = (x_1, x_2)^t$ and $V = (y_1, y_2)$, the equation is $x_1y_1 + x_2y_2 = 0$. The rank-$(1, 1)$ locus is dense in this three-dimensional quadric. In dimension $(1, 1, 1)$, however, the equation is $xy = 0$, and the two maximal rank pairs $(1, 0)$, $(0, 1)$ give two distinct irreducible components. Its coordinate ring $k[x, y]/(xy)$ has zero divisors. It cannot be a subalgebra of a Laurent polynomial ring and therefore cannot be the ring of a cancellative affine semigroup in a lattice. This small example prevents an incorrect universal assertion that every complex scheme is itself an irreducible toric variety.

10.1.1 Color chains and products

Suppose now that Q has no oriented cycle. Starting at an arrow without a forbidden predecessor, follow the unique forbidden successor as long as one exists. The gentle restrictions and acyclicity ensure that this procedure terminates and partitions the arrows into directed chains. Color each such chain separately. The ideal I is generated by the monochromatic consecutive compositions.

For fixed vector spaces at the vertices, each color uses its own collection of matrix variables. Relations involving one color do not involve variables of another color. The polynomial algebra on all arrow matrices is the tensor product of the color polynomial algebras, and the relation ideal is the sum of the extended color ideals. Therefore

$$k[\operatorname{Rep}(A, \beta)] \cong \bigotimes_{\text{colors } s} k[\text{complex scheme along } s]. \tag{G2}$$

This assertion holds at the scheme level, including any scheme structure specified by the composition equations. It explains why straightening on complexes is integral to gentle invariant theory [32, Sections 2–5]. However, two colors can meet at the same vertex, and the same $\mathrm{GL}(\beta_x)$ acts on both. Taking invariants of (G2) therefore does not mean taking the tensor product of the separate invariant rings. The coupling of the two Schur factors at each shared vertex is the central additional step.

When oriented cycles are allowed, a forbidden chain may close up. It then yields a *cyclic* complex, and the finite-chain statements must be reconsidered. Acyclicity in the quoted gentle-algebra theorems is a mathematical hypothesis, not an optional simplification of the picture.

10.2 A second example: Kronecker pencils and regular modules

The Kronecker quiver has two parallel arrows $a, b : 1 \to 2$, with no relations. Its path algebra is four dimensional because no two arrows compose. It is gentle. In dimension $(1, 1)$, its representation space is k^2, with coordinates (a, b), and

$$(g_1, g_2) \cdot (a, b) = (g_2 g_1^{-1} a, g_2 g_1^{-1} b).$$

All polynomial $\mathrm{GL}_1 \times \mathrm{GL}_1$-invariants are constants: invariance under common scalar multiplication forces each nonconstant homogeneous part to vanish. The two coordinates are nevertheless semi-invariants with the same point-evaluation character g_2/g_1. Since SL_1 is trivial, the total semi-invariant ring is $k[a, b]$. Taking the projective quotient of its positive-degree part gives $\mathbb{P}^1$. Thus an invariant ring that sees only constants and a semi-invariant ring that parametrizes a projective family are entirely compatible.

In dimension (n, n), a representation is a matrix pencil (A, B) with simultaneous left-right equivalence

$$(A, B) \longmapsto (g_2 A g_1^{-1}, g_2 B g_1^{-1}).$$

On the open locus $\det A \neq 0$, choose bases so that $A = I_n$. The residual equivalence is conjugation of $T = A^{-1}B$. Jordan decomposition consequently gives all indecomposable regular modules on this chart:

$$(I_n, J_n(\lambda)).$$

Their endomorphism ring is $k[t]/(t - \lambda)^n$, which is local, so they are indecomposable. The other chart includes $(J_n(0), I_n)$ at the point at infinity. For nonzero finite parameter, these are the usual band families; special zero and infinity presentations require the usual string/band conventions. The general string-algebra classification is due to Butler–Ringel and is not being reproved here [29].

The polynomial

$$D_{A,B}(s, t) = \det(sA + tB)$$

transforms by

$$D_{g_2 A g_1^{-1}, g_2 B g_1^{-1}}(s, t) = \frac{\det g_2}{\det g_1} D_{A,B}(s, t). \tag{G3}$$

Every coefficient is a semi-invariant of the same character. Ratios of nonzero coefficients are rational invariants. They describe spectral parameters but cannot recover every extension. For example, $(I_2, \lambda I_2)$ is decomposable and $(I_2, J_2(\lambda))$ is indecomposable, yet both have determinant $(s + t\lambda)^2$. This is an explicit instance of the distinction between semisimplification data and the module itself. A proposed arithmetic determinant comparison must say which of these data it retains.

10.3 Homological qualifications: a gentle algebra of infinite global dimension

Consider the oriented triangle $1 \to 2 \to 3 \to 1$, with every length-two path set to zero. This algebra is gentle and finite dimensional, with basis the three idempotents and three arrows. In the representation convention used here, the indecomposable projective P_i has top S_i and radical S_{i+1}, with indices taken modulo three. Hence

$$\Omega S_i = S_{i+1}, \qquad \Omega^3 S_i = S_i. \tag{G4}$$

Splicing the projective covers gives an infinite periodic resolution.

Proposition 10.2. *For this algebra and $n \geq 0$,*

$$\dim_k \operatorname{Ext}^n_A(S_i, S_j) = \begin{cases} 1 & j \equiv i+n \pmod 3, \\ 0 & \text{otherwise.} \end{cases}$$

In particular the global dimension is infinite.

Proof. The n-th projective in the minimal resolution of S_i is P_{i+n}. Applying $\operatorname{Hom}_A(-, S_j)$ kills all differentials, since the resolution maps land in radicals and the target is semisimple. The term $\operatorname{Hom}_A(P_{i+n}, S_j)$ is k if the top is S_j, and zero otherwise. These are therefore the cohomology groups. Nonzero groups in arbitrarily high degrees preclude finite global dimension. □

This example connects the periodic resolutions of orders to quiver algebra in a controlled way. It also shows why finite Euler sums and two-term determinant complexes require a projective-dimension hypothesis. One cannot use the hereditary formula $\dim \mathrm{Hom} - \dim \mathrm{Ext}^1$ for every gentle module merely because its algebra is tame or its quiver is easy to draw.

10.4 Schofield determinants, with the transformation law derived

Let Q now be a finite acyclic quiver without imposed relations. Let V, W be representations of dimension vectors α, β. Define

$$C^0(V,W) = \bigoplus_x \operatorname{Hom}_k(V_x, W_x), \qquad C^1(V,W) = \bigoplus_{a:x\to y} \operatorname{Hom}_k(V_x, W_y)$$

and

$$d^V_W(\phi)_a = W_a \phi_x - \phi_y V_a. \tag{G5}$$

The kernel is $\operatorname{Hom}_Q(V,W)$ by the definition of an intertwiner. The cokernel describes extensions.

Proposition 10.3 (The elementary extension complex)**.** *There is an exact sequence*

$$0 \longrightarrow \operatorname{Hom}_Q(V,W) \longrightarrow C^0(V,W) \xrightarrow{d^V_W} C^1(V,W) \longrightarrow \operatorname{Ext}^1_Q(V,W) \longrightarrow 0. \tag{G6}$$

Proof. Any extension $0 \to W \to E \to V \to 0$ splits as vector spaces at each vertex. Relative to these splittings its arrow maps have blocks

$$E_a = \begin{pmatrix} W_a & \eta_a \\ 0 & V_a \end{pmatrix}.$$

There are no path relations, so every collection $\eta_a : V_x \to W_y$ defines a representation and hence an extension. Changing the vertex splittings by $\phi_x : V_x \to W_x$ changes η by $d^V_W(\phi)$, up to the fixed sign convention for the chosen change of basis. Conversely an equivalence of extensions fixing W and V has exactly this upper triangular form. Thus extension classes are the cokernel, and block addition realizes their vector-space addition. The kernel assertion was already observed. □

Set

$$\langle \alpha, \beta \rangle_Q = \sum_x \alpha_x \beta_x - \sum_{a:x\to y} \alpha_x \beta_y.$$

If this is zero, (G5) is square and its determinant $c^V(W)$ is defined after bases have been chosen. It is well defined up to a nonzero scalar. Equation (G6) proves

$$c^V(W) \neq 0 \quad \Longleftrightarrow \quad \mathrm{Hom}_Q(V,W) = 0 = \mathrm{Ext}^1_Q(V,W). \tag{G7}$$

A zero of this determinant is therefore a concrete failure of exactness, rather than an unexplained spectral analogy.

Let $g = (g_x) \in \prod_x \mathrm{GL}(W_x)$. On C^0 let $T_0\phi = (g_x\phi_x)$, and on C^1 let $T_1\eta = (g_y\eta_a)$. Direct substitution gives $d^V_{gW}T_0 = T_1 d^V_W$. Taking determinants proves

$$c^V(gW) = \prod_x \det(g_x)^{\sum_{a:y\to x}\alpha_y - \alpha_x}\; c^V(W). \tag{G8}$$

The exponent is minus the coefficient of β_x in the Euler form. Under the coordinate-function convention $g \cdot f(W) = f(g^{-1}W)$, the weight changes sign, agreeing with the conventional weight $\langle \alpha, \cdot \rangle_Q$. Stating the action removes an otherwise frequent sign ambiguity.

Derksen–Weyman's spanning theorem says that, for an acyclic quiver, these determinantal constructions span the semi-invariant weight spaces; their work also proves saturation results underlying Littlewood–Richardson multiplicities [44]. Those global theorems are external inputs. The exactness test and character computation needed for our examples have been proved directly.

Example 10.4 (The Kronecker determinant)**.** Fix the representation $V = (1, \lambda)$ of dimension $(1, 1)$, and take $W = (A, B)$ of dimension (n, n). In compatible bases, (G5) is

$$d^V_W = \begin{pmatrix} A & -I_n \\ B & -\lambda I_n \end{pmatrix}, \qquad c^V(W) = \det(B - \lambda A). \tag{G9}$$

For invertible A, block elimination proves the determinant identity. Both sides are polynomials in the entries of A, B, λ, and the invertible-A locus is dense, so the identity holds everywhere. For $n = 1$, the determinant $b - \lambda a$ vanishes exactly when the two nonzero pairs have the corresponding regular parameter. For general n, it detects a nonzero homomorphism from the chosen regular simple and hence an eigenvalue of the pencil, with the qualifications at infinity specified by the other chart.

For a bound quiver algebra, arrow blocks must also satisfy linearized relations, so the unrestricted C^1 above no longer automatically computes extensions. A valid replacement is a projective resolution

$$0 \longrightarrow P_1 \xrightarrow{\delta} P_0 \longrightarrow V \longrightarrow 0$$

when $\mathrm{pd}\, V \leq 1$. Applying $\mathrm{Hom}_A(-, W)$ then gives a two-term map whose kernel and cokernel are $\mathrm{Hom}_A(V, W)$ and $\mathrm{Ext}^1_A(V, W)$. If δ is merely part of a noninjective presentation, its cokernel after applying Hom need not be Ext^1: the next syzygy imposes another condition. The periodic triangle makes that warning testable.

10.5 Schur factors and the semi-invariant matching condition

The standard-monomial filtration assigns to each arrow a a partition $\lambda(a)$, with length bounded by the relevant rank. At a vertex x, each color s contributes a rational irreducible $\mathrm{GL}(\beta_x)$-representation of a dominant integral highest weight $\gamma(x, s)$. Incoming partitions occur dualized and reversed; outgoing partitions occur with the opposite sign convention. The preceding chapter fixes the coordinate action and spells out these weights. A gentle vertex has at most two color factors, which is the decisive simplification.

Proposition 10.5 (The two-factor character test)**.** *Let L_γ, L_ν be irreducible rational representations of GL_d in characteristic zero. Then*

$$(L_\gamma \otimes L_\nu)^{\mathrm{SL}_d} \neq 0 \quad \Longleftrightarrow \quad \gamma_i + \nu_{d+1-i} = c \quad (1 \le i \le d) \tag{G10}$$

for some integer c. When nonzero this space is one dimensional, and GL_d acts on it by $\det^c$.

Proof. The invariant space identifies with

$$\mathrm{Hom}_{\mathrm{SL}_d}(L_\gamma^*, L_\nu).$$

The dual has highest weight $(-\gamma_d, \dots, -\gamma_1)$. Restriction to SL_d is irreducible, and two rational GL_d-highest weights restrict to the same SL_d-highest weight exactly when their difference is an integer multiple of $(1, \dots, 1)$. Schur's lemma now gives the criterion and dimension. The one-dimensional GL_d-subrepresentation has the stated determinant character. □

For a single color at a vertex, the condition reduces to a constant highest-weight tuple. Applying (G10) at every vertex defines an additive semigroup Λ_{SI} of admissible arrow-partition data. The one-dimensional statement applies to each *filtration factor*. Distinct partition data can have the same global character, so it does not assert that all semi-invariant weight spaces are one dimensional.

The strongest step is multiplicative. For an acyclic gentle algebra, a dimension vector β, and a maximal admissible rank sequence r specifying an irreducible component, Carroll–Weyman construct normalized semi-invariants m_λ such that

$$m_\lambda m_\mu = m_{\lambda+\mu}, \qquad k[\mathrm{Rep}(A,\beta)_r]^{\mathrm{SL}(\beta)} \cong k[\Lambda_{\mathrm{SI}}]. \tag{G11}$$

This is their semi-invariant theorem [32, Theorem 6.10]. It is stronger than saying that an associated graded vector space has Schur factors. Their construction uses the exterior diagonal maps splitting a top wedge into complementary wedges at a coupled vertex, then the maps that send wedge basis elements to minors. Straightening produces a distinguished standard multitableau with coefficient one, proving nonvanishing. A compatible column partition in the product yields the exact multiplicative normalization. The semigroup ring in (G11) is thus attached to actual semi-invariants on a specified component.

Subtracting adjacent entries in (G10) gives

$$\gamma_i - \gamma_{i+1} = \nu_{d-i} - \nu_{d+1-i}. \tag{G12}$$

These nonnegative simple-root differences are propagated along a partition-equivalence graph. Its edges express equality of differences either at a shared vertex or along one arrow's incoming and outgoing copies. Components that are paths impose endpoint equations; components that are cycles leave a free nonnegative integer parameter. Writing $u_a = \lambda(a)_{r(a)}$ for the last part, with zero-rank arrows omitted, the remaining semigroup has the form

$$U = \{u \in \mathbb{N}^{Q_1} : \varphi_e(u) = \varphi_{\Theta(e)}(u)\}. \tag{G13}$$

Each endpoint expression is zero, one u_a, or a sum $u_a + u_b$; each variable occurs at most twice. The cycle parameters contribute polynomial variables:

$$k[\Lambda_{\mathrm{SI}}] \cong k[U][z_1, \dots, z_b]. \tag{G14}$$

Here b counts cycles of the partition-equivalence graph. These combinatorial cycles should not be silently identified with band-module isomorphism classes or with handles of a Riemann surface; those objects have different definitions.

A matching graph depicts each equation as a dotted pairing of its two sides and each variable as a solid edge joining its occurrences, or as a solid loop if it occurs only once. Alternating walks account for balanced nonnegative exponent vectors. The reduced-equation condition means that a variable does not occur on *both* sides of the same equation; it permits its coefficient to be zero on both sides. In particular it should not be read as requiring every pair of zero-one coefficients to be unequal. The matching-walk theorem yields generators with coordinate bounds two, and the gentle invariant degree bounds

$$\deg(\text{generators}) \le 2 \sum_{a \in Q_1} \binom{r(a)+1}{2}, \qquad \deg(\text{relations}) \le 8 \sum_{a \in Q_1} \binom{r(a)+1}{2}. \tag{G15}$$

These bounds refer to polynomial degree in arrow matrix entries, not just to the number of named generators in a binomial. General proofs are in Carroll and Weyman [32, Sections 7–9]. The following example supplies a fully proved instance.

10.6 A matching semigroup with two explicit binomial relations

Consider the semigroup $S \subseteq \mathbb{N}^{10}$, with coordinates $(a_1, \dots, a_5, b_1, \dots, b_5)$, defined by

$$\begin{aligned} a_1 + a_2 &= a_4 + a_5, \quad & b_1 + b_2 &= b_4 + b_5, \\ a_2 + a_3 &= b_2 + b_3, \quad & a_3 + a_4 &= b_3 + b_4. \end{aligned} \tag{G16}$$

Use the same letters for independent polynomial variables when writing its monomials. Define

$$\begin{aligned} X_1 &= a_1 a_5, & X_2 &= b_1 b_5, \\ Y_1 &= a_1 a_4 b_4 b_1, & Y_2 &= a_5 a_2 b_2 b_5, \\ Z_1 &= a_2 b_3 a_4, & Z_2 &= b_2 a_3 b_4, \\ B_1 &= a_2 b_2 b_4 a_4, & B_2 &= a_3 b_3. \end{aligned} \tag{G17}$$

These are the alternating-walk generators of the worked example in Carroll and Weyman [32, Example 8.20]. Direct substitution requires a correction to its arXiv-v1 printed relation: the factor multiplying X_1X_2 is B_1, not B_2.

Theorem 10.6 (Corrected explicit presentation)**.** *The semigroup S is generated by the eight exponent vectors in (G17), and*

$$k[S] \cong \frac{k[X_1, X_2, Y_1, Y_2, Z_1, Z_2, B_1, B_2]}{(Y_1Y_2 - B_1X_1X_2,\ Z_1Z_2 - B_1B_2)}. \tag{G18}$$

It is a normal affine domain of dimension six.

Proof. First take an arbitrary element of S. The equations imply

$$d := a_1 - a_5 = b_1 - b_5 = a_4 - a_2 = b_4 - b_2.$$

If $d \ge 0$, subtract d copies of the exponent vector of Y_1; all four affected coordinates are at least d. If $d < 0$, subtract $-d$ copies of Y_2. The residual coordinates satisfy $a_1 = a_5$, $b_1 = b_5$, $a_2 = a_4$, and $b_2 = b_4$. Remove the remaining a_1 copies of X_1 and b_1 copies of X_2.

Put $A = a_2 = a_4$, $B = b_2 = b_4$, $C = a_3$, $D = b_3$ in what remains. The single condition is $A + C = B + D$. Subtract $\min(A, B)$ copies of B_1. If A remains positive, $B = 0$ and $D = A + C$, so subtract A copies of Z_1, followed by C copies of B_2. If B remains positive, the symmetric procedure

uses B copies of Z_2 and then D copies of B_2. If both are zero, use only B_2. This proves generation with nonnegative coefficients.

The two displayed binomials vanish by multiplication of the ten underlying variables. To prove there are no additional defining relations, compute the integer kernel of the ten-by-eight exponent matrix. If the coefficients of a kernel vector in the displayed generator order are $c_1, c_2, d_1, d_2, e_1, e_2, f_1, f_2$, the a_1, a_5, b_1, b_5 coordinates give

$$d_1 = d_2 = p, \qquad c_1 = c_2 = -p.$$

The a_3, b_3 coordinates give $e_1 = e_2 = q$, $f_2 = -q$, and the remaining coordinates give $f_1 = -p - q$. Thus every integer relation is

$$p(-1, -1, 1, 1, 0, 0, -1, 0) + q(0, 0, 0, 0, 1, 1, -1, -1).$$

Reduce monomials using $Y_1Y_2 \mapsto B_1X_1X_2$ and $Z_1Z_2 \mapsto B_1B_2$. The number of Y- and Z-factors strictly decreases, so reduction terminates. A reduced monomial contains neither both Y_1, Y_2 nor both Z_1, Z_2. If two reduced monomials had the same underlying exponent vector, the kernel formula would apply. A nonzero p would force one of them to contain both Y's after common factors are canceled, and a nonzero q would force both Z's. Hence $p = q = 0$, and the reduced monomials are identical. Their images are therefore linearly independent. This proves the quotient presentation for arbitrary polynomials as well as for binomials.

The exponent matrix has rank six, since its integer kernel has rank two. The group algebra of its exponent lattice is a Laurent polynomial ring in six independent variables, so $k[S]$ has dimension six. To check normality, write $L = \mathrm{gp}(S) \subseteq \mathbb{Z}^{10}$. From (G16),

$$k[S] = k[L] \cap k[a_1, \dots, a_5, b_1, \dots, b_5]$$

inside the ambient Laurent polynomial ring: an exponent in L with all coordinates nonnegative satisfies (G16), hence lies in S. Both rings in the intersection are integrally closed. An element of the fraction field of $k[S]$ integral over the intersection is integral over both rings and belongs to both fraction fields, hence belongs to the intersection. This proves normality. The monomial embedding proves that the ring is a domain. □

The first relation has degree eight in the underlying arrow-style variables, because $\deg Y_i = 4$, $\deg B_1 = 4$, and $\deg X_i = 2$. The second has degree six, since $\deg Z_i = 3$ and $\deg B_2 = 2$. Treating all eight symbols as degree one would destroy this grading and misstate the degree information. The companion verification checks every solution of (G16) whose ten coordinates lie in $\{0, 1, 2\}$ by the constructive decomposition above, and separately verifies the two exponent identities. The finite check is additional evidence; the preceding proof establishes the unbounded assertion.

10.7 Quotients and the information retained by semi-invariants

Carroll–Chindris prove, for an acyclic gentle algebra over an algebraically closed field of characteristic zero, that the rational invariant field on each irreducible module-variety component is purely transcendental. Its transcendence degree is the sum of the multiplicities of the indecomposable regular components in the generic decomposition. For a regular irreducible component and a weight with nonempty semistable locus whose stable summands are regular, the corresponding moduli space is a product of projective spaces [31, Theorem 1]. The qualifications about the component and stable summands are part of the theorem. They do not assert that every module variety, every affine invariant ring, or every stability quotient is a projective-space product.

Ratios of generalized Schofield determinants give their explicit rational coordinates. On the elementary Kronecker chart, this becomes familiar: the coefficients of $\det(B - \lambda A)$, normalized by a nonzero coefficient, recover the symmetric functions of the generic regular parameters. Repeated eigenvalues and nonsplit extensions explain why these rational coordinates describe a generic quotient rather than a complete isomorphism classification at all points.

There are consequently three levels of geometry useful for this project. A representation variety remembers actual matrices satisfying relations. A semi-invariant ring remembers prescribed determinant characters and defines graded quotients. A rational invariant field describes a generic orbit quotient. The associated graded algebra of a straightening filtration adds another construction with its own specialization map. Passing between these objects is justified by maps and hypotheses, never by the fact that all four admit combinatorial descriptions.

Chapter 11

Exterior traces, zeta functions and the spectral comparison problem

Evidence status. The finite-dimensional determinant and trace identities are proved here. The cohomological description of elliptic-curve zeta functions and the Weil bound are external theorems. The transition to completed arithmetic L-functions is formulated as an obligation of the research program.

11.1 Exterior powers supply determinant coefficients

Let $T : V \to V$ be an endomorphism of an n-dimensional vector space over a characteristic-zero field. Its action on $\bigwedge^j V$ is the special Schur functor associated to a column of height j.

Theorem 11.1 (Exterior determinant identity)**.**

$$\det(I - tT) = \sum_{j=0}^{n} (-t)^j \operatorname{Tr}(\bigwedge^j T). \tag{Z1}$$

Proof. For a diagonal matrix with eigenvalues $\lambda_1, \ldots, \lambda_n$, the wedge basis has eigenvalues $\lambda_{i_1} \cdots \lambda_{i_j}$. Its trace is therefore the j-th elementary symmetric polynomial, which is exactly the coefficient in $\prod_i (1 - t\lambda_i)$. For an arbitrary matrix, extend to an algebraic closure. The identity holds on matrices with distinct eigenvalues, a Zariski-dense open set. Each coefficient on both sides is a polynomial in the matrix entries, so the identity holds identically and descends to the original field. □

In particular the minors that generate standard monomials also compute the exterior traces in a characteristic polynomial. The scalar determinant is the top exterior character. General Schur functors organize more elaborate polynomial characters, but the existence of those characters by itself places no sign or location restriction on their roots.

For a T-invariant short exact sequence $0 \to U \to V \to W \to 0$, a basis adapted to U makes T block upper triangular, proving

$$\det(I - tT_V) = \det(I - tT_U) \det(I - tT_W). \tag{Z2}$$

No equivariant splitting is needed. In the Grothendieck group of finite-dimensional vector spaces equipped with an endomorphism, with relations from equivariant short exact sequences, this becomes

multiplicativity on virtual differences:

$$D_{[V]-[W]}(t) = \frac{\det(I - tT_V)}{\det(I - tT_W)}. \tag{Z3}$$

The denominator must be retained. An alternating determinant of a complex is naturally a rational function before additional cancellations are established.

Proposition 11.2 (Cancellation through cohomology)**.** *Let $(C^\bullet, d)$ be a finite complex of finite-dimensional vector spaces and let T be a chain endomorphism. Then*

$$\prod_i \det(I - tT_{C^i})^{(-1)^i} = \prod_i \det(I - tT_{H^i})^{(-1)^i}. \tag{Z4}$$

Proof. Let $Z^i = \ker d^i$ and $B^i = \operatorname{im} d^{i-1}$. Both are T-stable. Apply (Z2) to $0 \to Z^i \to C^i \to B^{i+1} \to 0$ and $0 \to B^i \to Z^i \to H^i \to 0$. Each boundary determinant appears once with exponent $+1$ and once with exponent -1 in the alternating product, leaving exactly the cohomology factors. □

This is the elementary determinant counterpart of the Euler characteristic identity. It states exactly what finite-complex exactness can do for a determinant: it cancels acyclic factors. It does not supply analytic continuation, an arithmetic Euler product or a positive inner product.

11.2 Traces, closed walks and a finite zeta model

In the formal power-series ring $k[[t]]$, where characteristic zero allows division by positive integers,

$$-\log\det(I - tT) = \sum_{m\geq 1} \frac{\operatorname{Tr}(T^m)}{m} t^m. \tag{Z5}$$

To prove this, first diagonalize on the dense open set and use $-\log(1 - t\lambda) = \sum_{m\geq 1}(t\lambda)^m/m$. Each coefficient is a polynomial identity, so the same density argument proves it for all matrices. Alternatively differentiate the logarithm formally and expand $(I - tT)^{-1}$, then integrate with constant term zero. Over $\mathbb{C}$, the same expression converges analytically for sufficiently small $|t|$, for example when $|t|\,\|T\| < 1$.

For a finite directed graph with adjacency matrix B, the entry $(B^m)_{vv}$ counts length-m walks from v to itself, and the trace counts closed walks with a marked starting point. Thus

$$Z_B(t) := \exp\left(\sum_{m\geq 1} \frac{\operatorname{Tr}(B^m)}{m} t^m\right) = \det(I - tB)^{-1}. \tag{Z6}$$

An Euler-product version follows by decomposing a closed walk into repetitions of a primitive cyclic word. A primitive cycle of length ℓ has ℓ marked starts, and its r-fold repetition contributes $\ell/(r\ell) = 1/r$, yielding the factor $(1 - t^\ell)^{-1}$. These identities are formal: every coefficient receives finitely many contributions.

For a ribbon graph or dessin, choose B to be the nonbacktracking transition matrix on oriented edges if that is the walk model of interest. Its entries are one when two directed edges concatenate without immediately reversing, and zero otherwise. Formula (Z6) then defines the corresponding finite graph zeta function directly. The rotation system defining the Riemann surface and the transition rule defining B are separate inputs. Changing the former need not change the ordinary underlying graph, while changing the latter changes the trace counts. The interactive dessin therefore offers a precise combinatorial laboratory without identifying this finite rational function with a completed number-field L-function.

11.3 The elliptic-curve numerator as a cohomological determinant

Let $E/\mathbb{F}_q$ be a smooth projective elliptic curve. With $N_m = \#E(\mathbb{F}_{q^m})$, define

$$Z(E,t) = \exp\left(\sum_{m\geq 1} \frac{N_m}{m} t^m\right).$$

The finite-field cohomological theorem and the Weil bound give

$$Z(E,t) = \frac{1 - at + qt^2}{(1-t)(1-qt)}, \qquad a = q + 1 - N_1, \quad |a| \leq 2\sqrt{q}. \tag{Z7}$$

These are deep external facts; the elementary determinant identities above do not prove them. The endomorphism and Frobenius framework is described by Tate and the standard elliptic-curve sources already used in this monograph [152, 159, 164].

In a Frobenius convention compatible with $N_m = 1 + q^m - \mathrm{Tr}(F^m|H^1)$, the numerator is $\det(I - tF|H^1)$, the two denominator factors come from H^0 and H^2, and (Z5) recovers the point-counting exponential. Equivalently, the roots α, β of $X^2 - aX + q$ satisfy

$$\alpha + \beta = a, \quad \alpha\beta = q, \quad N_m = q^m + 1 - \alpha^m - \beta^m.$$

The zeros in the t-plane are α^{-1}, β^{-1}. Confusing eigenvalues with reciprocal zeros would reverse the radius.

For the companion curve $y^2 = x^3 + 2x + 2$ over $\mathbb{F}_{17}$, exact counting gives $N_1 = 19$, so $a = -1$. Its numerator is

$$1 + t + 17t^2,$$

whose zeros are $(-1 \pm i\sqrt{67})/34$ and have modulus $1/\sqrt{17}$. The recurrence gives $\alpha^2 + \beta^2 = a^2 - 2q = -33$, hence $N_2 = 323$. The code independently counts the points over $\mathbb{F}_{17^2} = \mathbb{F}_{17}[u]/(u^2 - 3)$ and verifies this value. This finite-field example directly connects the order generated by Frobenius, exterior traces, point counts, the ECC group and a zeta numerator. It does not turn the finite-field Weil theorem into the generalized Riemann hypothesis for number-field L-functions.

11.4 What standard monomials can contribute to the arithmetic program

The constructions now align in a useful sequence. A complex or quiver presentation specifies relations among maps. Exterior products provide minors, determinant characters and straightening generators. Schur filtrations keep track of the representation content of coordinate functions. At gentle vertices, complementary highest weights identify semi-invariants. Matching semigroups make a class of their rings and relations computable. Determinants of appropriate homomorphism complexes detect the failure of exactness. Exterior traces then turn an actual endomorphism into characteristic coefficients and a finite trace generating function.

Each step has a definite object and a proved or cited theorem. None yet identifies the result with an arithmetic L-function. To make that final identification, the inherited spectral program must still specify the arithmetic operator or complex, the exact Euler factors, local conductor and archimedean normalizations, the treatment of poles and zero multiplicities, and the analytic domain in which determinant identities are meaningful. In infinite dimension, an ordinary determinant usually does not exist; trace-class Fredholm determinants, zeta regularization and relative determinants have

different hypotheses. A filtration identity or a formal power series at one point cannot silently supply equality after analytic continuation.

There is an elementary obstruction to any claim that gentle geometry alone forces the desired zero set. In the Kronecker example, take $A = I_n$, $B = \lambda I_n$ for an arbitrary complex λ. All representation-theoretic relations remain valid, but $\det(B - zA) = (\lambda - z)^n$ has its root at any chosen λ. If one imposes $B = B^*$ in a specified positive Hermitian metric, the roots become real by the spectral theorem. If a generalized pencil is $K - zM$ with $M > 0$ and $K = K^*$, it reduces by $M^{-1/2}$ to a self-adjoint eigenvalue problem. Positivity is additional structure on the representation, not a consequence of its Schur decomposition. The earlier positive-pencil implications state the corresponding conditional arithmetic conclusion precisely.

The same distinction controls the Yang–Mills comparison. The deformation complex and its determinant line describe zero modes, orientations and local moduli. A mass gap requires an inequality on the quantum Hamiltonian after the continuum and infinite-volume constructions. Standard monomial bases can simplify finite algebraic sectors and organize multiplicities, but they do not create a cutoff-uniform coercive estimate. The surface-code examples make the finite version visible: the binary chain condition proves stabilizer commutation, while code distance depends on the shortest nontrivial homology representatives, and a decoder's performance depends on a noise model and algorithm. These are compatible layers of one construction, with different proof obligations.

11.5 A concrete program for further formalization

The next formalization tasks can therefore be specified without suggesting that the cited theories are already encoded. One tractable target is the explicit monoid in (G16): define its ten integer coordinates, encode the constructive generator decomposition, and prove the two-binomial normal form. A second is the functorial determinant identity (Z4) for finite complexes and commuting chain maps. A third is the Schofield transformation law (G8), with the action convention in the type signature. These would extend the current verified chain-map/Pauli bridge with actual exterior and invariant-theory interfaces.

A much larger task is a reusable standard-monomial formalization: partitions and tableau orders, wedge shuffle maps over a stated base ring, a termination measure for straightening, standard-basis independence, the resulting filtration, and its equivariant Schur factors. In positive characteristic the distinction between a filtration and a semisimple decomposition must remain explicit. Formalizing only the rank identities or declaring the straightening theorem as a new axiom would not accomplish that task. The present repository supplies mathematical specifications and independently checked examples for it; its current verified Lean inventory is the 121 selected declarations compiled in the seven-module package.

Part III

Cryptography and quantum algorithms

Chapter 12

Elliptic-curve arithmetic and quantum period finding

12.1 The group problem and its arithmetic representation

Let $p > 3$ be prime and consider the nonsingular projective completion of

$$E: \quad y^2 = x^3 + ax + b, \qquad 4a^3 + 27b^2 \neq 0 \pmod{p}.$$

The point at infinity is denoted by O. The group law on $E(\mathbb{F}_p)$ is the usual algebraic group structure of an elliptic curve. Its existence and associativity can be established through divisor classes, identifying P with $[P - O] \in \mathrm{Pic}^0(E)$; this is a standard external theorem of elliptic-curve theory, not a consequence of a drawing of chords [152]. For computations, if $P = (x_1, y_1)$ and $Q = (x_2, y_2)$ have distinct x coordinates, set

$$m = (y_2 - y_1)(x_2 - x_1)^{-1}, \quad x_3 = m^2 - x_1 - x_2, \quad y_3 = m(x_1 - x_3) - y_1.$$

Then $P + Q = (x_3, y_3)$. For doubling, replace the slope by $m = (3x_1^2 + a)(2y_1)^{-1}$ when $y_1 \neq 0$. The exceptional cases are $P + O = P$, $P + (-P) = O$, and $2P = O$ when $y_1 = 0$. All inverses are in $\mathbb{F}_p$.

These formulas follow by substituting the line through the points into the cubic. The resulting monic cubic in x has roots x_1, x_2, x_3 and coefficient of x^2 equal to $-m^2$, so their sum is m^2. Reflection in the x axis gives the group-law output rather than the third point of intersection. Tangency replaces the two distinct roots by a repeated root. This derivation explains the formula, while the divisor-class construction supplies associativity.

An implementation must distinguish formulas valid on a dense open subset from a total group operation. A denominator that is generically nonzero can vanish on exceptional points. A reversible quantum circuit must implement a bijection on its entire represented state space, or provide a precise bound for any omitted amplitudes. Calling the exceptional cases "rare" without a probability or operator-norm analysis is not a correctness proof. This distinction has been central to the evolution from early resource analyses to more explicit quantum point-addition constructions [91, 123, 130].

Fix $P \in E(\mathbb{F}_p)$ of prime order q and $Q \in \langle P \rangle$. The elliptic-curve discrete-logarithm problem is to find $d \in \mathbb{Z}/q\mathbb{Z}$ such that $Q = dP$. The parameter controlling generic group difficulty is the subgroup order q, not automatically the field size p. For a cryptographic subgroup these sizes are usually comparable in logarithmic scale, but the mathematical statements should preserve their different roles.

Example 12.1 (A completely reproducible finite curve)**.** Take $p = 17$, $a = b = 2$ and $P = (5, 1)$. The discriminant condition holds because $4 \cdot 2^3 + 27 \cdot 2^2 = 140 \equiv 4 \pmod{17}$. The companion enumerates every affine solution, adds O, and checks that the group has 19 elements and that P generates it. For instance,

$$2P = (6, 3), \qquad 3P = (10, 6), \qquad 7P = (0, 6).$$

Every pairwise addition and every associativity triple in this finite model is checked. These exhaustive computations certify this small example, not an asymptotic complexity theorem or a deployed cryptographic implementation.

12.2 The relation subgroup

The structural step in Shor's discrete-logarithm algorithm is to replace the unknown scalar by a hidden relation. Define

$$f : (\mathbb{Z}/q\mathbb{Z})^2 \longrightarrow \langle P \rangle, \qquad f(a, b) = aP + bQ.$$

Writing $Q = dP$ gives $f(a, b) = (a + db)P$.

Proposition 12.2 (Kernel and cosets)**.** *The kernel of f is*

$$H_d = \{(-dt, t) : t \in \mathbb{Z}/q\mathbb{Z}\}.$$

Two inputs have the same image exactly when they lie in the same coset of H_d. The subgroup has q elements and index q.

Proof. Since P has order q, $f(a, b) = O$ if and only if $a + db = 0$ in $\mathbb{Z}/q\mathbb{Z}$. This is equivalent to $(a, b) = (-db, b)$, proving the kernel formula. For any homomorphism, $f(x) = f(y)$ is equivalent to $f(x - y) = O$, hence to $x - y \in H_d$. The parametrization by t is injective because its second coordinate is t, so $|H_d| = q$. The ambient group has q^2 elements, giving the index. □

The kernel is an algebraic object, while the algorithm also needs an efficient reversible implementation of f. Knowing that a homomorphism exists is not enough. Its output representation, the point arithmetic, its workspace and the cost of uncomputing ancillary registers all enter the complexity analysis. The reduction to a subgroup is general; the cost of realizing it depends on the group presentation [123, 151].

12.3 Fourier sampling with every normalization explicit

Write $G = (\mathbb{Z}/q\mathbb{Z})^2$, and let $\zeta = \exp(2\pi i/q)$. Its characters are

$$\chi_{u,v}(a, b) = \zeta^{ua+vb}, \qquad (u, v) \in G.$$

Use the Fourier transform convention

$$\mathcal{F}|a, b\rangle = \frac{1}{q} \sum_{u,v \in \mathbb{Z}/q\mathbb{Z}} \zeta^{ua+vb} |u, v\rangle.$$

The factor $1/q$ is $1/\sqrt{|G|}$. A normalized coset state is

$$|x + H_d\rangle = \frac{1}{\sqrt{q}} \sum_{t \in \mathbb{Z}/q\mathbb{Z}} |x + (-dt, t)\rangle.$$

Such a state can be obtained by evaluating f coherently on a uniform input and measuring the output. Discarding the output gives the corresponding mixture over cosets, which has the same Fourier outcome probabilities.

Lemma 12.3 (Character cancellation)**.** *If $z \in \mathbb{C}$ satisfies $z^q = 1$, then*

$$\sum_{t=0}^{q-1} z^t = \begin{cases} q, & z = 1, \\ 0, & z \neq 1. \end{cases}$$

Proof. The first case is a sum of q ones. In the second case, multiplying the sum by $1 - z$ telescopes to $1 - z^q = 0$. Since $1 - z \neq 0$ in the field $\mathbb{C}$, the sum is zero. □

Theorem 12.4 (Exact annihilator distribution)**.** *Measuring $\mathcal{F}|x+H_d\rangle$ gives the uniform distribution on*

$$H_d^{\perp} = \{(u, v) : v = du \pmod q\}.$$

If q is prime, any observed pair with $u \neq 0$ determines $d = vu^{-1} \pmod q$. The probability of this event in one ideal sample is $1 - 1/q$.

Proof. The amplitude of (u, v) is

$$\frac{\chi_{u,v}(x)}{q\sqrt{q}} \sum_{t=0}^{q-1} \zeta^{t(-ud+v)}.$$

The lemma makes it zero unless $v = du$. If equality holds, the amplitude is $\chi_{u,v}(x)/\sqrt{q}$, whose squared absolute value is $1/q$. There are exactly q such pairs, one for each u, so the probabilities sum to one. A nonzero u is invertible in $\mathbb{F}_q$. Exactly one of the q possible values is zero, giving the success probability. Finally, the candidate can be verified by checking $Q = dP$. □

This proof gives the exact finite Fourier mechanism. A practical circuit often uses transforms over power-of-two register sizes rather than exactly over $\mathbb{Z}/q\mathbb{Z}$. Its rounding, distributional approximation and classical reconstruction require additional arguments; they are addressed in the original algorithmic literature [123, 151]. The exact model in the companion deliberately isolates the clean mathematical kernel from those implementation details.

Example 12.5 (The seven-point frequency line)**.** For $q = 7$ and $d = 3$, the outcomes are

$$(0, 0), (1, 3), (2, 6), (3, 2), (4, 5), (5, 1), (6, 4),$$

each with probability $1/7$. The picture is a line over a finite field, not a straight line in a continuous Euclidean square. Connecting adjacent displayed points by a real segment is a visual convention. The companion instead labels modular coordinates and shows the probability at each actual discrete pair.

12.4 What an endomorphism order contributes

Suppose an explicitly computable endomorphism φ preserves $\langle P\rangle$ and acts there as multiplication by a known $\lambda \in \mathbb{Z}/q\mathbb{Z}$. Then

$$(a + b\lambda)P = aP + b\varphi(P).$$

The integer relation lattice

$$L_\lambda = \{(a, b) \in \mathbb{Z}^2 : a + b\lambda \equiv 0 \pmod q\}$$

has an elementary integral description. Choose an integer representative of λ.

Proposition 12.6 (Relation lattice and covolume)**.** *The vectors* $(q, 0)$ *and* $(-\lambda, 1)$ *form a* $\mathbb{Z}$*-basis of* L_λ*. Its index in* $\mathbb{Z}^2$*, and hence its Euclidean covolume in the standard metric, is* q*.*

Proof. Both displayed vectors satisfy the relation. If $(a, b) \in L_\lambda$, there is an integer k with $a+b\lambda = kq$. Hence $(a, b) = k(q, 0) + b(-\lambda, 1)$. The two vectors are independent, and the determinant of their column matrix is q. The determinant-index formula gives the index and covolume. Alternatively the homomorphism $(a, b) \mapsto a + b\lambda \pmod q$ is surjective and has this kernel, so the quotient has q elements. □

The relation lattice supports short scalar decompositions and structured arithmetic. It does not reveal the unknown scalar relating an arbitrary Q to P. Indeed, the lattice depends on the known λ and q, not on that unknown scalar. The potential saving is in the implementation of group operations or scalar multiplication; a theorem about the endomorphism ring is not itself a theorem about the number of Toffoli gates.

For a proposed quantum improvement, the relevant comparison is explicit. If a baseline group evaluation uses N_a additions and N_d doublings, and the revised evaluation uses N'_a, N'_d plus N_φ endomorphism applications, a necessary gate-budget comparison is

$$N'_a C_a + N'_d C_d + N_\varphi C_\varphi + C_{\text{conversion}} + C_{\text{cleanup}} < N_a C_a + N_d C_d.$$

The constants refer to the same gate set and error model. The corresponding live-qubit and depth inequalities must also be checked. Removing an arithmetic operation while adding a costly coherent lookup or inversion can increase the total cost. The order-theoretic view is useful because it provides a principled source of candidate decompositions; the circuit evidence still decides whether they help.

12.5 A faithful comparison of recent algorithmic developments

Shor's theorem establishes polynomial-time quantum algorithms for factoring and discrete logarithms with efficient group operations. Regev's multidimensional construction instead organizes modular products through an integer relation lattice and reduces the size of an individual coherent circuit under an arithmetic hypothesis in its original version. It uses roughly $\sqrt{n}$ independently repeated circuits for an n-bit input, followed by classical lattice processing [129, 151]. Thus "smaller circuit" and "smaller total gate count" are distinct claims.

Pilatte's 2026 result removes the relevant conjectural dependence for modified factoring and multiplicative discrete-logarithm algorithms. In the formulation of Theorems 1.1 and 1.2, the individual circuit has

$$O(n^{3/2} \log^3 n) \text{ gates}, \qquad O(n \log^3 n) \text{ qubits},$$

and is called $O(\sqrt{n})$ times with classical randomized polynomial-time processing and constant success probability. The discrete-logarithm theorem is stated for $g, y \in (\mathbb{Z}/N\mathbb{Z})^\times$, with $y \in \langle g \rangle$. This group and representation hypothesis must be retained [121]. Multiplying the number of calls by the per-call gate count gives an $O(n^2 \log^3 n)$ upper comparison for total gates at this level of abstraction. The coherent circuit size improvement can still matter physically, but it should not be described as a universal $n^{3/2}$ total-time algorithm.

The reason arithmetic enters is especially instructive. A relation lattice built from small integers can have very different geometry from one built from arbitrary residues. Pilatte proves the necessary short-basis statements using analytic number theory, including zero-density estimates, and modifies the generator selection to handle subgroup obstructions. An unconditional theorem for this setting is not an unconditional theorem about every elliptic-curve lift. It is precisely the kind of example in

which analytic number theory supplies a computationally meaningful bridge, with its domain stated correctly.

Barbulescu, Bârcău and Paşol extend the multidimensional strategy to elliptic curves through small-height rational data. Their published account explicitly retains number-theoretic conjectural hypotheses for the claimed asymptotic improvement. The primary publisher record and author conference presentation were inspected; the complete subscription chapter was not independently proof-audited in this review [20, 21]. A legitimate research direction is to identify concrete curve families and height bounds meeting those hypotheses. Merely replacing multiplication by elliptic addition in a pseudocode template does not prove the same cost bound.

Result	Precise contribution used here	Boundary retained
Shor; Proos–Zalka	Period finding and elliptic-curve group-operation realization	Fault-tolerant quantum computation and explicit arithmetic cost
Roetteler et al.	Reversible network resource estimates over prime fields	A specific circuit family and gate model
Regev	Multidimensional relation-lattice sampling and shorter coherent circuits	Original number-theoretic hypothesis; repeated calls
Ekerå–Gärtner	Extensions to logarithms, orders and related multiplicative tasks	Representation and classical processing assumptions
Pilatte	Unconditional modified factoring and multiplicative-DLP guarantees	The theorem's stated group, generator distribution and logarithmic factors
Barbulescu et al.	Elliptic-curve extension using arithmetic small-height structure	Its additional number-theoretic conjectures
Luo et al., v3	Explicit space-efficient inversion and point-addition construction	Logical resources; current version and lower-order terms

The 2017 Roetteler et al. bound gives at most $9n+2\lceil \log_2 n\rceil +10$ logical qubits and $448n^3 \log_2 n + 4090n^3$ Toffoli gates for the specified construction [130]. Luo et al.'s September 2026 version reports the leading space expression $3n + 6\lfloor \log_2 n\rfloor + O(1)$, a concrete 835-logical-qubit estimate at 256 bits, and leading Toffoli expression $1056n^3/\log_2 n + O(n^2)$ [91]. The version pin matters: earlier metadata gave different coefficients. The asymptotic $O(1)$ and $O(n^2)$ terms must not be silently set to zero when comparing concrete instances.

The 2026 PRX Quantum resource study by Babbush and collaborators reports different low-gate tradeoffs and supplies a zero-knowledge certificate for committed, undisclosed circuit details [18]. The certificate checks resource bounds and 9,024 pseudorandom point-addition cases generated through a Fiat–Shamir-style hash procedure; it is not a universal correctness proof on all computational basis states. Its false-pass estimate is conditional on the stated sampling and ideal-hash assumptions. This is relevant published evidence, but it is a different reproducibility category from an openly inspectable circuit. Its architectural runtime estimates additionally assume an error-correction architecture and physical error rates. Logical-qubit counts, physical-qubit counts and elapsed wall-clock predictions are not interchangeable units.

12.6 What could count as a new improvement

A rigorous new result could take several forms. It could reduce a proved gate count at equal live-qubit and error budgets. It could reduce qubits at a specified gate penalty. It could prove a family of elliptic curves satisfies a previously conditional small-height hypothesis. It could improve classical lattice reconstruction with a complete success analysis. It could formalize a previously unaudited arithmetic subroutine and expose a correctness defect in a cost estimate.

This monograph establishes none of these as a new asymptotic cryptanalytic breakthrough. It proposes a route to investigate them: use known order actions to enumerate algebraically valid transformations, compile each into a reversible representation, prove its functional equivalence, then compare costs with an identical accounting method. The selected Lean lemmas and exhaustive small models provide reusable interfaces for that route. They do not replace a complete algorithm or its complexity analysis.

Chapter 13

Order lattices, noise and post-quantum cryptography

13.1 From a module to a geometry

Let $\Lambda \subset \mathbb{R}^n$ be a Euclidean lattice with basis matrix $B \in \mathrm{GL}_n(\mathbb{R})$, so $\Lambda = B\mathbb{Z}^n$. Its covolume is $|\det B|$, its dual is

$$\Lambda^* = \{y \in \mathbb{R}^n : \langle y, x\rangle \in \mathbb{Z} \text{ for every } x \in \Lambda\} = B^{-T}\mathbb{Z}^n,$$

and, for $n \geq 1$, its first minimum is $\lambda_1(\Lambda) = \min_{0 \neq x \in \Lambda} \|x\|_2$. The minimum exists because a lattice has finitely many points in every bounded region and has no accumulation at zero.

Proposition 13.1 (Dual-basis formula)**.** *For an invertible real matrix B, the dual of $B\mathbb{Z}^n$ is $B^{-T}\mathbb{Z}^n$.*

Proof. A vector y belongs to the dual exactly when $y^T Bz \in \mathbb{Z}$ for every integer vector z. Testing the standard basis vectors shows that this is equivalent to $B^T y \in \mathbb{Z}^n$. Solving for y gives the formula. It also yields $\det(\Lambda^*) = 1/\det(\Lambda)$ for covolumes. □

An arithmetic order $\mathcal{O}$ in a number field K has finite rank as a $\mathbb{Z}$-module. The archimedean embedding sends a fractional ideal or an order-module into a real vector space. Real embeddings contribute one coordinate and a complex pair contributes two real coordinates, conventionally scaled by $\sqrt{2}$ when using the standard trace-compatible Euclidean norm. The chosen metric is part of the cryptographic instance. A coefficient embedding and the canonical embedding are related by a real linear map, whose condition number controls the distortion of lengths and error distributions [93].

For a fractional ideal I of the maximal order of a number field, the trace dual is

$$I^\vee = \{x \in K : \mathrm{Tr}_{K/\mathbb{Q}}(xI) \subset \mathbb{Z}\}.$$

When I is invertible, this equals $\mathfrak{D}^{-1}I^{-1}$, where $\mathfrak{D}^{-1} = \mathcal{O}_K^\vee$ is the codifferent. To see the identity, the trace condition says exactly that $xI \subset \mathfrak{D}^{-1}$; multiplication by I^{-1} gives the result. For noninvertible ideals in a nonmaximal order, one should use the ideal quotient rather than assume the same inverse formula. This is one place where the book's local order theory directly prevents a cryptographic algebra mistake.

At complex places one must distinguish the trace pairing from the positive Euclidean pairing. Write $j(x)$ using real embeddings and the coordinates $(\sqrt{2}\,\mathrm{Re}\,\tau(x), \sqrt{2}\,\mathrm{Im}\,\tau(x))$ for each chosen

complex embedding. Let J negate every imaginary coordinate and fix the other coordinates. Direct multiplication of complex numbers gives $\langle j(x), Jj(y)\rangle = \mathrm{Tr}_{K/\mathbb{Q}}(xy)$. Consequently

$$j(I)^* = Jj(I^\vee).$$

Indeed $j(I^\vee)$ is exactly the lattice of vectors whose pairing through J with $j(I)$ is integral. The involution J acts on the real coordinate space; it need not come from an automorphism of an arbitrary number field. Thus trace duality and Euclidean duality should not be identified without this correction.

13.2 An algebraic equivalence need not preserve short vectors

Proposition 13.2 (Unimodular changes can have unbounded metric distortion)**.** *For each positive integer N, the matrix*

$$U_N = \begin{pmatrix} 1 & N \\ 0 & 1 \end{pmatrix}$$

is an automorphism of the abelian group $\mathbb{Z}^2$ and has determinant one. Nevertheless $\|U_N(0,1)^T\|_2 = \sqrt{N^2+1}$, so its operator norm is unbounded as $N \to \infty$.

Proof. The inverse is $\left(\begin{smallmatrix} 1 & -N \\ 0 & 1 \end{smallmatrix}\right)$, which also has integer entries. Thus the map and its inverse preserve the integer lattice as a set. Its value on the unit vector $(0,1)$ is $(N,1)$, giving the norm formula and the lower bound on the operator norm. □

The example has two complementary readings. If U_N is regarded as a change of basis of the same Euclidean lattice, the lattice itself has not changed, but the coordinate distribution described as "independent small coefficients" changes its geometric meaning. If it is regarded as a linear map between two metric module presentations, it is not an isometry. In either reading, categorical equivalence alone does not transfer the sampling or short-vector problem with its parameters intact.

For a random vector e with mean zero and covariance Σ, the transformed vector has covariance

$$\mathrm{Cov}(Te) = T\Sigma T^T.$$

This follows immediately by expanding $\mathbb{E}[(Te)(Te)^T] = T\mathbb{E}[ee^T]T^T$. For a square invertible T and spherical covariance $\sigma^2 I$, it becomes $\sigma^2 TT^T$, which, for $\sigma > 0$, is spherical exactly when $TT^T = cI$ for some $c > 0$, equivalently when T is a scaled orthogonal transformation. This exact calculation is a basic obligation in any proposed transfer of LWE to a new order or Morita-equivalent module category.

13.3 Noisy linear equations

For a modulus q, a dimension n, a secret $s \in (\mathbb{Z}/q\mathbb{Z})^n$ and a chosen error law χ, a learning-with-errors sample has the form

$$a \leftarrow (\mathbb{Z}/q\mathbb{Z})^n, \qquad e \leftarrow \chi, \qquad b = \langle a, s\rangle + e \pmod q.$$

The search task asks for s from many samples; the decisional task asks to distinguish this law from uniform samples. The error law, number of samples, modulus, dimension and distribution of the secret are part of the problem, not optional details. An assertion that "LWE is hard" suppresses parameters that a precise reduction must state.

Regev's theorem gives a quantum reduction from appropriate worst-case lattice problems to specified LWE instances. In a standard formulation, for a suitable modulus and Gaussian error family with $\alpha q > 2\sqrt{n}$, an efficient LWE solver yields a quantum algorithm for worst-case GapSVP and SIVP at approximation factor $\widetilde{O}(n/\alpha)$, with the exact parameter restrictions supplied by the original theorem [128]. This is a conditional hardness implication. It is not a proof that no quantum algorithm can solve those lattice problems, and it is not a blanket theorem for every small-error distribution used in every implementation.

Ring-LWE packages several coordinates into algebraic multiplication in an order, and module-LWE uses vectors of ring elements. The foundational ring-LWE reduction uses ideal lattices and a carefully specified canonical-embedding error family [93]. A noncommutative order can likewise support noisy module equations as a definition, but defining the equations is easier than proving a secure distribution, a worst-case reduction or useful concrete parameters. Noncommutativity should be treated as extra structure to analyze, not as an automatic source of security.

13.4 A complete elementary correctness calculation

The following is an educational linear-encryption skeleton. It is not a complete KEM and does not claim chosen-ciphertext security. Let $A \in (\mathbb{Z}/q\mathbb{Z})^{m\times n}$, $s \in (\mathbb{Z}/q\mathbb{Z})^n$, and $b = As + e$. Choose a short vector $r \in \mathbb{Z}^m$, short errors $e_1 \in \mathbb{Z}^n, e_2 \in \mathbb{Z}$, and an encoded message $\mu \in \mathbb{Z}/q\mathbb{Z}$. Form

$$u = A^T r + e_1, \qquad v = b^T r + e_2 + \mu \pmod q.$$

The receiver computes $v - s^T u$.

Proposition 13.3 (Noise cancellation identity)**.** *With the notation above,*

$$v - s^T u = \mu + e^T r + e_2 - s^T e_1 \pmod q.$$

If the message encoding has decision radius ρ in circular distance modulo q, decryption is correct whenever the centered integer representative of the total error has absolute value less than ρ.

Proof. Substitute $b = As + e$:

$$v - s^T u = (As + e)^T r + e_2 + \mu - s^T(A^T r + e_1).$$

The two terms $s^T A^T r$ cancel, leaving the displayed expression. If the residual error lies within the decision radius of the message's encoding point, the nearest encoding rule returns that message by the definition of the radius. □

For binary encoding at 0 and $\lfloor q/2 \rfloor$, a conservative open decoding radius is half their minimum circular separation. Compression, rounding and a ring coefficient representation add further error terms in a real scheme. The correctness proof should include them rather than appeal only to this simplified identity. NIST's ML-KEM standard specifies a complete mechanism and three parameter sets, and its current publication page records an errata note; this monograph does not replace that specification [105].

For the following real-valued bound, choose and fix integer lifts of s, e, e_1, e_2; r is already an integer vector. Norms refer to those lifts, and reducing their residual modulo q yields the preceding modular identity.

Proposition 13.4 (A deterministic bound on the toy error)**.** *If* $\|e\|_\infty \le B_e$*,* $\|r\|_1 \le B_r$*,* $|e_2| \le B_2$*,* $\|s\|_1 \le B_s$*, and* $\|e_1\|_\infty \le B_1$*, then*

$$|e^T r + e_2 - s^T e_1| \le B_e B_r + B_2 + B_s B_1.$$

Proof. The triangle inequality bounds the absolute value by $\sum_i |e_i||r_i| + |e_2| + \sum_j |s_j||(e_1)_j|$. Use the coordinatewise bounds in the two sums. □

This is a worst-case sufficient condition. A failure-probability estimate for Gaussian or bounded random errors can be much sharper, but must specify independence and tails. A numerical sample in which every trial decrypts correctly does not prove a negligible failure probability over the full randomness space.

13.5 Why the noiseless hidden-subgroup argument does not transfer directly

The quantum Fourier argument in the previous chapter required an efficiently evaluable deterministic function constant on exactly the cosets of a subgroup. LWE samples are noisy observations. For a fixed secret, the conditional law of b given a is a translated error distribution, not a deterministic group homomorphism.

One may enlarge the domain to include a sampled error variable, but then the map $(a, e) \mapsto \langle a, s\rangle + e$ has a different domain distribution and requires access to the unknown s if one tries to evaluate it coherently. The public samples are not automatically an oracle for that map on arbitrary inputs. This explains the failure of a direct substitution into the earlier proof. It does not rule out all quantum algorithms for LWE; the computational hardness question remains an assumption supported by its reduction framework and cryptanalytic study.

The relation lattice used in Regev's factoring algorithm and the random lattices underlying LWE also play different roles. The former is a hidden relation object whose approximate dual can be sampled by a specially available group oracle. The latter is part of a noisy recovery problem whose public data does not provide the same oracle. Sharing a geometric language is useful for designing reductions, but does not itself supply one.

13.6 Orders, local structure and a responsible design criterion

The Roggenkamp framework suggests asking how an order behaves locally at primes dividing the modulus or discriminant. Does the module remain projective? Are there small quotients exposing unusually concentrated error? Does a decomposition of the algebra split the instance into easier components? Are coefficient and canonical norms comparable with a uniform bound? Is the reduction map efficiently computable and distribution preserving? These are concrete mathematical questions arising from order theory.

For example, a product algebra acts componentwise. If both the public equations and the error law split into independent factors, solving the components separately solves the original instance. A designer cannot count the total vector-space dimension as though the factors were computationally inseparable. Conversely, an algebraic decomposition may be difficult to obtain from a chosen encoding; that computational obstacle must be stated and analyzed rather than assumed from the abstract decomposition.

Proposition 13.5 (Product decomposition of a noisy module equation)**.** *Let $R = R_1 \times R_2$ and let all vectors and matrices be over R. Projecting $b = As + e$ to either factor gives $b_i = A_i s_i + e_i$. If algorithms recover s_1 and s_2 from the two projected instances, their pair recovers s in the original instance.*

Proof. The coordinate projections are ring homomorphisms, hence commute with every matrix addition and multiplication. Equality in a direct product is coordinatewise. The pair of recovered vectors is therefore exactly the original secret vector. □

The proposition makes no claim that every order has such an efficiently accessible product decomposition. Its use is diagnostic: extra algebraic structure can help computation on both sides of a security reduction. The correct research objective is to prove which structures preserve the desired hardness and improve implementation, not simply to maximize algebraic sophistication.

13.7 Other cryptographic families and the limits of the common framework

Family	Relevant mathematical interface	What does not follow automatically
RSA and multiplicative Diffie–Hellman	Integer orders, unit groups, order finding, relation lattices	A theorem about ECC gate counts does not supply RSA resource counts
Elliptic-curve signatures and key exchange	Finite abelian subgroups and endomorphism arithmetic	Extra endomorphisms do not protect the subgroup from Shor's general mechanism
Isogeny-based constructions	Ideal classes, quaternion orders, oriented isogenies	Hardness of one isogeny task does not imply security of every protocol with auxiliary data
Lattice-based KEMs and signatures	Noisy modules, duality, norm bounds and reductions	An order-module equivalence need not preserve errors or concrete security
Code-based constructions	Linear maps, syndrome kernels and decoding metrics	Chain-complex notation does not turn noisy decoding into noiseless period finding
Hash-based signatures	Finite function assumptions, combinatorial verification	No order-theoretic proof of a hash assumption is supplied here
Quantum key distribution	Quantum states, measurements and authenticated classical communication	Field-theoretic existence or a Yang–Mills gap is not a security proof for a QKD protocol

The broad connection is through rigorously represented maps, kernels, distributions and verification. It is useful to compare these families within that language, but their security definitions differ. A key-establishment mechanism, a signature scheme and a quantum communication protocol require different games, interfaces and adversary models. The original proofs and standards remain the authoritative sources for their security claims.

13.8 A quantitative requirement for any proposed generalization

Suppose a map T converts one family of metric module instances into another. A useful reduction needs at least three bounds: a polynomial bound for computing T and recovering answers, a bound on the distortion of the relevant geometric quantity, and a bound on the statistical distance between the transformed input law and the target law. The first is computational, the second geometric and the third probabilistic. None is implied by a bare equivalence of module categories.

For norm distortion, if T is invertible then

$$\|T^{-1}\|^{-1}\|x\| \le \|Tx\| \le \|T\|\|x\|.$$

The upper bound is the definition of operator norm. The lower bound follows from $\|x\| = \|T^{-1}Tx\| \le \|T^{-1}\|\|Tx\|$. Thus an approximation factor may deteriorate by the condition number $\kappa(T) = \|T\|\|T^{-1}\|$. If this grows exponentially in the input dimension, an algebraically simple transformation may be useless for a polynomial-factor hardness transfer. The elementary shear family shows why recording the metric is essential.

Chapter 14

Implemented cryptography and verifiable artifact contracts

Evidence status. Executable library-backed encryption, signature and key-encapsulation interfaces, with 31 passing regression-test methods and preserved runtime evidence. The framing and tests are local contributions. The cryptographic primitives are external implementations; no new security reduction or Lean verification of those implementations is claimed.

14.1 From mathematical examples to executable interfaces

The finite curve, hidden-relation and noise calculations in Chapters 12 and 13 explain mathematical mechanisms. The new `extension/crypto/` directory supplies a separate operational layer: authenticated artifact encryption, signatures and post-quantum key encapsulation. Its Python interface delegates every cryptographic primitive to `cryptography` 50.0.0 and its native backend. This choice makes the implementation boundary inspectable without presenting an educational arithmetic example as a deployed cryptosystem [125–127].

Function family	Primitive	Contract exercised
Artifact encryption	AES-256-GCM; ChaCha20-Poly1305	Byte confidentiality and authentication of ciphertext plus public version, algorithm and context
Artifact signatures	Ed25519	Verification of context-framed bytes relative to a separately trusted public key
Key encapsulation	ML-KEM-768; ML-KEM-1024	Creation and recovery of a shared secret through the backend's standard encapsulation API

The algorithm sources are the GCM recommendation, the ChaCha20–Poly1305 and EdDSA RFCs, and FIPS 203 for ML-KEM [48, 73, 104, 106]. ML-KEM's module-lattice basis connects conceptually to the preceding order-lattice discussion. Its security is not derived here from a finite determinant, a Hodge gap, a tropical model, GRH, or the fluid construction. A standard algorithm and a FIPS-validated cryptographic module are also different claims: this release records no module-validation certificate.

14.2 Authenticated serialization and expected context

The encryption envelope contains exactly five JSON fields: `version`, `algorithm`, `context`, `nonce` and `ciphertext`. The parser rejects unknown or duplicate fields, incorrect types, unsupported versions, noncanonical base64 and invalid lengths. Version is the integer 1; even the JSON Boolean `true` is rejected instead of being accepted through Python's integer comparison. The plaintext limit is 16 MiB and the context contains 1–1024 UTF-8 bytes. Context is authenticated public metadata, so confidential information belongs in the encrypted payload.

Associated data consists of a fixed application-domain prefix and the compact, sorted-key JSON encoding of version, algorithm and context. The recipient supplies the expected algorithm and expected context independently of the incoming envelope. These expectations cannot safely be copied from attacker-controlled metadata. Altering a header and simultaneously selecting the altered expectation still fails the library's authentication check. The full 16-byte authentication tag is retained, and no plaintext is returned before verification succeeds. These are concrete wrapper contracts around the external AEAD APIs, not an independent proof of those APIs' computational security [125].

Each encryption generates a fresh 96-bit random nonce internally. The interface does not accept caller-selected nonces. For m ideal independent uniform draws, the union bound over pairs gives

$$\Pr[\text{a nonce collision}] \leq \binom{m}{2} 2^{-96} = \frac{m(m-1)}{2^{97}}.$$

For $m \leq 2^{20}$ this particular probability is below 2^{-57}. This calculation addresses nonce collisions only; it is not the complete AEAD forgery or privacy bound. Random sampling does not guarantee perpetual uniqueness. The stateless wrapper cannot enforce a distributed key-use quota or retain a counter across process restarts. Applications need their own key lifecycle and usage policy, with independent keys for different algorithms and purposes. A fresh key for a single artifact is the simplest example, while secure storage and distribution of that key remain application responsibilities.

14.3 Signatures, replay and scientific meaning

The signature wrapper signs a separate domain prefix, the four-byte big-endian length of the UTF-8 context, the context itself, and the exact artifact bytes. The explicit length prevents the ambiguity between, for example, context `a` followed by message `bc` and context `ab` followed by message `c`. This is ordinary Ed25519 applied to the framed byte string; it is not the distinct Ed25519ctx variant specified in the EdDSA family [73, 126].

Verification requires an authenticated association between the public key and the intended signer. A valid signature authenticates bytes relative to that key; it does not establish that the signed mathematical assertion is true. Signing a verification manifest can preserve provenance, while its actual proof or test evidence must still be checked. The example keys are ephemeral and do not constitute an author identity or an archival trust anchor.

An authentic old envelope or signature remains authentic when replayed. Putting an artifact identifier or release number in the expected context binds that value, but preventing a second acceptance requires receiver state. Likewise, byte-exact signatures do not automatically identify semantically equivalent JSON encodings; any required canonicalization is an application contract. No secret key belongs in a public repository or client-side static site, and Python object destruction is not a guarantee of memory zeroization.

14.4 The post-quantum boundary

The two ML-KEM interfaces return and recover 32-byte shared secrets using the native library. The exercised ML-KEM-768 public-key and ciphertext lengths are 1184 and 1088 bytes; ML-KEM-1024 uses 1568 bytes for each. The tests serialize and reload public keys, encapsulate, decapsulate and check these sizes against the documented API [127].

ML-KEM is a key-encapsulation mechanism, not a signature scheme or complete authenticated transport protocol [104]. It does not establish either party's identity, specify file encryption or prevent replay. In particular, an invalid ciphertext of the correct length normally yields a rejection secret through implicit rejection; a returned byte string is not evidence that the ciphertext is valid. The regression checks this behavior and checks that malformed lengths are rejected. Combining these functions into a KEM/DEM transport would require separately specified authentication, transcript binding, derivation and lifecycle rules; this release makes no such protocol claim. Ed25519 remains a classical signature, so its presence beside ML-KEM does not make every component post-quantum.

14.5 Recorded execution and formal coverage

The recorded execution used Python 3.12.14, `cryptography` 50.0.0 and OpenSSL 4.0.1. All five algorithm choices in the table were exercised. The 31 test methods passed with no failures, errors or skips; parameterized cases additionally cover empty, binary, Unicode and larger payloads; ciphertext, tag and nonce changes; wrong keys, contexts and algorithms; parser ambiguity; signature changes; context framing; both KEM parameter sets; implicit rejection; and incorrect KEM lengths. Permanent boundary tests cover the full 16 MiB payload and 1024-byte context limits, oversized inputs, UTF-8 byte counting, noncanonical base64 pad bits and unpaired Unicode surrogates. Private keys and shared secrets were generated in memory and were neither persisted nor printed.

```
python extension/crypto/verify_crypto.py
```

The reproducible evidence is `extension/crypto/verification.json` together with `extension/crypto/verification.log`. The JSON record identifies the runtime versions and hashes the exact implementation, test driver, dependency pin and test log. A successful finite test run is not a proof for all inputs, a security reduction, an independent backend audit or a certification claim.

The companion Lean noise-cancellation and hidden-relation lemmas concern the finite algebra stated in their types. They do not formalize this Python code, the native backend, signature unforgeability, chosen-ciphertext security, side-channel resistance or authenticated key distribution. Keeping those interfaces explicit permits useful connections between arithmetic research and cryptographic engineering without changing the evidentiary meaning of either.

Chapter 15

Resource-certified elliptic arithmetic

15.1 A complete specification before a resource claim

The original section's criterion is a comparison between complete resource budgets. The implementation fixes a nonsingular short-Weierstrass curve $E : y^2 = x^3 + ax + b$ over a declared odd prime field, a public fixed generator G, an affine coordinate encoding, a gate alphabet and an ancillary-register policy. It generates ordinary fixed-base multiplication and a signed endomorphism decomposition under the same compiler. A shorter arithmetic expression is not counted as an improvement until every operation needed to realize it has been charged.

For $b_p = \lceil \log_2 p \rceil$, the point register has $2b_p + 1$ bits: b_p each for x, y, and a unique infinity word with the top bit set and both coordinates zero. A valid finite point has top bit zero, coordinates below p, and satisfies the curve equation. Every other bit string is invalid. Each in-place point translation acts as the group translation on valid words and fixes every invalid word. This extends the mathematical permutation to a total reversible operation on the register.

Proposition 15.1 (Total translation and coherent extension)**.** *For fixed $Q \in E(\mathbb{F}_p)$, the above map is a permutation of all point words. Its inverse is the extension of translation by $-Q$. Its permutation matrix is unitary, with no input-dependent phase.*

Proof. The valid words and invalid words are disjoint invariant subsets. Translation is a bijection on the first with inverse translation by $-Q$, and the identity is a bijection on the second. Their disjoint union is a permutation. Its columns are distinct standard basis vectors and therefore form an orthonormal basis. Linear extension establishes the action on every superposition. □

The scalar circuit acts on an arbitrary initial valid point, not only infinity:

$$|k, P, 0\rangle \longmapsto |k, P + kG, 0\rangle.$$

It preserves the scalar register and restores all declared scratch and decomposition registers. Invalid P remains fixed. The output is deterministic under the ideal logical-gate model.

15.2 An emitted gate construction with clean work

Every finite permutation is decomposed into cycles. A cycle is realized by a sequence of transpositions between one anchor and the remaining cycle entries. To swap two bit strings, choose a Gray path between them. Apply its adjacent transpositions in forward order and then reverse order omitting

the last transposition. The resulting permutation swaps only the endpoints; the internal path entries return to their initial locations.

An adjacent Gray transposition flips one target bit conditional on all other data bits matching a pattern. Negative controls are implemented by an X conjugation on the corresponding control wire. A k-controlled NOT with $k \geq 3$ is lowered to an AND ladder, one final controlled flip, and the reverse ladder. It uses $k-2$ clean work bits and $2k-3$ CCX gates. For zero, one or two controls the primitive is X, CX or CCX directly.

Proposition 15.2 (Clean controlled flip)**.** *On inputs whose ladder work bits are zero, the emitted ladder toggles the target exactly when every control is one, leaves the controls unchanged and restores every work bit to zero.*

Proof. The first CCX stores the conjunction of the first two controls. Inductively each following ladder bit stores the conjunction of the controls accumulated so far. The last CCX therefore toggles the target exactly under the full conjunction. The target is not a control in the ladder. Reversing the ladder toggles each stored conjunction off in reverse dependency order while its input controls still have their original values. Negative-control conjugation changes only which input pattern selects the flip and is undone afterward. □

The compiler emits every primitive rather than treating the controlled operation as a free black box. An independent parser counts the serialized primitives and reconstructs their all-to-all dependency depth. This count is separate from the generator's own count. Physical routing, fault tolerance, logical-error targets and T-gate synthesis are not inferred from an X/CX/CCX table. Those require another explicitly selected compilation architecture.

15.3 Field and variable-point oracles

The field reference supplies XOR-output addition, multiplication and inversion. For inversion, zero has a separate false validity flag rather than an invented inverse. Invalid field words contribute zero oracle output. The general reversible specification is

$$|x, z, 0\rangle \longmapsto |x, z \operatorname{xor} f(x), 0\rangle.$$

Because XOR by a fixed word is an involution for each retained x, arbitrary initial output words are supported.

The variable-point addition oracle retains both input points and writes the valid result together with a validity bit. Invalid point input pairs contribute zero to the output register. All valid pairs, all invalid input encodings, and independent nonzero output words are tested. This complements the in-place fixed-base translation circuit; it is not presented as the same register layout or the same resource budget.

The lookup/Gray-path compiler is deliberately explicit and bounded. Its worst-case table cost grows exponentially in the bit width. This is a complete finite reproducibility reference and a test bed for replacing each primitive by a scalable field-arithmetic circuit. It is not a claim of having improved cryptographic-size resource estimates such as [E1]. A future efficient replacement must pass the same total-domain and cleanup contracts before its resource table is comparable.

15.4 A fully paid endomorphism comparison

For the tested $j = 0$ curves, choose a nontrivial cube root of unity $\beta \in \mathbb{F}_p$ and define

$$\phi(x, y) = (\beta x, y), \qquad \phi(O) = O.$$

The generator's cyclic subgroup has $\phi(G) = [\lambda]G$. The finite decomposition table chooses signed pairs (a_k, b_k) satisfying

$$k \equiv a_k + \lambda b_k \pmod{\text{ord}(G)}.$$

The circuit coherently computes the signed pair into clean registers, performs controlled additions of the signed powers of G and $\phi(G)$, and uncomputes the pair. Two's-complement sign bits use negative powers in their corresponding controlled addition. Every scalar register word is interpreted modulo the declared subgroup order.

Here the base is public and fixed, so $\phi(G)$ is an explicitly recorded classical constant. This does not make evaluation of $\phi(P)$ on an unknown variable point free. Nor does a fixed-base table justify replacing a general lattice-decomposition algorithm by a costless oracle. The current comparison pays for the exact finite decomposition it actually executes.

Field	Path	Qubits	X	CCX	Depth
7	binary	16	2328	2904	3390
7	endomorphism	22	3808	4766	5608
13	binary	21	7040	8850	9832
13	endomorphism	27	9204	11768	13186

The endomorphism path costs more in both measured instances. For $p = 7$, it uses 4,766 rather than 2,904 CCX gates; for $p = 13$, it uses 11,768 rather than 8,850. The complete register counts rise from 16 to 22 and from 21 to 27, respectively. These are decisive answers for this compiler and these instances. They neither rule out a better compiler nor establish a saving at a different field size.

15.5 Formal and numerical verification layers

The Python suite checks the group operation on the complete small point sets, every declared scalar/point data word, invalid encodings, reverse-circuit replay and clean work. The $p = 7$ fixed-base domain contains 2,048 scalar/point words; the $p = 13$ domain contains 16,384. The gates are phase-free permutations, so complete basis semantics fixes their coherent linear extension. The separately executed GPU test verifies that extension on complex amplitudes across the entire $p = 7$ data domain.

The generated Lean certificate uses a direct modular affine-addition formula to check the emitted small-curve table, then checks all table associativity, identity, inverse, endomorphism and scalar-decomposition instances. These are exact finite propositions elaborated by the kernel, not statistical tests. General cleanup and normalized-map lemmas are formalized separately. Neither an exhaustive finite table nor a successful GPU comparison is represented as a formal asymptotic complexity theorem.

This completes the original subsection's concrete artifact program at its declared finite scope: a generator, reversible specification, total-domain tests, actual state-vector execution, independent gate counts and an assumption manifest. The next optimization question now has a reproducible baseline against which a claimed improvement can be falsified.

Chapter 16

Executable quantum codes and neural circuits on dessins

16.1 From a drawing to a quantum computation

A dessin becomes relevant to a quantum circuit only after the underlying combinatorics determine concrete Hilbert spaces, operators, and gates. This chapter carries out that construction. It supplies oriented cellulations, their clean-dessin permutations, binary boundary matrices, commuting Pauli checks, an executable encoder, and a small trainable circuit whose output can be independently derived. The smallest live example has eighteen data qubits, rather than a picture of eighteen unrelated two-level systems. Its state has entanglement enforced by the same incidence relations that describe the torus.

The construction is the homological surface-code construction introduced by Kitaev, within the Calderbank–Shor–Steane framework [30, 78, 154]. Here the contribution is a reproducible connection to the manuscript's orders, complexes and dessins, together with exact examples and verification artifacts. Existing hypermap-code literature already studies broader embedded combinatorial objects; the mere use of a dessin does not establish a new code family or better parameters [84, 136].

There are three computational representations. The browser executes the eighteen-qubit gate sequence as a full complex state vector. Larger Clifford computations can use stabilizer or Pauli-frame representations. GPU state-vector experiments use the resources stated in their recorded reports. These representations answer different computational questions. The size of a stabilizer simulation is not evidence that the same machine stored an arbitrary state on that many qubits.

16.2 Oriented cells and clean-dessin permutations

Let Σ be a connected, closed, oriented surface with a finite cellular embedding of a graph. Let V, E, F denote its numbers of vertices, edges and disk faces. An oriented edge has a tail and a head; reversing it produces the other dart of the same unoriented edge. There are therefore $2E$ darts. Put a black vertex at every original vertex and a white vertex at the midpoint of every edge. The resulting bipartite graph is clean: every white vertex has degree two.

Write σ_1 for dart reversal and φ for the permutation that follows the oriented boundary of a face. Our convention is

$$\sigma_0 = \varphi\sigma_1, \qquad \sigma_\infty = (\sigma_0\sigma_1)^{-1} = \varphi^{-1},$$

with products acting from right to left. The cycles of σ_0 are the rotations at black vertices, the transpositions of σ_1 are white vertices, and the cycles of σ_∞ are faces with reversed boundary

orientation. The code checks these identities as permutations; it does not infer them from the apparent shape of a rendered mesh.

Proposition 16.1 (Genus from permutation cycles)**.** *For the connected oriented cellulation above,*

$$2 - 2g = c(\sigma_0) + c(\sigma_1) + c(\sigma_\infty) - 2E = V - E + F,$$

where c denotes the number of cycles.

Proof. Subdividing each edge inserts one white vertex and replaces that edge by two bipartite edges. The new numbers of vertices and edges are $V + E$ and $2E$; the disk faces are unchanged. Its Euler characteristic is therefore $(V + E) - 2E + F = V - E + F$. The three cycle counts are V, E, F, by their definitions. The classification formula for a connected closed oriented surface gives the final value $2 - 2g$. □

An abstract rotation system needs a manifold check. The boundary cycles must use every dart exactly once, and all darts based at a given vertex must form one rotation cycle. Multiple rotation cycles at one declared vertex would create a nonmanifold identification. The implementation checks these conditions, checks graph connectivity, and requires every unoriented edge to occur twice in face boundaries with opposite signs. Thus the Euler characteristic is attached to an actual closed oriented cellulation, not merely to three integers satisfying an equation.

The Riemann existence theorem realizes a transitive permutation triple with product one as monodromy of a branched covering of the Riemann sphere. For a dessin, its branch values can be chosen as $0, 1, \infty$. Belyi's theorem relates this situation to curves defined over algebraic numbers [23]. These are external existence theorems. The generated permutations are explicit; algebraic coefficients of the higher-genus Belyi maps are not computed by this implementation. Nor does a Euclidean rendering specify the uniformizing metric of that Belyi pair.

16.3 An explicit elliptic Belyi model

The algebraic curve

$$E_0 : \quad y^2 = x^3 - x$$

provides a direct connection with elliptic-curve theory. Its smooth projective completion has genus one, and the meromorphic function $\beta = x^2$ has degree four.

Proposition 16.2 (A clean degree-four elliptic dessin)**.** *The branch values of $\beta : E_0 \to \mathbb{P}^1$ lie in $\{0, 1, \infty\}$. Its ramification partitions over these values are (4), $(2, 2)$ and (4), respectively. Its dessin has one black vertex, two white vertices and one face.*

Proof. At $(0, 0)$, the curve equation gives $x = -y^2 + O(y^6)$, so β has a zero of order four. At $(r, 0)$ for $r = 1$ or $r = -1$, the derivative of $x^3 - x$ is two. Thus $x - r = y^2/2 + O(y^4)$, and $\beta - 1 = 2r(x - r) + O((x - r)^2)$ has order two. At the point at infinity, x has a pole of order two, so β has a pole of order four. These fibers already account for degree four. Away from them, a critical point would require $x = 0$ or a critical point of the degree-two projection to the x-line; the latter occur at $x = 0, 1, -1, \infty$, already listed. The multiplicities give the claimed partitions and cycle counts. □

Topologically the dessin is a two-loop graph on a torus with disk complement. Choose its two loop classes as generators of the fundamental group. The connected unramified cover associated with $L_x\mathbb{Z} \oplus L_y\mathbb{Z} \subseteq \mathbb{Z}^2$ has degree L_xL_y. Pulling back the dessin gives the square torus cellulation

used below. The covering becomes holomorphic for the pulled-back complex structure, so composing it with β gives a clean Belyi map of degree $4L_xL_y$. This supplies a genuine elliptic-surface interpretation of the torus family without requiring numerical evaluation of elliptic functions during circuit execution.

For comparison, an explicit higher-genus family is

$$C_g: \quad y^2 = 1 + x^{2g+1}, \qquad \beta_g = -x^{2g+1} = 1 - y^2, \quad g \geq 1.$$

The double cover of the x-sphere has $2g+2$ simple branch points, including infinity, so its smooth completion has genus g. The displayed Belyi map has two black vertices of valency $2g+1$, $2g+1$ white vertices of valency two, and one face. Locally this follows from the two points over $x = 0$, the simple roots of $1 + x^{2g+1}$, and the pole orders at infinity. Its degree is $2(2g+1)$ and its Euler characteristic is $2 - (2g+1) + 1 = 2 - 2g$.

This last dessin is an instructive warning. Collapsing its white subdivision produces two vertices, $2g+1$ parallel edges and one face. The face boundary is zero modulo two, so the corresponding code has no nontrivial Z checks and distance one. Genus alone does not make an error-correcting code useful. The executable higher-genus cases instead use the nondegenerate connected-sum cellulations constructed below.

16.4 The cellular chain complex is the stabilizer specification

Work over $\mathbb{F}_2$. Give the cellular vector spaces their cell bases:

$$C_2 = \mathbb{F}_2^F \xrightarrow{\partial_2} C_1 = \mathbb{F}_2^E \xrightarrow{\partial_1} \mathbb{F}_2^V = C_0.$$

The boundary of an edge is the sum of its endpoints; a loop contributes twice and hence zero. The boundary of a face is the parity of its edge occurrences. Since the boundary of a closed polygon has each vertex an even number of times, $\partial_1\partial_2 = 0$. Define

$$H_X = \partial_1, \qquad H_Z = \partial_2^{\mathsf{T}}.$$

Their rows are the vertex stars and face boundaries. In the actual files the rows are sparse lists of qubit indices. They are not dense floating matrices subsequently rounded to binary values.

Place one qubit on each original edge. For $a, b \in \mathbb{F}_2^E$, put

$$X(a) = \prod_e X_e^{a_e}, \qquad Z(b) = \prod_e Z_e^{b_e}.$$

The single-qubit identity $XZ = -ZX$ gives

$$X(a)Z(b) = (-1)^{a\cdot b} Z(b)X(a).$$

Let A_v be the X product on the star of vertex v, and B_f the Z product on the boundary of face f.

Theorem 16.3 (Commutation from the chain identity)**.** *All A_v and B_f commute. Their common $+1$ eigenspace is a CSS code with*

$$k = E - \operatorname{rank} H_X - \operatorname{rank} H_Z$$

logical qubits.

Proof. Two products of X operators commute, as do two products of Z operators. The commutation sign of A_v with B_f is the (v, f) entry of $H_X H_Z^{\mathsf{T}} = \partial_1 \partial_2 = 0$. Choose independent rows in the two matrices. Their Pauli generators are independent jointly: an identity product must have zero X mask and zero Z mask separately. Ordering all X factors before all Z factors introduces no sign because these groups commute. Thus the generated stabilizer group contains 2^r elements, where $r = \operatorname{rank} H_X + \operatorname{rank} H_Z$, and does not contain $-I$.

The projector onto the common $+1$ eigenspace is $P = 2^{-r} \sum_{S \in \mathcal{S}} S$. Multiplication in the finite group proves $P^2 = P$, and Hermiticity follows because its elements are Hermitian involutions. Every nonidentity Pauli has trace zero, while I has trace 2^E. Therefore $\dim \operatorname{im} P = \operatorname{tr} P = 2^{E-r} = 2^k$. □

Theorem 16.4 (Logical qubits on a closed oriented surface)**.** *For a connected closed oriented surface of genus g,*

$$\operatorname{rank} H_X = V - 1, \qquad \operatorname{rank} H_Z = F - 1, \qquad k = 2g.$$

Proof. A vector in $\ker \partial_1^{\mathsf{T}}$ assigns equal binary values to the two endpoints of every edge. Connectivity forces it to be constant. Its kernel has dimension one, so $\operatorname{rank} \partial_1 = V - 1$.

Assign a coefficient a_f to each face. In $\partial_2 \sum_f a_f f$, the coefficient of an edge separating faces f and f' is $a_f + a_{f'}$. The dual graph is connected: a path between face interiors can be perturbed to cross edges transversely and avoid vertices. Consequently the kernel consists exactly of the constant face assignments, and $\operatorname{rank} \partial_2 = F - 1$. The all-face sum really has zero boundary because each edge occurs twice. Substitution in the CSS formula gives $k = E - V - F + 2 = 2 - (V - E + F) = 2g$. □

This is a quantitative use of topology, but its quantity is the dimension of a finite code space. It does not compute a classical cryptographic secret or a continuum field-theory mass. The theorem also explains why one must retain the dependent checks: the products of all star checks and of all face checks are each identity, although the checks remain useful local observables.

16.5 Logical operators, distance and syndrome

A Z string commutes with every star precisely when its edge mask belongs to $\ker \partial_1$. Multiplying it by a face stabilizer adds a boundary. Thus its logical class lies in

$$H_1(\Sigma; \mathbb{F}_2) = \ker \partial_1 / \operatorname{im} \partial_2.$$

Similarly X strings are classified by $\ker H_Z / \operatorname{row} H_X$, the corresponding cohomological quotient. The dot product is their commutation pairing. A nonzero pairing means that the associated logical X and Z anticommute. This provides the operator interpretation of the $2g$ homology generators.

For an error $X(a)Z(b)$, its syndrome is

$$s_X = H_X b, \qquad s_Z = H_Z a.$$

Each syndrome bit records an anticommuting check. A single Z error creates star defects at the edge endpoints, and a single X error creates defects on its two adjacent faces. A recovery with the same syndrome can still differ from the error by a nontrivial logical loop. Therefore displaying an error syndrome is not by itself a demonstration of a decoder. Repeated noisy syndrome extraction and decoding require the additional protocol assumptions studied in surface-code memory analyses [43].

For the $L_x \times L_y$ square torus with $L_x, L_y \geq 3$, index the horizontal and vertical edges by

$$h(x, y) = yL_x + x, \qquad v(x, y) = L_x L_y + yL_x + x,$$

using coordinates modulo the two lengths. A face at (x, y) has oriented boundary

$$h(x, y) + v(x + 1, y) - h(x, y + 1) - v(x, y).$$

There are $L_x L_y$ vertices and faces, and $2L_x L_y$ qubits.

Proposition 16.5 (Distance of the rectangular torus code)**.** *The square torus construction has parameters*

$$[[2L_x L_y,\ 2,\ \min(L_x, L_y)]].$$

Proof. The logical dimension is the preceding theorem. A binary cycle is a union of closed edge paths. If its homology class is nonzero, at least one constituent path has odd winding in a periodic direction. Lift that path to the square lattice in the plane. Its endpoint displacement has a nonzero coordinate equal to an integral multiple of L_x or L_y. Each edge changes one coordinate by one, so the path has length at least $\min(L_x, L_y)$. The same argument applies to the dual square lattice. Straight horizontal and vertical loops attain the corresponding lengths, so the minimum is attained.

For completeness, a mixed Pauli logical operator cannot improve this minimum. Its X and Z masks satisfy their separate centralizer conditions. If both were stabilizer masks, their product would also be a stabilizer. At least one is a nontrivial logical mask, and the weight of their union is at least the weight of that component. □

On the eighteen-qubit torus, examples of conjugate logical pairs are

$$Z_{\mathrm{h}} = Z_0 Z_1 Z_2, \quad X_{\mathrm{h}} = X_0 X_3 X_6, \qquad Z_{\mathrm{v}} = Z_9 Z_{12} Z_{15}, \quad X_{\mathrm{v}} = X_9 X_{10} X_{11}.$$

The displayed matched pairs overlap once; unmatched pairs have zero overlap. All commute with the stabilizers. The backend also checks distance independently by exhaustively testing increasing Pauli-support weights modulo the stabilizer row spaces. For CSS codes it suffices to test the two pure types, by the last paragraph of the proof.

16.6 Higher genus by a verified connected sum

To obtain genus two, take two copies of the 3×3 torus and remove one open square face from each. If their boundary vertices are $a_0, \ldots, a_3$ and $b_0, \ldots, b_3$ in their original face orientations, identify a_i with b_{-i}, with indices modulo four. This reverses the boundary orientation. The boundary edge $a_i a_{i+1}$ is paired with the reverse of $b_{-i-1} b_{-i}$.

The new counts are

$$V = 9 + 9 - 4 = 14, \quad E = 18 + 18 - 4 = 32, \quad F = 9 + 9 - 2 = 16.$$

Thus $\chi = -2$ and $g = 2$. Pairing edges explicitly is important: identifying every edge with the same unordered endpoint pair would incorrectly merge distinct parallel edges in more general maps. The implementation retains edge identities, reverses their incidence signs when required, and reconstructs every face and vertex rotation cycle.

A second connected sum with a third torus gives $(V, E, F) = (19, 46, 23)$ and genus three. The selected remaining face has four distinct vertices, so its closure is an embedded disk even when it shares vertices with the earlier seam. The same manifold checks are repeated after gluing. The counts are not accepted as a substitute for these checks.

Cellulation	E	g	$\operatorname{rank} H_X$	$\operatorname{rank} H_Z$	k	d
3×3 torus	18	1	8	8	2	3
3×5 torus	30	1	14	14	2	3
Two tori, connected sum	32	2	13	15	4	3
Three tori, connected sum	46	3	18	22	6	3
32×32 torus	2048	1	1023	1023	2	32

The first four distances are checked by exhaustive search through weight three, including proof by enumeration that smaller nontrivial logical supports do not exist. The last distance uses the rectangular-torus theorem, not a search over 2^{2048} masks. At a genus-two seam, some star checks have weight six; the genus-three example has star weights up to eight. All their face checks have weight four. These constructions increase the number of logical qubits, but the chosen small cells retain distance three. They do not establish an improved asymptotic rate-distance tradeoff.

16.7 A deterministic encoder with a complete proof

Let $W = \operatorname{row} H_X \subseteq \mathbb{F}_2^E$ and let $r = \dim W$. The prepared code state is

$$|\Omega\rangle = 2^{-r/2} \sum_{w \in W} |w\rangle. \tag{16.1}$$

Distinct binary strings label orthonormal computational basis vectors, so the normalization follows from $|W| = 2^r$.

Proposition 16.6 (Stabilizers and logical sector of the prepared state)**.** *Every star and face stabilizes $|\Omega\rangle$. Every $Z(b)$ with $b \in \ker H_X$ also has eigenvalue $+1$ on this state. Thus the state chooses the simultaneous $+1$ sector of a full set of logical Z operators.*

Proof. An X star adds an element of W to each summation label, permuting W and leaving its uniform sum invariant. For $b \in \ker H_X$, every $w \in W$ satisfies $b \cdot w = 0$, so $Z(b)|w\rangle = (-1)^{b \cdot w}|w\rangle = |w\rangle$. The face masks lie in $\ker H_X$ by the chain identity. The remaining homologically nontrivial masks provide the logical Z conditions. □

Reduce H_X to a binary reduced row-echelon basis $r_1, \ldots, r_r$, with pivot positions $p_1, \ldots, p_r$. On an all-zero register, apply a Hadamard to each pivot. Then, for every nonpivot position q with $(r_i)_q = 1$, apply $\mathrm{CX}_{p_i \to q}$. This is the actual gate sequence in the JSON files, not merely a reference to an abstract encoding unitary.

Theorem 16.7 (Correctness of the encoder)**.** *The specified Hadamard–CNOT circuit prepares* (16.1) *exactly in the ideal gate model.*

Proof. After the Hadamards, the pivot bits range uniformly over $u = (u_1, \ldots, u_r) \in \mathbb{F}_2^r$, with amplitude $2^{-r/2}$, and all other bits are zero. Reduced row-echelon form makes the pivot submatrix the identity. No later CNOT targets a pivot. Therefore its control retains u_i, and the final nonpivot bit q is $\sum_i u_i (r_i)_q$. The entire output string is $\sum_i u_i r_i$. Independence of the rows makes this a bijection from $\mathbb{F}_2^r$ to W. CNOTs permute basis states without phases, which proves the claimed uniform positive sum. □

The circuits have respectively $42, 74, 82, 134$ gates for the eighteen, thirty, thirty-two and forty-six-qubit cases. The eighteen-qubit circuit uses eight Hadamards and thirty-four CNOTs. This is an algebraic encoder; some CNOT endpoints are distant in the surface drawing. A nearest-neighbor hardware compilation would need routing, scheduling and extra gates. Local stabilizer supports do not automatically make this particular preparation circuit local.

The implementation uses little-endian basis indices: qubit q is bit q of the array index. Each gate specifies its name and target, a control for CNOT, and an angle for rotations. This convention is shared by the browser and the CPU/GPU interfaces. An inconsistent bit convention could produce a normalized state while assigning its amplitudes to the wrong qubits, so normalization alone is not an adequate validation.

16.8 A finite Hamiltonian with a provable gap

The same checks define the positive Hamiltonian

$$H_{\mathrm{code}} = \sum_{v} \frac{I - A_v}{2} + \sum_{f} \frac{I - B_f}{2}.$$

Every term is an orthogonal projector and all terms commute. The ground space is exactly the CSS code space.

Proposition 16.8 (Gap in the specified projector normalization)**.** *For a connected closed oriented cellulation containing at least one edge with distinct endpoints, the smallest positive eigenvalue of H_{code} is two. In particular, this applies to all four simulated torus and connected-sum cellulations in this chapter.*

Proof. Commuting Hermitian involutions admit a simultaneous eigenbasis. In it, the energy is the number of violated checks. The relations $\prod_v A_v = I$ and $\prod_f B_f = I$ force the number of violations of each type to be even. Consequently any nonzero energy is at least two. Choose an edge with distinct endpoints and apply its Z operator to $|\Omega\rangle$. Exactly its two endpoint stars change sign; all face checks retain eigenvalue $+1$. This nonzero vector has energy two, which proves that the lower bound is attained. □

This statement adds operator and normalization data that an Euler characteristic does not contain. It is a gap theorem for an explicit finite commuting-projector $\mathbb{Z}_2$ gauge model. It is not the mass-gap theorem for nonabelian continuum Yang–Mills theory. Even the physical questions within the code model require care: a growing error string can move its endpoints without increasing the number of defects. A gap alone therefore does not prove a passive, thermally self-correcting memory or a fault-tolerance threshold. Those conclusions need dynamical noise and recovery assumptions [43].

16.9 A surface-local variational neural circuit

Cong, Choi and Lukin introduced QCNNs using repeated local unitary layers, measurement-conditioned pooling, and progressively reduced system size. Their applications include phase recognition and optimization of quantum error correction [37]. The present demonstrator is more specific: it is a trainable four-parameter quantum correlator on a star of an encoded surface state. It has executable non-Clifford rotations and a measured output, but does not reproduce that paper's multiscale QCNN, learned decoder, or claimed performance comparisons.

Let the first star in the eighteen-qubit cellulation have ordered support

$$(q_0, q_1, q_2, q_3) = (0, 2, 9, 15).$$

The classical input is $s = (s_0, s_1, s_2, s_3) \in \{-1, 1\}^4$, and the trainable vector is $\theta \in \mathbb{R}^4$. Define

$$\phi_i = s_i \pi/2 + \theta_i, \qquad R_Y(\phi_i) = \exp(-i\phi_i Y/2).$$

Apply these rotations to the four listed qubits of $|\Omega\rangle$. Then apply CNOTs with controls q_1, q_2, q_3 and common target q_0. The output is the expectation of Z_{q_0}. Four-edge stars make this circuit local in the cellulation's edge-incidence sense; a physical chip still needs a specified coupling graph to realize those interactions.

Theorem 16.9 (Exact output of the variational correlator)**.** *For the specified state, patch and gates,*

$$f_\theta(s) = \left(\prod_{i=0}^{3} s_i\right)\left(\prod_{i=0}^{3} \cos\theta_i\right). \tag{16.2}$$

Proof. Conjugating Z_{q_0} backward through the three CNOTs gives $Z_{q_0} Z_{q_1} Z_{q_2} Z_{q_3}$. Direct multiplication of the two-by-two rotation matrix gives

$$R_Y(\phi)^\dagger Z R_Y(\phi) = \cos\phi\, Z - \sin\phi\, X.$$

Expand the product of these four expressions. Any term other than the all-X term contains a nonempty subset of Z factors on the star. The four neighbors of its center are distinct. A chosen Z edge is incident to a neighboring vertex with no other edge in that subset. Thus its Z mask has nonzero incidence and is not in $\ker H_X$.

More explicitly, for a Pauli $X(a)Z(b)$, the expansion of $\langle\Omega|X(a)Z(b)|\Omega\rangle$ is zero unless $a \in W$; when $a \in W$, it is $|W|^{-1}\sum_{w\in W}(-1)^{b\cdot w}$. If b is not orthogonal to W, pairing w with $w + w_0$ for a vector satisfying $b \cdot w_0 = 1$ makes this sum zero. Therefore all the terms just identified have zero expectation.

The all-X product is exactly the star stabilizer and has expectation one. Its coefficient is $\prod_i(-\sin\phi_i) = \prod_i \sin\phi_i$, because there are four factors. Finally $\sin(s_i\pi/2 + \theta_i) = s_i \cos\theta_i$, proving the formula. □

This theorem makes the neural experiment independently checkable. A state-vector implementation must execute the rotations and CNOTs, and its result can be compared with the product formula. The formula itself is an efficient classical baseline. Training this circuit therefore does not demonstrate a quantum computational advantage or solve a previously unknown classification problem.

16.10 Training, parameter shifts and the meaning of a held-out input

The target is the four-bit parity $y(s) = \prod_i s_i$. For a nonempty training set T, use mean squared error

$$\mathcal{L}(\theta) = |T|^{-1}\sum_{s\in T}(f_\theta(s) - y(s))^2.$$

Let $C(\theta) = \prod_i \cos\theta_i$. The exact output theorem gives

$$\mathcal{L}(\theta) = (C(\theta) - 1)^2, \qquad \partial_j \mathcal{L} = -2(C-1)\sin\theta_j \prod_{i \neq j} \cos\theta_i. \tag{16.3}$$

These identities include points where $\cos\theta_j = 0$, unlike an expression written using $\tan\theta_j$ and division by that cosine. They supply an independent reference gradient.

The quantum parameter-shift identity for this circuit is

$$\partial_j f_\theta(s) = \tfrac{1}{2}(f_{\theta+(\pi/2)e_j}(s) - f_{\theta-(\pi/2)e_j}(s)).$$

To prove it directly, hold the other parameters fixed in (16.2) and apply $[\cos(t+\pi/2) - \cos(t-\pi/2)]/2 = -\sin t$. More generally this two-evaluation rule follows when a parameter appears once in a gate generated by a Pauli operator divided by two. Repeated occurrences require a sum of gatewise derivative contributions, and different generator spectra require an appropriate rule [148]. The shift formula applies to the expectation; the loss derivative then uses the ordinary chain rule. Shifting the entire squared loss by $\pm\pi/2$ is generally not the same calculation.

The recorded training specification uses initial angles $(0.65, -0.4, 0.5, -0.6)$, four sign inputs $++++, -+++, --++, +-++$, twenty-four gradient steps, and learning rate 0.2. All sixteen possible sign inputs are then evaluated, including twelve not used in the numerical updates. The theorem explains why the same learned scale controls every input: the loss is independent of which nonempty subset of the truth table was selected. Consequently the held-out results are useful software checks, but they cannot support a statistical generalization claim for an unknown data distribution.

If the output were measured physically, each shot would produce $+1$ or -1, with mean $f_\theta(s)$. For M independent shots, the sample mean has variance $(1 - f_\theta(s)^2)/M$. A full-state simulator can compute the expectation directly without sampling noise. Its gradient agreement therefore verifies numerical circuit implementation, not the sample complexity or robustness of training on a noisy device. The report distinguishes these modes explicitly.

16.11 Execution evidence and scaling

The CPU validation executes the forty-two encoder gates on a complex128 array of length $2^{18} = 262144$. It compares the result with an independently constructed row-space superposition and evaluates all nine star and nine face expectations. The recorded norm squared is 0.9999999999999989, the maximum amplitude discrepancy is below 4.2×10^{-17}, and the state has exactly 256 nonzero amplitudes at the expected locations. The sixteen variational input computations agree with (16.2) to below 1.1×10^{-15}. These are floating-point execution results; the encoder and output theorems above provide the exact mathematical assertions.

The topology checks are independent of floating-point geometry. They verify edge incidences, oriented face closure, connectivity, rotation cycles, Euler characteristic, commutation and binary matrix ranks. The genus-two and genus-three certificates therefore establish nontrivial CSS codes even when a browser displays their charts without storing their exponentially large state vectors. A visual drawing can be changed without changing the checked combinatorial data.

The sparse grid interface constructs star and face supports in $O(L_x L_y)$ storage. It can supply much larger toric instances to a stabilizer backend without allocating a dense parity-check matrix or performing generic binary elimination. The rank formula is then a proved topological result, and the manifest identifies whether ranks were computed by row reduction or obtained from that theorem. The 2048-qubit archived case additionally has exact row-reduction checks.

A complex64 state vector requires $8 \cdot 2^n$ bytes before workspaces. Thirty qubits require 8 GiB; thirty-one require 16 GiB. A measurement ancilla can therefore double the dominant allocation. Forty-six qubits would require 512 TiB for the raw state alone. These arithmetic facts explain why different backends are necessary. A compressed Clifford state or batch of Pauli errors can be processed at a much larger qubit count, but arbitrary R_Y rotations generally leave the Clifford class.

16.11.1 Recorded GPU execution

The executable GPU evidence was obtained on 13 November 2026 on an NVIDIA GeForce RTX 4090. The recorded environment reports 25,250,627,584 bytes of device memory and 23,960,354,816 initially free bytes. It uses Python 3.12.3, cuQuantum Python 26.11.0, cuStateVec 1.15.1, cuStabilizer 0.5.0, and CuPy 14.2.0. The original record is `extension/quantum/results/gpu-results.json`; it contains nine successful experiments, the numerical outputs, timestamps, package versions, and the SHA256 of the executed script. The source and the record are retained separately so that a later, larger test does not silently replace earlier evidence.

Dense state-vector results. A two-qubit Bell-state calculation first checked amplitude ordering and XX, ZZ expectations. The surface-code runs then executed the emitted Hadamard/CNOT encoder using actual cuStateVec matrix applications. They computed all star and face expectations, injected an X error on a specified edge, and compared every resulting expectation with the predicted incident-face signs. One thirty-qubit run appended a fresh ancilla and four data-to-ancilla CNOTs to extract one face syndrome. This final experiment explicitly stored a thirty-one-qubit state.

Data and ancillas	Raw state	Encoder gates	Preparation (s)	Timed total (s)
$18 + 0$	2 MiB	42	0.000312	0.000537
$30 + 0$	8 GiB	74	0.918804	1.204211
$30 + 1$	16 GiB	74	1.838463	2.528538

All three cases use complex64 amplitudes. Their recorded norm squared is approximately 0.9999997616 and their maximum pre-error stabilizer expectation discrepancy is 2.385×10^{-7}. The extra ancilla has Z expectation approximately -0.9999997616, as predicted by the injected error. The routine evaluates these expectations directly; it does not simulate a finite collection of measurement shots. The listed raw-state sizes are exact allocations implied by the precision and qubit count. The required gate workspace was zero in these runs.

Preparation time covers encoder application and synchronization. The timed total begins after construction of the state-vector object and includes preparation, the reported expectation checks, and the inserted error and extraction gates. It excludes state allocation and final cleanup. These are measured intervals for this particular implementation, not complete application launch times or portable hardware benchmarks.

Training a circuit that has an exact analytic baseline. The eighteen-qubit neural experiment executed the encoded preparation, rotations, fan-in CNOTs, and output expectation on the GPU. Each gradient component was obtained from the parameter-shift expectation evaluations; the analytic product formula was used as a validation reference, not as a replacement for the quantum gates. Twenty-four updates at learning rate 0.2 changed the initial mean squared error 0.21987940665826008 to 0.0027618471230164243, a reduction by a factor of approximately 79.6. The recorded final angles are

$$(0.1918860335, -0.1237818552, 0.1522945504, -0.1791698376).$$

The largest discrepancy from the exact output formula along the recorded training history was 3.093×10^{-7}. All sixteen final truth-table entries had the correct sign, including the twelve inputs omitted from the updates. The training and evaluation interval was approximately 0.045774 seconds, excluding the initial encoded-state construction.

The initial angles already give $C(\theta) > 0$. Consequently the thresholded sign classifier is correct on all sixteen inputs before training as well as afterwards. The measured improvement is the calibration of expectation magnitude and mean squared error; it is not an increase from poor classification accuracy to perfect classification. The magnitude changes from approximately 0.531087 to 0.947447. If a single ideal measurement were used to predict parity, the exact Born-rule success probability $(1 + C)/2$ would accordingly change from approximately 0.76554 to 0.97372. These probabilities are deductions from computed expectations, not empirically sampled success frequencies.

Pauli frames and the reference-state assumption. For the larger experiments, cuStabilizer stores batched Pauli errors relative to a known ideal encoded reference. It does not store the 2^n amplitudes of that reference. The mathematical encoder theorem supplies an ideal state with every specified check equal to $+1$. The frame experiments assume that state and propagate its errors through ideal syndrome extraction. They do not execute the large encoder or construct a generic stabilizer tableau for an arbitrary unknown state.

Each experiment uses 32,768 error realizations with seed 20261113. Every data qubit receives an independent X error with probability 0.001 and an independent Z error with the same probability. Both can occur, giving the corresponding Y error up to an irrelevant phase. Syndrome ancillas and extraction gates are noiseless. A star ancilla is reset, Hadamard transformed, used as the control for its incident CNOTs, Hadamard transformed again, and measured. A face ancilla receives CNOTs controlled by the boundary data qubits and is then measured. Measurement randomization is disabled because the stored output is the error-induced sign change from the assumed deterministic ideal outcome.

Let $x, z \in \mathbb{F}_2^E$ be the data error masks. The exact predicted syndrome is

$$s_X = H_X z, \qquad s_Z = H_Z x.$$

With these noiseless, freshly reset ancillas, the extraction circuit preserves the data error masks. Hence the implementation can compare its measured sign-change bits with sparse binary products computed from the returned data frames. It does so for every stored syndrome byte, not merely for a random sample. Every comparison in the recorded runs has zero mismatches. This checks the consistency of the extraction circuit, qubit indexing, packed-bit representation and binary incidence calculation. It does not test a recovery algorithm or estimate a logical failure probability.

Case	Data	Ancillas	Apply interval (s)	Timed total (s)
Genus two	32	30	0.004756	3.025362
Torus $L = 32$	2,048	2,048	0.005072	0.013088
Torus $L = 128$	32,768	32,768	0.073878	0.165669
Torus $L = 512$	524,288	524,288	1.147181	2.693085

Here the apply interval is a synchronized wall-clock interval around noise and extraction circuit application, including host dispatch. It is not a CUDA-event measurement of kernel time alone. The timed total includes simulator setup and exhaustive syndrome verification, but excludes topology generation before entry into the timer and final cleanup afterwards. First-use compilation and initialization make the genus-two total much larger than its apply interval. No cross-backend speedup is inferred from these differently scoped measurements.

For the $L = 512$ experiment, all 2,147,483,648 packed syndrome bytes were checked. There were 68,489,748 violated checks over all shots, giving an aggregate syndrome rate of approximately 0.00398663. For a weight-four check under independent relevant Pauli errors of probability p, the exact odd-parity probability is

$$\frac{1-(1-2p)^4}{2}.$$

At $p = 0.001$ this is approximately 0.00398802. This agreement is a useful noise-model sanity check. Correlations among neighboring checks and their global parity constraints mean that the aggregate count should not be interpreted as independent Bernoulli samples at every check.

Memory accounting and the scope of a capacity claim. For S shots, the stored row stride in bytes is $4\lceil S/32 \rceil$. With N total data-plus-ancilla qubits and A measured ancillas, the implementation estimates the explicit packed bit tables as

$$B = (2N + A)\, 4\lceil S/32 \rceil.$$

The factor two accounts for the X and Z frame components. This is an analytic allocation estimate, not an instrumented peak-memory measurement. The first run's $L = 512$ case has $B = 10$ GiB. Before allocation, the program requires $1.25B$ to fit within available device memory after retaining a fixed 2 GiB reserve. This policy leaves room for temporary arrays and the desktop; it does not establish a universal memory bound for every library configuration.

A separate followup script, `extension/quantum/gpu_capacity.py`, retains the original experiment source and result unchanged. It hashes its own source, the imported experiment source, the topology source, all archived case JSON files, and the original result. Its capacity rule fixes 32,768 shots and chooses the largest multiple of 32 satisfying

$$1.25\,(10L^2)(4096) \le \text{free device bytes} - 2\,\text{GiB}.$$

The resulting test is a bounded workload selected under stated memory assumptions. A different shot count, error representation, precision, number of observables, or available-memory state changes the achievable qubit count. The dense thirty-one-qubit result and the much larger frame results therefore answer different computational questions.

Separate capacity followup. The followup ran from 10:58:12 to 10:58:17 UTC on the same date and device. It successfully exercised the genus-three code with 46 data qubits and 42 syndrome ancillas, and then selected $L = 640$ from the observed free-memory value of 23,922,606,080 bytes. The latter case has 819,200 data qubits and 819,200 ancillas, for a total of 1,638,400 qubits in each of 32,768 Pauli frames. Its estimated packed tables occupy 16,777,216,000 bytes, or 15.625 GiB. The budget including the 25 percent overhead allowance and 2 GiB reserve is 23,119,003,648 bytes, below the recorded available amount.

Case	Apply (s)	Timed body (s)	Full call (s)	Checked bytes
Genus three	0.002198	0.081878	0.209550	172,032
Torus $L = 640$	1.791833	4.327947	4.813268	3,355,443,200

Both cases had zero syndrome-byte mismatches. The larger test observed 107,029,374 violated checks, an aggregate rate of approximately 0.00398715. The additional full-call timer includes topology generation and cleanup; it still excludes Python startup, library imports, and output-file serialization. These boundaries make the measurements reproducible without calling a host wall-clock interval a pure kernel time.

The separate evidence file is `extension/quantum/results/capacity-results.json`. Its source and case hashes agree with the archived local inputs, and it records that those inputs were unchanged during execution. A byte-exact copy of the original first-run JSON is retained as `extension/quantum/results/gpu-results-remote-original.json`; the ordinary first-run JSON has identical parsed values but a few different numeric formatting choices. This distinction explains their different file hashes without conflating it with a difference in experiment data. The largest completed frame workload is therefore the stated $L = 640$ case under the stated memory policy. This observation establishes neither an absolute maximum qubit count nor the feasibility of arbitrary non-Clifford computation at that scale.

16.12 Extensions with explicit mathematical obligations

For a prime p, the cellular construction extends to qudits of dimension p. Retain signed integer incidence matrices and reduce them modulo p. Define $X|j\rangle = |j+1\rangle$ and $Z|j\rangle = e^{2\pi ij/p}|j\rangle$. Their products commute exactly when the corresponding signed row pairing is zero modulo p. The integer identity $\partial_1\partial_2 = 0$ supplies this condition, and the connected oriented-surface rank argument again gives $2g$ logical qudits. Over composite dimensions, field rank arguments must be replaced by module or Smith-normal-form calculations. This is an actual point where integral order and module theory helps prevent an invalid inference. The present executable gate engine remains a qubit implementation; the qudit construction is a mathematical extension.

Nonabelian finite-group gauge models require richer local Hilbert spaces and constraints. Kitaev's anyon construction develops that direction, but the qubit star-and-face specification cannot simply be relabelled as a universal nonabelian simulator [78]. Likewise an arbitrary hypermap encoding may admit a circuit transformation to a surface code, rather than introducing a fundamentally different source of protection; constructive comparisons are available in the literature [136].

For learned quantum circuits, a substantive next step would specify a noise channel, a family of input states, a training distribution, a decoder or observable, and comparison with classical and analytic baselines. Multi-layer pooling and measurement feedback would require additional circuit semantics beyond the present four-parameter model. Any claim of improved decoding should measure logical failure after recovery, not just a reduced training loss or an attractive surface visualization.

The connection to the rest of the monograph is now concrete. A clean dessin supplies a rotation system and an integral cell complex. Reduction modulo two supplies compatible binary matrices. Their zero composite becomes commutation of quantum operators. Homology becomes logical Pauli operators. A row-space construction becomes a finite gate list, and a local observable becomes a trainable circuit with a provable output. Every transition has specified data and a checkable statement. These are useful new executable interfaces within the project, without identifying them with a solution of GRH, an improved ECC attack, or a continuum Yang–Mills mass-gap proof.

Part IV

Typed tensor diagrams and quantum graph learning

Chapter 17

Tensor diagrams, invariant theory and quantum graph grammars

Evidence status. The finite-dimensional identities in this chapter are proved below. The accompanying executable laboratory checks concrete instances. The new Lean module formalizes a precisely stated sequential linear-map fragment and its contextual rewrite soundness; it is not a formalization of every tensor category, every graph grammar, or cosmology.

17.1 A diagram must have a type and an interpretation

The common language needed for this project is more exact than a resemblance between pictures. A dessin is an embedded bipartite graph with a cyclic order at its vertices. A cell complex has boundary maps. A tensor network has labelled vector spaces and contracted indices. A circuit has an ordered causal composition and a measurement rule. A grammar has productions and an acceptance convention. One underlying graph may carry several of these structures, but the additional data have to be supplied. A surface embedding alone does not specify a Hamiltonian, a grammar, a quantum state, or a physical communication channel.

We work initially over $\mathbb{C}$ with finite-dimensional Hilbert spaces. Fix an orthonormal basis e_i of each space H, and let H^* denote its linear dual. A boundary type is a finite ordered tensor word in such spaces and their duals. The empty word denotes $\mathbb{C}$. A box $A : H \to K$ has components $A^j{}_i$; a wire joins an output of one box to an input of the same type of another. Repeated matching indices are summed over their full stated ranges. In particular,

$$(BA)^k{}_i = \sum_j B^k{}_j A^j{}_i, \qquad (A \otimes C)^{j\ell}{}_{ik} = A^j{}_i C^\ell{}_k.$$

The drawing convention in this chapter places inputs on the left and outputs on the right, with the order of tensor factors specified at every boundary. Reflection accompanied by complex conjugation denotes the adjoint; reflection alone is not the adjoint of a complex matrix.

This is a concrete finite-dimensional instance of the graphical language surveyed by Selinger, while the invariant-tensor calculations below follow the purpose of Cvitanović's birdtracks: lines carry representation indices and vertices carry specified invariant tensors [39, 149]. We use component proofs for the identities actually needed, without importing an unrestricted diagrammatic coherence assertion.

Proposition 17.1 (Elementary diagram semantics)**.** *For composable linear maps the following identities hold:*

$$(B \otimes D)(A \otimes C) = BA \otimes DC,$$
$$(BA)^\dagger = A^\dagger B^\dagger,$$
$$(A \otimes C)^\dagger = A^\dagger \otimes C^\dagger.$$

The swap $S_{H,K} : H \otimes K \to K \otimes H$, defined by $S(x \otimes y) = y \otimes x$, is unitary, has inverse $S_{K,H}$, and satisfies $S_{H',K'}(A \otimes C) = (C \otimes A)S_{H,K}$.

Proof. The $(j, \ell; i, k)$ component of the first left-hand side is $\sum_{a,b} B^j{}_a D^\ell{}_b A^a{}_i C^b{}_k$. The finite sum factors into $(BA)^j{}_i (DC)^\ell{}_k$. For the adjoint, $(A^\dagger)^i{}_j = \overline{A^j{}_i}$; taking the conjugate of the finite matrix-product sum reverses its order. The tensor assertion follows from conjugating the product of entries. The swap permutes an orthonormal tensor basis, hence is unitary. Applying the remaining identities to $x \otimes y$, a spanning set, proves them. □

Evaluation and coevaluation are the maps

$$\epsilon_H : H^* \otimes H \to \mathbb{C}, \quad f \otimes x \mapsto f(x), \qquad \eta_H(1) = \sum_i e_i \otimes e^i \in H \otimes H^*.$$

Although displayed in coordinates, η_H is basis independent: under $e'_i = \sum_j M^j{}_i e_j$, the dual basis uses M^{-1} and their coefficients contract to a Kronecker delta. These maps are the cup and cap of the diagram language. They are algebraic maps, not automatically normalized state preparation and measurement operations.

Proposition 17.2 (Snake identity and the closed loop)**.** *The composite* $(\mathrm{id}_H \otimes \epsilon_H)(\eta_H \otimes \mathrm{id}_H)$ *is* id_H*. Closing the output of $A \in \mathrm{End}(H)$ against its input gives* $\operatorname{Tr} A$*. An empty closed H-wire evaluates to* $\dim H$*.*

Proof. For $x = \sum_j x_j e_j$, the first composite is $\sum_i e_i e^i(x) = x$. The closed diagram for A is $\sum_i e^i(Ae_i) = \sum_i A^i{}_i$. Taking A to be the identity gives the dimension. The other snake identity follows by evaluating a dual basis vector in the same way. □

For $\dim H = d > 0$, the cup has squared norm d. A physical maximally entangled state uses $d^{-1/2}\eta_H$ after a chosen identification of the dual factor with a laboratory register. Using normalized cups at both ends changes a snake contraction by d^{-1}. Ignoring that scalar would change success probabilities. This small distinction is essential when a birdtrack identity becomes an actual quantum circuit.

17.2 Wedge powers, Young symmetry and determinants as diagrams

The symmetric group acts unitarily on $H^{\otimes r}$ by

$$P_\sigma(x_1 \otimes \cdots \otimes x_r) = x_{\sigma^{-1}(1)} \otimes \cdots \otimes x_{\sigma^{-1}(r)}.$$

Crossings represent these permutations. In characteristic zero define

$$A_r = \frac{1}{r!} \sum_{\sigma \in S_r} \operatorname{sgn}(\sigma) P_\sigma, \qquad S_r^{\mathrm{sym}} = \frac{1}{r!} \sum_{\sigma \in S_r} P_\sigma.$$

These are alternating and symmetric projectors. Over an integral ring the denominators cannot simply be inserted: the Schur constructions and straightening relations of the earlier chapters are precisely what permits a formulation before inverting factorials.

Proposition 17.3 (Alternating projector and its rank)**.** *$A_r^2 = A_r = A_r^\dagger$. Its image has an orthogonal basis indexed by strictly increasing r-tuples in $\{1,\dots,d\}$, so $\operatorname{Tr} A_r = \binom{d}{r}$, interpreted as zero if $r > d$. For $r = 2$,*

$$A_2 = \tfrac{1}{2}(I - S), \qquad S_2^{\mathrm{sym}} = \tfrac{1}{2}(I + S), \quad \operatorname{Tr} A_2 = \tfrac{1}{2}(d^2 - d).$$

Proof. In the double sum for A_r^2, for each τ there are $r!$ pairs (σ, ρ) with $\sigma\rho = \tau$, and their sign product is $\operatorname{sgn}(\tau)$. This proves idempotence. Since $P_\sigma^\dagger = P_{\sigma^{-1}}$ and inversion preserves sign, it is self-adjoint. A tensor basis vector with repeated entries alternates to zero: exchanging the repeated positions pairs equal terms with opposite signs. Otherwise its alternating image is a signed copy of the image associated to its sorted tuple. Distinct sorted tuples have disjoint tensor-basis supports, so these nonzero images are orthogonal and span the image. Their number is $\binom{d}{r}$. Finally S fixes exactly the d diagonal tensor basis vectors in the trace, so $\operatorname{Tr} S = d$ and $\operatorname{Tr} I = d^2$. □

General Young diagrams use row symmetrization and column antisymmetrization in a specified order. Their product is generally not an orthogonal projector as written, and normalizing or orthogonalizing it is a separate operation. We do not replace the integral Schur module by that product in characteristics dividing its normalization. The standard monomial theory already developed supplies the required relations and bases; the diagram notation now makes the permutations and contractions visible. In particular, a shuffling relation is a linear identity among permutation diagrams, not a license to exchange arbitrary noncommuting circuit gates.

For a map $A : H \to K$, define $\bigwedge^r A$ on decomposables by $x_1 \wedge \cdots \wedge x_r \mapsto Ax_1 \wedge \cdots \wedge Ax_r$. Alternation makes this well-defined. In the increasing-tuple bases, its entries are the $r \times r$ minors of A.

Proposition 17.4 (Exterior functoriality and Cauchy–Binet)**.** *For $A : H \to K$ and $B : K \to L$,*

$$\bigwedge^r(BA) = (\bigwedge^r B)(\bigwedge^r A).$$

Equivalently, for increasing r-tuples I, J of output and input indices,

$$\det(BA)_{I,J} = \sum_{|K|=r} \det B_{I,K} \det A_{K,J}.$$

If $BA = 0$ and $r > 0$, the displayed composite is zero. For $r = 0$ it is the identity of the ground field, so the last conclusion would be false.

Proof. Both maps in the first equation send a decomposable r-vector to $BAx_1 \wedge \cdots \wedge BAx_r$. Such vectors span. Expanding the intermediate image in its increasing-tuple basis yields exactly the finite minor sum. Alternatively, in the multilinear expansion of $BAx_1 \wedge \cdots \wedge BAx_r$, repeated intermediate indices vanish and the $r!$ orderings of each distinct set assemble into a determinant. When $BA = 0$ every factor is zero and $r > 0$; the empty wedge in degree zero is 1. □

The determinant is thus the action on the one-dimensional top exterior power. In coordinates it is an epsilon-tensor contraction. The Plücker relation for two vectors in a four-dimensional space expresses decomposability:

$$p_{12}p_{34} - p_{13}p_{24} + p_{14}p_{23} = 0, \qquad p_{ij} = x_i y_j - x_j y_i.$$

Each of the twelve expanded terms cancels with another after commuting scalar coefficients. This is the same exact relation checked by the existing algebra laboratory. Its graphical use preserves the typed external legs and coefficients. The minors, Schur modules, and invariant contractions thereby refer to the same algebraic constructions, rather than three unrelated kinds of pictures.

17.3 A color contraction with every normalization fixed

Let $N \geq 2$ and choose traceless Hermitian matrices T^a, $1 \leq a \leq N^2 - 1$, with $\operatorname{Tr}(T^a T^b) = \delta_{ab}/2$. They form a real basis of traceless Hermitian matrices and a complex basis of traceless matrices. An adjoint-index edge joins two generators; each generator also carries one incoming and one outgoing fundamental index.

Proposition 17.5 (Fundamental completeness and Casimir)**.** *For $1 \leq i, j, k, l \leq N$,*

$$\sum_a (T^a)^i{}_j (T^a)^k{}_l = \frac{1}{2}\left(\delta^i{}_l \delta^k{}_j - \frac{1}{N}\delta^i{}_j \delta^k{}_l\right), \qquad \sum_a T^a T^a = \frac{N^2-1}{2N} I.$$

Proof. The matrices $I/\sqrt{N}$ and $\sqrt{2}T^a$ form an orthonormal basis of $\operatorname{End}(\mathbb{C}^N)$ for $\langle A, B\rangle = \operatorname{Tr}(A^\dagger B)$. Expanding an arbitrary matrix X in this basis gives $X = \operatorname{Tr}(X) I/N + 2\sum_a T^a \operatorname{Tr}(T^a X)$. Compare the coefficient of the independent entry $X^l{}_k$ in its (i, j) entry. This yields the first identity. Set $k = j$ and sum over j: the right-hand side becomes $\frac{1}{2}(N - 1/N)\delta^i{}_l$. □

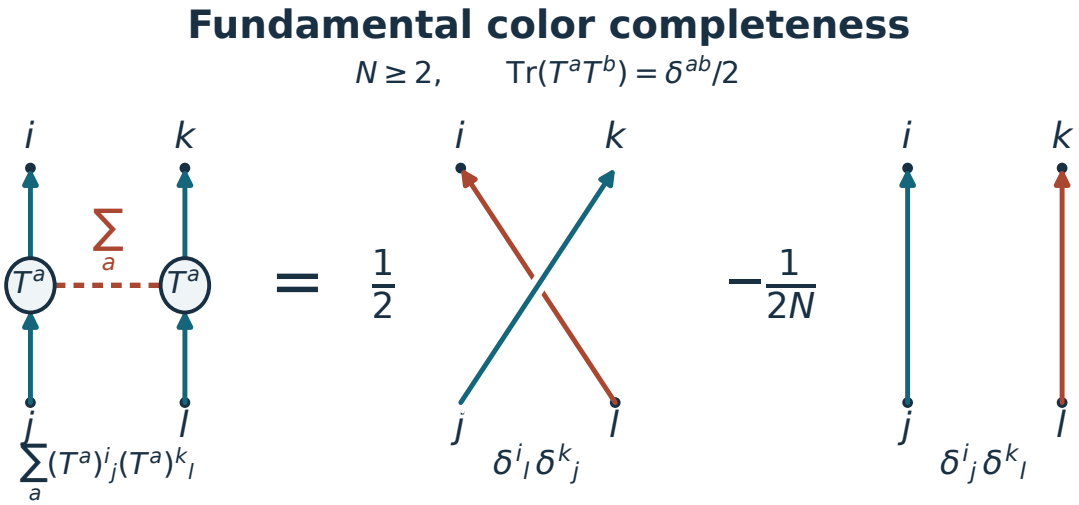


Figure 17.1: The fundamental completeness relation with its four external indices and every coefficient. In this figure fundamental wires run vertically, from the explicitly labelled inputs to outputs. The dotted adjoint edge denotes summation over the normalized generators.

The first term on the right joins opposite fundamental legs; the second joins each pair without exchanging them. The factor $1/N$ and the factor $1/2$ are part of the theorem. A diagram that suppresses them is not the same invariant. This contraction is relevant to color indices in Yang–Mills calculations, but it evaluates a finite representation-theoretic factor. It supplies no continuum measure, renormalization estimate, or lower spectral bound. The earlier analytic obligations remain in force.

17.4 Channels as positive tensors

To discuss noisy circuits, learned measurements, and any proposed cross-aeon message, a linear amplitude map is insufficient. A density operator ρ is positive semidefinite with trace one. A channel $\mathcal{E} : \mathrm{End}(H) \to \mathrm{End}(K)$ is completely positive and trace preserving. In the input-first convention its Choi tensor is

$$J_{\mathcal{E}} = \sum_{i,j} |i\rangle\langle j| \otimes \mathcal{E}(|i\rangle\langle j|).$$

Here $|i\rangle\langle j|$ is a matrix unit; the outer input leg is conjugated when interpreted as a transformation rule under basis change. The convention is explicit so a transpose cannot be silently replaced by a dagger.

Theorem 17.6 (Finite Choi–Kraus criterion)**.** *Let $\mathcal{E} : \mathrm{End}(H) \to \mathrm{End}(K)$ be complex-linear. It is completely positive exactly when $J_{\mathcal{E}} \succeq 0$. In that case there are matrices $K_\alpha : H \to K$ such that*

$$\mathcal{E}(X) = \sum_{\alpha} K_\alpha X K_\alpha^\dagger.$$

It is trace preserving exactly when $\sum_\alpha K_\alpha^\dagger K_\alpha = I_H$, equivalently $\mathrm{Tr}_K J_{\mathcal{E}} = I_H$.

Proof. If $\mathcal{E}$ is completely positive, apply $\mathrm{id}_H \otimes \mathcal{E}$ to the positive operator $|\Omega\rangle\langle\Omega|$, where $\Omega = \sum_i e_i \otimes e_i$ in a basis-labelled copy of $H \otimes H$. Its image is $J_{\mathcal{E}}$. Conversely diagonalize the positive matrix as $J = \sum_\alpha v_\alpha v_\alpha^\dagger$, incorporating square roots of eigenvalues into the vectors. Write $v_\alpha = \sum_i e_i \otimes K_\alpha e_i$; this defines every matrix entry of K_α uniquely. The (i, j) block of J is then $\sum_\alpha K_\alpha |i\rangle\langle j| K_\alpha^\dagger$. Equality on matrix units proves the Kraus formula. Each map $\mathrm{id}_m \otimes \mathcal{E}$ sends a positive Y to $\sum_\alpha (I_m \otimes K_\alpha) Y (I_m \otimes K_\alpha)^\dagger$, which is positive, proving complete positivity. Finally cyclicity of the finite trace gives $\mathrm{Tr}\, \mathcal{E}(X) = \mathrm{Tr}(X \sum_\alpha K_\alpha^\dagger K_\alpha)$. The trace pairing is nondegenerate, so equality with $\mathrm{Tr}\, X$ for all matrix units is equivalent to the stated matrix identity. Taking the partial trace of the Choi formula gives the transpose of this sum; equality to I is unchanged by transpose. □

The positive cone intersected with the affine partial-trace constraint gives a concrete feasible set for a learned channel. It differs from the affine variety $BA = 0$: positivity is an inequality, not a polynomial equality over $\mathbb{C}$. With real and imaginary coordinates, the channel condition can be written as polynomial equalities together with nonnegative principal minors. This is a semialgebraic constraint. The geometry of complexes may constrain other blocks of a model, but it cannot replace the channel's positivity test.

An isometry or unitary circuit followed by a partial trace automatically produces a channel. A free matrix-valued neural output does not. One practical parameterization is to output a unitary dilation and trace an environment. Another is to parameterize Kraus operators and explicitly enforce their completeness relation. Renormalizing arbitrary operators without checking invertibility or the remaining constraints is not a proof of trace preservation.

17.5 A precise grammar for the executable circuits

We now separate three different questions. *Generation* asks which finite graphs or circuit strings a grammar produces. *Interpretation* maps each well-typed expression to a matrix or channel. *Acceptance* asks what decision is made from a measurement or a classical computation on that interpretation. Their complexities can differ greatly.

For the finite laboratory let the wire type be $Q = \mathbb{C}^2$, with $1 \le n \le 6$. A gate record consists of a symbol from $\{H, X, Z, S, S^\dagger, T, T^\dagger, \mathrm{CX}, \mathrm{SWAP}\}$ and a valid ordered list of distinct wire indices of

the correct arity. A circuit is a sequence of such records; time runs from left to right and its matrix is the corresponding right-to-left product. Parameters such as $R_y(\theta)$ belong to a separate numerical experiment unless an angle alphabet or a computable encoding has been specified. Arbitrary real numbers are not a finite grammar alphabet.

A graph representation replaces each wire by a chain of directed edges and each gate by a typed vertex. Its input and output ports retain their order. In surface-constrained mode, allowed two-qubit gates must follow a supplied adjacency relation between data qubits on the cellulation. Relabelling vertices changes index names consistently, including the adjacency relation. A gate sequence is not a new dessin merely because it was rendered near a torus.

There are two useful kinds of production. A *generative* production appends an allowed gate and changes the represented unitary. An *equational* production replaces a subdiagram by an equal map and preserves the represented unitary. The laboratory's cancellation rule removes adjacent $U, U^\dagger$ on the same ordered wires. It also exposes commuting disjoint gates as an independently checked identity. No rule treats an approximate learned equality as exact.

Theorem 17.7 (Typed rewrite soundness)**.** *Interpret a well-typed expression using composition, tensor product, and the specified generator maps. Suppose every equational rule $L \rightsquigarrow R$ has the same input and output type and $[\![L]\!] = [\![R]\!]$. Replacing any subexpression by such a rule preserves the full interpretation. For unitary circuit interpretations, every finite sequence of these replacements preserves every output measurement probability on every input density operator.*

Proof. An expression with one hole is built from the hole by successively composing on either side with a fixed map or tensoring on either side with a fixed map. Equality is preserved by each operation, using bilinearity or simply substitution. Induction on this construction proves equality in every context. Induction on the number of rewrites proves equality after a finite derivation. Equal circuit matrices U, V give equal operators $U\rho U^\dagger = V\rho V^\dagger$, and hence equal numbers $\mathrm{Tr}(MU\rho U^\dagger)$ for every measurement effect $0 \preceq M \preceq I$. □

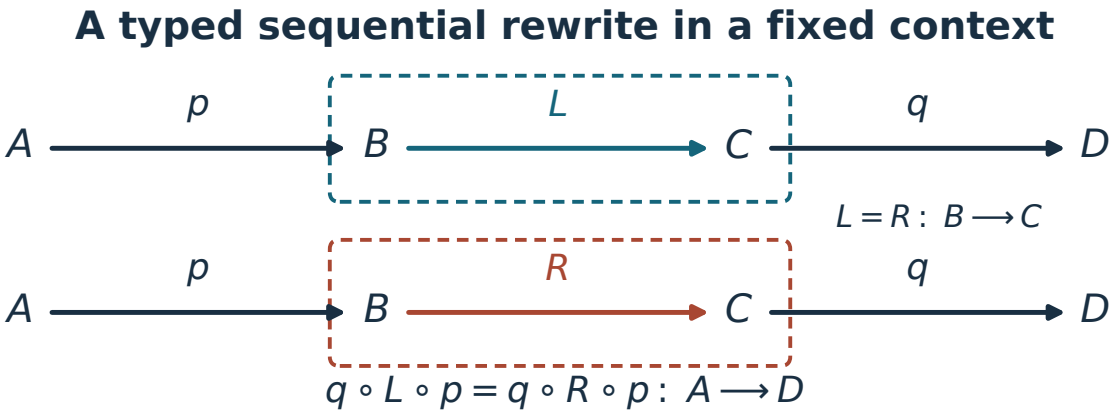


Figure 17.2: A typed sequential context. An equality of maps with the same domain and codomain can be substituted between fixed maps. The formal module checks this sequential syntax and finite derivations; tensor extensions are proved in the text.

The Lean file verifies the sequential linear-map part of this argument with a typed syntax and one-hole contexts. Tensor-context cases and density-operator consequences are proved in the

text. This is an exact scope statement: checking the sequential theorem does not certify a parser, graphical embedding, tensor coherence theorem, or browser floating-point arithmetic.

Proposition 17.8 (Termination and accumulated approximation error)**.** *Repeated adjacent inverse cancellation terminates. If U_j, V_j are unitary operators on the same Hilbert space and $\|U_j - V_j\| \le \epsilon_j$, then*

$$\|U_m \cdots U_1 - V_m \cdots V_1\| \le \sum_{j=1}^{m} \epsilon_j.$$

For a unit input vector and a measurement effect of norm at most one, the resulting probability difference is at most $2\sum_j \epsilon_j$.

Proof. Every cancellation reduces the nonnegative integer number of gates by two. For the norm bound, insert products whose first j factors have been replaced and telescope. Each difference is a left unitary product times $U_j - V_j$ times a right unitary product, so its norm is at most ϵ_j. For output vectors u, v, $\langle u, Mu\rangle - \langle v, Mv\rangle = \langle u - v, Mu\rangle + \langle v, M(u - v)\rangle$. Cauchy–Schwarz and $\|u\| = \|v\| = 1$ give the stated bound. □

Termination is not completeness. A circuit with no adjacent inverse pair may still equal the identity; for example $(HZ)^8$ does. The chosen local simplifier is allowed to miss an equality. Its guarantee is that every equality it does apply is valid. A learned policy can rank applicable valid rewrites without changing this guarantee. If it is also allowed to propose new identities, a separate certificate or numerical-error bound must be checked before acceptance.

17.6 How the grammar preserves a complex

The varieties-of-complexes constraint can be carried through a learned representation without pretending that any graph mutation preserves it. Let $C_2 \xrightarrow{d_2} C_1 \xrightarrow{d_1} C_0$ satisfy $d_1 d_2 = 0$.

Proposition 17.9 (Transport of a complex and its metric)**.** *Let $F_i : C_i \to C_i'$ be invertible linear maps and put $d_i' = F_{i-1} d_i F_i^{-1}$. Then $d_1' d_2' = 0$, and F_1 induces an isomorphism $\ker d_1 / \operatorname{im} d_2 \cong \ker d_1' / \operatorname{im} d_2'$. If the spaces carry Hilbert metrics and all F_i are unitary, the middle Laplacians satisfy*

$$d_1'^{\dagger} d_1' + d_2' d_2'^{\dagger} = F_1 (d_1^{\dagger} d_1 + d_2 d_2^{\dagger}) F_1^{-1}.$$

Proof. The middle inverse cancels in $d_1' d_2' = F_0 d_1 F_1^{-1} F_1 d_2 F_2^{-1} = 0$. The same identities show $F_1(\ker d_1) = \ker d_1'$ and $F_1(\operatorname{im} d_2) = \operatorname{im} d_2'$, giving a well-defined invertible quotient map. When $F_i^{-1} = F_i^{\dagger}$, taking adjoints of the transported differentials and cancelling the adjacent unitary factors gives the displayed conjugacy term by term. □

Invertibility alone gives homology invariance; it does not give this spectral equality for fixed metrics. This is the same algebra-versus-metric distinction encountered in lattice cryptography. A learned basis change should therefore specify whether it preserves the inner product, changes the metric along with the basis, or deliberately changes the spectrum.

Proposition 17.10 (A controlled enlargement)**.** *Adjoin the two-term complex $0 \to W \xrightarrow{I} W \to 0$ in any adjacent degrees, with zero couplings to the original complex. The direct-sum complex has the same homology as the original. With orthogonal direct-sum metrics its Laplacian gains an eigenvalue equal to one with multiplicity $\dim W$ in each of those two degrees.*

Proof. The added identity has zero kernel and full image, so its two homology groups vanish. Kernels and images of a block-diagonal differential decompose into direct sums, yielding the homology assertion. On each new W summand precisely one of $d^\dagger d$ and $dd^\dagger$ is I and the other is zero. Thus its Laplacian block is I. □

This is an elementary grammar rule for increasing representation size with a proof of homological preservation. It is not automatically a geometric subdivision of a given surface; constructing the required cellular chain equivalence is an additional task. It also shows why preserving homology does not preserve a numerical gap unchanged: when $W \neq 0$, a former gap greater than one in an affected degree becomes one after this enlargement.

17.7 Expressivity: four different classifications

The Chomsky hierarchy classifies languages of finite strings over a finite alphabet. A graph language enters that hierarchy only after a serialization and acceptance relation have been specified. The label "context-free graph grammar" does not assert that every natural string yield is a context-free language. Hyperedge replacement has its own gluing semantics and expressive distinctions [46].

Proposition 17.11 (Bounded interfaces give bounded treewidth)**.** *Consider a finite hyperedge replacement grammar. A production replaces one nonterminal hyperedge by a fixed finite graph with an ordered external interface; the only identifications are the prescribed attachment vertices, distinct external vertices remain distinct, and new internal vertices are fresh. Suppose every right-hand side and the initial graph has at most b vertices. For every finite complete derivation, the primal graph of its terminal result has a tree decomposition of width at most $b - 1$.*

Proof. Use the derivation tree, with a root for the initial graph and one node for each replacement. At each node place in a bag the images in the final graph of all vertices in its right-hand side, including its external vertices; at the root use the initial vertices. Each bag has at most b vertices. Every final vertex was either initial or introduced by a production and lies in its creating bag. Every pair of vertices made adjacent by a terminal hyperedge occurs together in the bag in which that hyperedge was introduced. Finally, a vertex can enter a descendant production only through the attachment interface of a hyperedge already containing it. Along the unique ancestry path from its creation to any later occurrence, it therefore remains in every intervening bag. Its bags form a connected subtree. These are exactly the covering, edge, and connectedness conditions for a tree decomposition. □

The bound applies to a fixed grammar with fixed finite production templates. It does not cover an arbitrary new gate layer supplied as an unbounded template, unbounded interfaces, or later identifications of distant vertices. In particular, a fixed bound cannot generate all increasingly large square grids: their treewidth is unbounded. This limits what a narrow-interface grammar can express on larger and higher-genus surface-code families.

Proposition 17.12 (Context-free graph generation can have a non-context-free string yield)**.** *A fixed-interface graph grammar with distinct attachment vertices can generate path graphs with edge labels $a^n \# b^n \# c^n$, $n \geq 1$. Their left-to-right string language, and its yield after erasing the two separators, are not context free.*

Proof. Use a nonterminal with six distinct external vertices, the endpoints of three ordered paths. Its base production makes three one-edge paths labelled a, b, c. Its recursive production extends the first path by one a, the second by one b, and the third by one c, passing their new endpoints as the

interface of one recursive occurrence. An initial production adds one separator-labelled edge from the end of the first path to the start of the second and another from the end of the second to the start of the third, producing one path without identifying any attachment vertices. Equivalently its three-component yield obeys $(x, y, z) \mapsto (ax, by, cz)$ from (a, b, c), followed by separator insertion. Induction gives the stated graph language and, after erasure, $\{a^n b^n c^n : n \geq 1\}$. If the latter language were context free, let p be a pumping length and consider $a^p b^p c^p$. In any pumping decomposition with the central three-piece block of length at most p, that block meets at most two adjacent constant-letter blocks. Pumping the two designated pieces changes at least one count, leaves at least one of the three counts unchanged, and either destroys the letter order or makes the counts unequal. The pumped string is outside the language, contradicting the pumping property. Finally, erasing a terminal in the productions of a context-free grammar yields a context-free grammar for the erased language, allowing empty productions. Thus the separator-containing language cannot have been context free either. □

The example's final graphs are paths, so their treewidth is only one. Graph width, language yield complexity, and quantum simulation complexity are different quantities. The separator-erasing homomorphism is specified explicitly rather than hidden in an informal reading of the graph.

Proposition 17.13 (Unrestricted effective graph rewriting can express Type-0 recognition)**.** *Allow finite labelled graphs, local finite rewrite rules, unbounded graph growth, and acceptance when a distinguished accepting state is reached. This model can simulate every Turing machine. Its recognized input-string languages are exactly the recursively enumerable languages.*

Proof. Represent a machine tape by a directed path of cells with symbol labels and a distinguished cell carrying the machine state. A transition changing symbol a to b and moving right is a local rewrite of the current head cell and its right neighbour: replace a by b and move the state marker to that neighbour. A second rule handles the end of the represented tape by creating a fresh blank cell. Left moves have analogous rules, with a fixed left-end marker for a one-sided tape. Since state and tape alphabets and the transition table are finite, there are finitely many rules. Induction on machine steps proves exact correspondence between configurations and reachable labelled paths. An accepting state is reached exactly when the machine accepts. Conversely, finite labelled graphs and finite rule matches have effective finite encodings. Breadth-first enumeration of finite derivations semi-decides whether an accepting graph is reachable. This gives the reverse inclusion. □

No general terminating algorithm decides all acceptance questions for that unrestricted model: such an algorithm would decide Turing-machine acceptance, contradicting diagonalization. Adding quantum amplitudes does not justify omitting computability assumptions on coefficients or providing a finite simulator with unbounded resources.

Model	Precisely classified object	Consequence
Bounded browser laboratory	Finite circuits on at most six wires and a bounded gate count	A finite set; every serialization language is regular.
Fixed-width unbounded instruction strings	All well-typed instructions from a finite gate alphabet	Syntactic validity is regular; semantic equality is a different problem.
Finite-interface replacement grammar	Graphs glued through specified ports	A bounded-width graph class under the proposition's assumptions; string yields need not be context free.
Unrestricted effective rewriting	Existence of a finite accepting derivation	Type-0 recognition; general termination is undecidable.
Uniform polynomial quantum circuits	Bounded-error measurement decisions as input length grows	BQP, a computational complexity class, not an extra level of the Chomsky hierarchy.

For the last row, "uniform" means that a classical algorithm produces the size-n circuit in polynomial time, using a fixed finite gate set with effectively computable entries. Bounded error means acceptance probability at least $2/3$ on yes instances and at most $1/3$ on no instances. Noncomputable real parameters would hide an oracle in the model. A fixed six-qubit browser demonstration is a finite numerical example, not evidence that an asymptotic BQP-hard problem has been solved efficiently.

17.8 When the tensor grammar is also efficiently classical

Proposition 17.14 (Dynamic contraction along a derivation tree)**.** *Suppose every tensor index has dimension at most D, each production has at most q distinct local index variables including its external interface, and there are m production instances. If child tensors are contracted and summarized at their prescribed interfaces, the final tensor can be computed with at most $O(mD^q)$ scalar operations, up to a fixed factor depending on the grammar. The bound assumes that the required final interface also fits within q.*

Proof. Process the finite derivation tree from its leaves upward. At a production, the already computed child tensors and the finitely many local generator tensors have entries indexed by the at most q local variables. Enumerate all their assignments, multiply the relevant entries, and add to the corresponding external-interface entry. There are at most D^q assignments, and a fixed grammar bounds the number of factors per assignment. This computes exactly the contracted subdiagram by associativity and distributivity of finite sums. Summing the work over m nodes proves the estimate. □

This explains an important tradeoff for quantum learning on diagrams. A bounded-interface family may be interpretable, trainable, and classically tractable for the same structural reason. It should be chosen for those benefits when appropriate; its tractability is not evidence of a quantum speedup. Conversely, allowing wide entangling interfaces can remove this contraction bound and make state-vector simulation exponential in the qubit count. The earlier GPU results distinguish dense amplitudes from stabilizer frames precisely for this reason.

The graph's usual treewidth is also not automatically the contraction width. A tensor of very high valence can have exponentially many entries even when the vertex graph is a star. The bound above controls the actual index interface. This supplies a concrete complexity budget for a quantum graph grammar rather than inferring computational efficiency from an attractive surface drawing.

17.9 A learned rule selector with a checked semantic core

A quantum or classical neural network may assign scores $s_\theta(r, G)$ to applicable rules r on a current graph G. Mask ill-typed rules before forming a softmax. The learned policy chooses among the remaining rules; the certificate checker verifies the local algebraic condition. If all permitted moves are the equational moves of Theorem 17.7, every accepted derivation preserves the represented function regardless of the policy's weights. If moves instead generate candidate models, the learned score supplies a proposal distribution and carries no equality guarantee.

These two uses support different research tasks. For a toric code, a selector can simplify a stabilizer circuit using verified identities. For a variety of complexes, it can propose a basis transport with the invertibility or unitarity certificate required above. For a CMB inference problem, it can generate a family of source templates and rank their likelihoods, but the observational null calibration must include that selection process. For finite tokenization DAGs, the quantum ToricGT experiment below proposes paths and checks them against an exact shortest-path certificate. In every case the mathematical rule, the learned search heuristic, and the final certificate have distinct jobs.

An L-function or zeta function requires an additional arithmetic or spectral definition. A tensor determinant may be a characteristic polynomial, a finite partition function, or a local factor when the required operator has been specified. Merely labelling a closed diagram by ζ does not turn its contraction into the Riemann zeta function. Similarly, graph rewriting can preserve a determinant under a proved similarity or multiply it by a known factor under a block extension; a statement about its zeros then follows from that actual identity. The earlier convergence and nonvanishing conditions remain necessary for any passage to an infinite analytic object.

17.10 Executable obligations

The companion's grammar laboratory displays a circuit before and after inverse cancellation, the typed wire lists, the gate count, and the numerical maximum-entry difference between their full matrices. It shows the alternating projector's rank and closed-loop dimension, and verifies the two sides of the fundamental color contraction. Its language panel runs an exact integer-state recursive recognizer and contrasts bounded prefixes with the unbounded theorem proved in the next chapter. The executable checks have independent matrix or combinatorial references; a displayed small residual is floating-point evidence for those instances, while the propositions establish the general identities under their hypotheses.

This gives the requested birdtracks-level discipline: every wire has a space, every vertex has a tensor, every contraction has a summation, every scalar normalization is retained, and every rewrite has a precise soundness obligation. The interpretation remains equally valid when the graph is drawn in three dimensions or projected from four-dimensional coordinates; those visual choices do not alter its algebra.

Chapter 18

A trainable quantum interface for ToricGT

Evidence status. This chapter supplies an original, bounded implementation and elementary proofs for a quantum feature interface to a graph transformer. It reports nine completed GPU training runs and independent numerical checks. It does not identify this implementation with the complete upstream ToricGT program, a biomolecular language model, or a quantum computer. The simulation uses three qubits per graph token. Its exact path certificates are classical mathematical certificates, with their own proofs and verification code.

18.1 Which architecture is being extended?

A useful connection between algebra and learning begins with a precise interface. A theorem about a ring, a circuit, or a chain complex does not automatically become a theorem about a trained predictor. We therefore give each part of the construction a separate role. Graph structure specifies legal objects; a finite circuit supplies trainable features; a transformer predicts a value function; and an independent checker determines what a proposed solution actually establishes.

The implementation sources reviewed for this extension are pinned to `amelie-iska/ToricGT` commit `a5199bc989a6fb170e21220e5e87bf95bc2fc0b4` and `amelie-iska/ToricGT-Program` commit `f91ba1811b5a4f70d53c7120acf8ca2a0049f85b`. The four source modules inspected in full implement graph tokenization, tropical attention, finite Weyl matrices, and Soft-MoE routing. The program overview requires empirical evaluation of proposed algebraic auxiliaries; it does not make their utility follow from their names. [145, 146]

Our model retains vertex and edge tokens, endpoint incidence, a mixture of ordinary and max-plus attention, and four experts with two slots each. Its new component is a jointly trained circuit feature encoder. The older repository example under `examples/quantum_tokengt` used a frozen three-qubit feature construction and trained a classical attention head. The present experiment trains both the input-to-angle map and twelve circuit offsets. This is a concrete change in the optimization problem, not an increase in the number of qubits. The larger surface-code circuits described elsewhere in this monograph are separate experiments.

There are also deliberate changes to the upstream graph encoding. The reviewed tokenizer contains learned identifier embeddings. Such a table is not automatically invariant under arbitrary renaming of vertices. Here the identifiers enter only through equality of endpoints, equivalently through inner products of orthonormal identifier vectors. Intrinsic byte positions remain ordinary

input attributes. A permutation of storage order transports these attributes; it does not change the underlying string. This distinction will be used explicitly in the equivariance proof.

The ToricBLM source located during this review is the author's June 2026 proposal, which discusses graph tokens and affine steering regions. It is useful as a design source. An independently pinned, trained ToricBLM repository was not located in the reviewed material. Consequently this chapter makes no claim to reproduce its proposed 100-million-parameter biomolecular model, its biological validation program, or broad claims of generalization.[143]

The selected task is exact weighted tokenization on small directed acyclic graphs. It gives the learner a meaningful combinatorial problem while keeping the teacher and certificate transparent. This choice also makes a limitation explicit: dynamic programming already solves the task efficiently. The experiment tests whether a particular trainable representation works; it does not establish a computational advantage over that algorithm.

18.2 Tokenization graphs and exact solution certificates

18.2.1 The fixed-vocabulary problem

Let a string have length n. Its byte boundaries are vertices $0, 1, \dots, n$. An allowed token occupying the substring from boundary i to boundary j gives an edge $e : i \to j$, where $i < j$. An edge has integer cost c_e. Every one-byte edge is available, so a path to n exists from every vertex. A complete tokenization is a directed path from 0 to n. In the experiment, candidate tokens have length at most three and the alphabet has four symbols. Costs depend on a recorded, bounded control parameter; the full graph and its costs are supplied to the model.

The colored-DAG formulation follows the tokenization viewpoint of ConvexTok. Its vocabulary-selection problem and its decoding problem must be distinguished. Choosing a globally shared vocabulary under a size budget is a discrete optimization problem. Once the allowed vocabulary is fixed, the remaining path problem on a DAG admits the elementary solution below. [165]

Proposition 18.1 (Bellman recursion on a finite token DAG)**.** *Suppose all edges point from smaller to larger vertex indices and every nonterminal vertex reaches n. Set $d(n) = 0$ and, in decreasing order of i, define*

$$d(i) = \min_{e:i\to j} \{c_e + d(j)\}.$$

Then $d(i)$ is the minimum cost of a path from i to n. Choosing a minimizing edge at every visited vertex produces an optimal path.

Proof. At n the only path of length zero costs zero. Suppose the assertion holds for vertices larger than i. Every path beginning at i has a first edge $i \to j$ and a remaining path beginning at j. By the induction hypothesis its cost is at least $c_e + d(j)$, hence at least the displayed minimum. Conversely choose an edge attaining the minimum and append an optimal path from its target. The result attains $d(i)$. The graph is finite and every edge strictly increases the index, so this procedure terminates. Backward induction proves every assertion. □

The recursion works even with negative edge costs because there are no directed cycles. Our sampled costs are positive, which makes their interpretation as token costs convenient but is not needed for this proof. With outgoing adjacency lists, the arithmetic work is linear in the number of edges. This is the exact teacher used during dataset generation.

Proposition 18.2 (A checkable primal–dual certificate)**.** *Let $p : V \to \mathbb{R}$ satisfy $p(n) = 0$ and*

$$p(i) - p(j) \le c_e \qquad (e : i \to j).$$

For every path P from 0 to n, $p(0) \le C(P) := \sum_{e \in P} c_e$. If equality holds for one legal path, that path is optimal. The Bellman distances in Proposition 18.1 give such a potential and equality on every path obtained by choosing minimizing edges.

Proof. Write the successive vertices of P as $i_0 = 0, i_1, \dots, i_m = n$. Adding its m inequalities gives

$$\sum_{r=0}^{m-1} c_{i_r i_{r+1}} \ge \sum_{r=0}^{m-1} (p(i_r) - p(i_{r+1})) = p(0) - p(n) = p(0).$$

Thus $p(0)$ is a lower bound for every feasible solution. A path attaining the bound is optimal. The Bellman minimum implies $d(i) \le c_e + d(j)$ for each outgoing edge, and a minimizing edge makes this inequality an equality. Summing along the resulting path proves the last assertion. □

The certificate checker verifies the token substrings, continuity of the path, its endpoint, all potential inequalities, and equality of the final cost and lower bound. It uses literal integers, rather than a numerical tolerance. A learned value function $\widehat{d}$ supplies a candidate path by minimizing $c_e + \widehat{d}(j)$ at each step. Rounded predictions supply an initial integral potential. Backward replacement

$$p(i) \leftarrow \min\bigl(p(i), \min_{i \to j}(c_{ij} + p(j))\bigr)$$

makes that potential feasible while preserving $p(n) = 0$. Feasibility does not force the proposed path to attain its lower bound. This is why the experiment reports the frequency of an immediately certified learned proposal separately from the frequency of a merely optimal learned path.

If the certificate does not close, an exact dynamic-programming repair is available. The final output is then certified, but its certainty belongs to the checker and repair algorithm. Calling this a perfect neural success rate would conceal the distinction between prediction and verification. The stored results retain the original proposal, its regret, its initial certificate, whether repair was required, and the final certificate.

18.2.2 A separate vocabulary-selection relaxation

For completeness, consider several token DAGs with a shared finite set of token colors $\mathcal{C}$. For each document let f_e indicate whether an edge is used, and let y_c indicate whether a priced token color is included. Byte edges need no color budget. With the incidence convention $Be = e_{\mathrm{source}(e)} - e_{\mathrm{target}(e)}$, impose

$$Bf = e_s - e_t, \qquad f_e \le y_{c(e)} \text{ on priced edges}, \qquad \sum_c y_c \le K.$$

The integer program uses binary variables and minimizes the number, or specified weighted cost, of used edges. Replacing binary variables by intervals $[0, 1]$ gives a linear relaxation. A rounded vocabulary must be evaluated by decoding again; a fractional flow is not itself a vocabulary. Our companion implements this small formulation on the three strings `abca`, `bcab`, and `cabc` with budget two.

Lemma 18.3 (Integral unit flows on an ordered DAG)**.** *Let every edge of a finite graph increase a fixed vertex order. If $f \in \mathbb{N}^E$ and $Bf = e_s - e_t$, with distinct vertices s first and t last, then f is the incidence vector of one directed s–t path.*

Proof. At s, incoming flow is zero and outgoing flow is one. Integrality selects exactly one outgoing edge of multiplicity one. Follow it. At an internal vertex reached by a unit of flow, conservation supplies an outgoing edge with positive integral multiplicity unless the vertex is t. Increasing order prevents repetition and forces eventual arrival at t; otherwise a last visited nonterminal vertex would violate conservation. Subtract this path's flow from f. The remainder is nonnegative and has zero boundary. A nonzero nonnegative flow with zero boundary on a DAG is impossible: choose its first vertex incident to a positive outgoing edge. No positive incoming edge can enter that vertex, contradicting conservation. Therefore the remainder is zero. □

Proposition 18.4 (An exact lower bound from a rational LP dual)**.** *For the real linear program*

$$\min c^T x, \qquad Ax = b, \quad Hx \le h, \quad 0 \le x \le \mathbf{1},$$

suppose y is unrestricted, $z \le 0$, $u \le 0$, and $A^T y + H^T z + u \le c$ coordinatewise. Then every feasible x satisfies

$$c^T x \ge b^T y + h^T z + \mathbf{1}^T u.$$

If a feasible integer solution attains this lower bound, it solves both the relaxation and the integer problem.

Proof. Multiply the coordinatewise dual inequality by $x \ge 0$ and sum. This gives $c^T x \ge y^T Ax + z^T Hx + u^T x$. The equality constraint supplies $y^T b$. Because $z \le 0$ and $Hx \le h$, one has $z^T Hx \ge z^T h$. Because $u \le 0$ and $x \le \mathbf{1}$, one has $u^T x \ge u^T \mathbf{1}$. Combining these bounds proves the claim. Integer feasible points form a subset of the relaxation's feasible set, so equality proves optimality in both problems. □

The numerical LP solver proposes multipliers. The checker converts them to rational numbers and explicitly checks the displayed inequalities. If rounding a proposed multiplier changes feasibility, the coordinate correction $u_j = \min(0, c_j - (A^T y + H^T z)_j)$ restores the required inequality. The resulting exact lower bound for the recorded example is 6. Exhaustive enumeration of the allowed two-color vocabularies also gives cost 6. The equality certifies this particular finite example. It does not assert that the relaxation is integral for arbitrary corpora.

18.3 A finite Weyl encoder with trainable gates

18.3.1 The noncommutative algebra is explicit

Let $q \ge 2$, let $\omega = \exp(2\pi i/q)$, and index the standard basis of $\mathbb{C}^q$ cyclically. Define

$$Ue_j = \omega^j e_j, \qquad Ve_j = e_{j-1}.$$

The backward-shift convention fixes the sign in the commutation relation. These are the finite matrices used as an algebraic ingredient in the reviewed ToricGT module.[147]

Proposition 18.5 (Finite Weyl relations and matrix basis)**.** *The operators U, V are unitary, $U^q = V^q = I$, and $VU = \omega UV$. The q^2 matrices $U^r V^s$, $0 \le r, s < q$, are an orthogonal basis of $M_q(\mathbb{C})$ for the Hilbert–Schmidt inner product.*

Proof. The diagonal entries of U have modulus one and V permutes an orthonormal basis, so both are unitary. Their qth powers are the identity. On e_j the two products satisfy $VUe_j = \omega^j e_{j-1}$ and $UVe_j = \omega^{j-1} e_{j-1}$, proving the relation, including the cyclic case $j = 0$ because $\omega^q = 1$.

If $s \not\equiv 0 \pmod q$, $U^r V^s$ has zero diagonal and trace zero. If $s = 0$, its trace is $\sum_{j=0}^{q-1} \omega^{rj}$, equal to zero unless $r = 0$, when it equals q. Indeed, for $r \neq 0$ the geometric sum multiplied by $1 - \omega^r$ is $1 - \omega^{rq} = 0$ and $1 - \omega^r \neq 0$. Using the commutation relation, the product $(U^r V^s)^* U^{r'} V^{s'}$ is a nonzero scalar times $U^{r'-r} V^{s'-s}$. Its trace is therefore zero for distinct pairs and is q for identical pairs. There are q^2 nonzero orthogonal matrices in a q^2-dimensional vector space, which proves the basis assertion. □

This finite algebra has a perfectly ordinary classical matrix representation. The word "noncommutative" establishes neither quantum hardware usage nor a computational speedup. In the implementation $q = 8$; one vector of dimension eight is represented by three qubits, not by eight qubits or by three independent eight-level systems. A finite representation at a rational phase parameter also does not by itself approximate every property of an irrational noncommutative torus.

18.3.2 Circuit and measurements

For a token feature vector $x \in \mathbb{R}^8$, the input encoder computes six angles

$$a(x) = \pi \tanh(W_a x + b_a).$$

There are two circuit layers. In layer $\ell \in \{0, 1\}$, for each $r \in \{0, 1, 2\}$ apply, in this order,

$$R_Y(a_r + \theta_{\ell r 0}), \qquad R_Z(a_{r+3} + \theta_{\ell r 1}),$$

to qubit r, then apply controlled-NOT gates $0 \to 1$ and $1 \to 2$. The initial state is $|000\rangle$ and qubit zero is the most significant bit. There are twelve learned offsets $\theta_{\ell r b}$. The six input angles are reused in both layers; this matters when differentiating the input encoder.

For the resulting unit vector $\psi(x; \theta)$ define

$$f(x; \theta) = (\Re\langle U\rangle, \Im\langle U\rangle, \Re\langle V\rangle, \Im\langle V\rangle, \Re\langle UV\rangle, \Im\langle UV\rangle), \quad \langle T\rangle = \psi^* T \psi.$$

These are real features. For a unitary T, its real and imaginary expectations are expectations of the Hermitian operators $(T + T^*)/2$ and $(T - T^*)/(2i)$. Thus the feature definition has a measurement interpretation, although estimating several noncommuting observables generally requires separate state preparations. The present experiments calculate expectations in a classical statevector simulator; they do not include shot noise or noisy physical gates.

Lemma 18.6 (Normalization and feature bounds)**.** *The displayed circuit preserves norm. Each of its six real features lies in* $[-1, 1]$*. Moreover, for unit vectors* ψ, φ *and an operator* T *with* $\|T\| \leq 1$*,*

$$|\psi^* T \psi - \varphi^* T \varphi| \leq 2\|\psi - \varphi\|.$$

Proof. Both Pauli rotations are unitary and controlled-NOT is a permutation of an orthonormal basis; their product preserves norm. Cauchy–Schwarz gives $|\psi^* T \psi| \leq \|\psi\| \, \|T\psi\| \leq 1$ for the three unitary operators under consideration. Real and imaginary parts have no larger absolute value. Finally write the expectation difference as

$$(\psi - \varphi)^* T \psi + \varphi^* T (\psi - \varphi)$$

and apply the triangle inequality and Cauchy–Schwarz to the two summands. □

18.3.3 What the parameter-shift formula actually differentiates

The familiar two-evaluation rule applies to an expectation whose parameter occurs in one gate with a two-eigenvalue generator. The general analytic gradient framework is developed by Schuld and collaborators. [148] The short proof here fixes the normalization and the composition rule needed by this implementation.

Proposition 18.7 (Single-occurrence Pauli parameter shift)**.** *Let $P = P^*$, $P^2 = I$, and $R(t) = \exp(-itP/2)$. For fixed state η and Hermitian H, put $F(t) = \eta^* R(t)^* H R(t)\eta$. Then*

$$F'(t) = \tfrac{1}{2}\{F(t + \pi/2) - F(t - \pi/2)\}.$$

The same identity holds for each real feature above when the varied offset occurs in exactly one gate of that token's circuit.

Proof. The identity $P^2 = I$ splits the exponential series into even and odd powers: $R(t) = I\cos(t/2) - iP\sin(t/2)$. On expanding R^*HR, the coefficients of H, PHP, and $i(PH - HP)$ are respectively $\cos^2(t/2)$, $\sin^2(t/2)$, and $\sin(t/2)\cos(t/2)$. Consequently $F(t) = A + B\cos t + C\sin t$ for real constants A, B, C. Its derivative is $-B\sin t + C\cos t$, and the displayed shifted difference has exactly this value. Gates preceding R can be absorbed into η; gates following it can be absorbed by conjugating H. Each real or imaginary Weyl feature uses a Hermitian observable as explained above, so the same argument applies to each feature. □

Proposition 18.8 (Hybrid gradient with shared parameters)**.** *Let $f_{tj}(\theta)$ be feature j of token t and let a differentiable classical computation give loss $L = \ell((f_{tj})_{t,j})$. If offset θ_k occurs once in each token circuit, then*

$$\frac{\partial L}{\partial \theta_k} = \sum_{t,j} \frac{\partial \ell}{\partial f_{tj}} \frac{f_{tj}(\theta_k + \pi/2) - f_{tj}(\theta_k - \pi/2)}{2},$$

where each difference changes only that occurrence in that token circuit and the classical derivatives are evaluated at the original features. If a parameter occurs in several gates, sum these contributions over its occurrences before applying its outer chain rule.

Proof. Temporarily replace every gate occurrence by an independent variable. The ordinary multivariable chain rule expresses the derivative of L as the sum of its feature adjoints multiplied by the feature derivatives. Proposition 18.7 evaluates each such derivative. Identifying occurrence variables with the shared parameter sums their derivatives, again by the chain rule. This proves both statements. □

It would generally be incorrect to replace the derivative of the entire nonlinear loss by half the difference of two shifted losses. For example, if $f(t) = \cos t$ and $L(t) = f(t)^2$, those two shifted losses are equal, whereas $L'(t) = -\sin(2t)$ is usually nonzero. The graph attention and expert networks are nonlinear functions of the features. Our verification therefore compares their feature adjoints composed with independent parameter-shift Jacobians against automatic differentiation of the full loss. At max-plus ties a classical derivative may fail to exist; the implementation uses its framework's subgradient convention. The analytic identity above is asserted at differentiability points, not as a unique derivative at every tie.

The circuit offsets each have a single occurrence per token. An input angle occurs once in each of the two layers, so its derivative has two occurrence contributions. The input-to-angle map then contributes its ordinary derivative $\pi\,\mathrm{sech}^2(W_a x + b_a)$. Automatic differentiation performs this composition during training; the independent shift check specifically audits the shared circuit offsets.

18.4 A graph transformer whose symmetries can be stated precisely

18.4.1 Tokens and incidence

For a vertex i, use endpoint pair (i, i); for an edge $i \to j$, use (i, j). Each token has eight scalar attributes: vertex and edge indicators, the two byte positions divided by twelve, token span divided by twelve, cost divided by thirty-two, the bounded control divided by two, and the edge's cost slope divided by two. Inapplicable entries are zero. The endpoint-incidence matrix is

$$R_{tu} = \tfrac{1}{4} \sum_{a,b=1}^{2} \mathbf{1}\{\mathrm{end}_a(t) = \mathrm{end}_b(u)\}.$$

Equivalently replace the indicator by the inner product of orthonormal endpoint vectors. This encoding uses equality structure, not a coordinate choice in the identifier space. Vertex and edge tokenization has a primary antecedent in TokenGT; its general expressivity results involve additional architecture hypotheses and are not claimed for this small model. [77]

The initial hidden table is

$$X_t = W_x x_t + b_x + W_f f(x_t; \theta) + b_f \in \mathbb{R}^{24}.$$

It has two graph blocks and two attention heads per block. Each block normalizes token rows, forms learned queries, keys, and values, and centers and normalizes each query and key before forming scores. For one head the scores have the form

$$S_{tu} = \frac{q_t^T k_u}{\sqrt{12}} + \rho R_{tu}, \qquad \rho > 0.$$

The two aggregation operations are

$$A^{\mathrm{soft}}_{tr} = \sum_u \frac{e^{S_{tu}}}{\sum_v e^{S_{tv}}} v_{ur}, \qquad A^{\max}_{tr} = \max_u (S_{tu} + v_{ur}).$$

A learned sigmoid gate mixes them. The head outputs are concatenated and mapped back to width twenty-four with a residual connection. The graph here is fully observed; no autoregressive next-byte probability is being reported. A causal language model would require a separate reveal-order mask and proof that endpoint attributes expose no future information.

Lemma 18.9 (Associativity of finite max-plus contraction)**.** *For finite compatible matrices with entries in $\mathbb{R} \cup \{-\infty\}$, define $(A \odot B)_{ik} = \max_j (A_{ij} + B_{jk})$, with the usual convention for $-\infty$ and $\max \varnothing = -\infty$. Then $(A \odot B) \odot C = A \odot (B \odot C)$. Splitting the contracted index set into blocks and taking the maximum of the block maxima gives the same result.*

Proof. Both parenthesizations at entry (i, l) equal $\max_{j,k}(A_{ij} + B_{jk} + C_{kl})$. Finite maxima can be taken in either order. A maximum over a finite disjoint union is the maximum of the maxima over its pieces, proving the second statement. □

This justifies an exact blockwise max-plus reduction when its scores and values are already fixed. It does not make the whole graph network piecewise affine in its raw input or in its parameters: normalized queries and keys, their dot products, quantum expectations, softmax, layer normalization, and GELU activations are nonlinear.

18.4.2 Soft experts and a permutation proof

For an $m \times d$ token table X and eight learned slot keys, form logits L_{ts} by scaled cosine similarity. In the implementation each vector v is normalized as $v/\max(\|v\|_2, 10^{-6})$, including the zero vector; this convention is also used for the centered queries and keys. Row layer normalization uses its population variance plus 10^{-5} under the square root. These positive constants make every displayed row operation defined at constant or zero rows. Define two differently normalized arrays

$$D_{ts} = \frac{e^{L_{ts}}}{\sum_{u=1}^{m} e^{L_{us}}}, \qquad C_{ts} = \frac{e^{L_{ts}}}{\sum_{r=1}^{8} e^{L_{tr}}}.$$

The slot inputs are $\widetilde{X} = D^T X$. Two slots are assigned to each of four learned feed-forward experts, giving output slots $\widetilde{Y}$. The token output is $C\widetilde{Y}$, multiplied by a learned residual coefficient. Dispatch normalizes over tokens; combination normalizes over slots. Confusing those axes changes the model. [124]

Proposition 18.10 (Equivariance of the implemented mathematical architecture)**.** *Let P be a permutation matrix on the valid tokens of a nonempty graph. Transport token attributes and incidence by $X \mapsto PX$ and $R \mapsto PRP^T$. Every displayed graph block and the rowwise readout commute with this action. Thus the predicted node values are permuted along with the corresponding node tokens. A coordinated renaming of endpoint identifiers also leaves the incidence data unchanged.*

Proof. Rowwise affine maps, rowwise normalizations, coordinatewise activations, and the identical token circuit commute with left multiplication by P. Queries, keys, and values therefore transform to PQ, PK, PV, so their score matrix transforms to PSP^T. Row softmax transforms by the same two-sided permutation, and its product with the permuted values is the permutation of the original output. For max-plus attention, relabeling the finite maximizing index merely reorders the same candidates, giving the same conclusion. Residual addition and concatenation preserve it.

For the expert layer, L transforms to PL. Consequently dispatch and combination transform to PD and PC. The slot inputs are unchanged: $(PD)^T(PX) = D^T X$. Each expert acts on those unchanged slots, so the output slots are unchanged. Multiplication by PC permutes the final token outputs. Induction over the blocks and the final rowwise readout proves the assertion. Finally, a bijective renaming preserves every endpoint equality appearing in R, which proves the last statement. □

The claim concerns attributed graphs. Changing a byte position while keeping a token in place changes an attribute and can change the answer. Likewise, erasing edge direction or swapping the semantic source and target is not the token-storage permutation in the proposition. Masks must be transported too when padding or causal restrictions are present. In exact mathematics, padded keys are absent from sums and maxima. The implementation checks the relevant finite-precision masking behavior numerically, including insertion of padding and token permutations.

18.5 Toric path relations, graph rewrites, and steering regions

18.5.1 A small commutative algebra with an exact graph meaning

Let G be a finite directed graph and give its edge variables the incidence degrees $d_e = e_{\mathrm{target}(e)} - e_{\mathrm{source}(e)} \in \mathbb{Z}^V$, the columns of a matrix D. Thus $D = -B$ for the flow matrix above, which does not change the kernel. Consider the algebra map

$$\Phi : k[z_e : e \in E] \longrightarrow k[t_v^{\pm 1} : v \in V], \qquad z_e \longmapsto t_{\mathrm{target}(e)}/t_{\mathrm{source}(e)}.$$

For a nonnegative exponent vector u, its image has exponent vector Du with this incidence convention. Multiplication forgets the order in which edges were listed; it retains their net boundary.

Proposition 18.11 (The binomial kernel of the boundary code)**.** *The kernel of Φ is the ideal generated by all $z^u - z^v$ with $u, v \in \mathbb{N}^E$ and $Du = Dv$. In particular, two directed paths with the same initial and final vertices have the same image under Φ.*

Proof. Every displayed binomial maps to zero. Conversely write a polynomial in the kernel as a finite sum $\sum_u c_u z^u$ and group its terms according to the integer vector Du. Distinct Laurent monomials are linearly independent over k. Thus within each group the coefficients sum to zero. Choose one exponent u_0 in a nonempty group. That group's contribution equals $\sum_{u \neq u_0} c_u (z^u - z^{u_0})$, a sum of the stated binomials. Adding the finitely many groups proves the kernel assertion. Along a path, the incidence columns telescope to $e_t - e_s$, proving the final claim. □

For three arrows $a : 0 \to 1$, $b : 1 \to 2$, and $c : 0 \to 2$, the two complete paths give the relation $z_a z_b - z_c$. In this example the whole kernel is $(z_a z_b - z_c)$. One may eliminate z_c, leaving the polynomial ring in z_a, z_b; their incidence vectors are linearly independent, so their images satisfy no further polynomial relation. The unit-flow polytope is

$$(f_a, f_b, f_c) = (\lambda, \lambda, 1 - \lambda), \qquad 0 \leq \lambda \leq 1.$$

The two integer endpoints are the two tokenizations. For edge costs c_a, c_b, c_c, the preferred endpoint changes on the wall $c_a + c_b = c_c$. The same comparison determines which side of the path binomial is selected by the corresponding weight. This is an explicit connection among a graph problem, a polytope, a binomial ideal, and a piecewise-linear optimal value function.

The kernel relation alone does not certify an arbitrary sequence of rewrite operations. The vector $u - v$ may have negative entries, and applying it to a current state may leave the nonnegative orthant. Even nonnegative boundary-preserving changes on a cyclic graph can add a disconnected circulation. The DAG and unit-flow hypotheses in Lemma 18.3 are what turn a nonnegative solution with the correct boundary into one legal path.

Proposition 18.12 (A constrained rewrite preserves legal tokenization)**.** *On the ordered DAG of Lemma 18.3, let x be a legal path flow and let $\delta \in \mathbb{Z}^E$ satisfy $B\delta = 0$. If $x + \delta \geq 0$ coordinatewise, then $x + \delta$ is another legal path flow with the same endpoints. Its cost change is exactly $c^T \delta$.*

Proof. Linearity gives $B(x + \delta) = Bx + B\delta = e_s - e_t$. The assumed nonnegative integral vector therefore satisfies Lemma 18.3 and is a path. The cost identity follows by distributivity: $c^T(x + \delta) - c^T x = c^T \delta$. □

A graph grammar can use such a local replacement when its interface and nonnegativity tests are explicit. If declared path relations are made columns of a map $\partial_2 : C_2 \to C_1$, their zero boundary gives $\partial_1 \partial_2 = 0$. A replacement by $\partial_2 w$ preserves the homology class of a cycle. This is a genuine chain-complex connection to the earlier chapters. It is weaker than proving that every algebraically equivalent trajectory is semantically interchangeable. Dependencies, labels, resource limits, and external facts need their own preservation conditions. The commutative algebra records the chosen invariant; it cannot reconstruct information intentionally discarded by abelianization.

18.5.2 Affine margins and nonlinear circuit coordinates

Suppose a finite decision system uses affine scores $\ell_j(z) = a_j^T z + b_j$ on a feature space. A fixed strict sign pattern of these scores determines an open polyhedral region. It may be empty, and there may

be many disconnected regions after pulling it back through a nonlinear encoder. The affine-space statement is the useful part of the steering proposal; identifying a region with a scientifically correct reasoning mechanism would require additional evidence.

Proposition 18.13 (Preserving a finite affine sign pattern)**.** *Let $s_j \in \{-1, 1\}$ and suppose $s_j \ell_j(z) \geq \gamma > 0$ for every j. If $|a_j^T \Delta z| < \gamma$ for every j, then $s_j \ell_j(z + \Delta z) > 0$ for every j. In particular it suffices that $\|\Delta z\|_2 < \gamma / \max_j \|a_j\|_2$ when the denominator is positive. If every $a_j = 0$, every perturbation preserves the signs.*

Proof. For each j, $s_j \ell_j(z + \Delta z) = s_j \ell_j(z) + s_j a_j^T \Delta z \geq \gamma - |a_j^T \Delta z| > 0$. Cauchy–Schwarz gives the sufficient norm bound. If all coefficients vanish, all scores are constant. □

The proof applies, for example, to a finite list of path-cost comparisons when the edge costs depend affinely on z. To use it in circuit parameters one also needs a bound on the feature change. Here is a directly proved bound for the twelve independent offsets, with the token input held fixed.

Lemma 18.14 (A conservative circuit-offset perturbation bound)**.** *For the token circuit above and two offset vectors θ, θ',*

$$\|\psi(\theta') - \psi(\theta)\| \leq \tfrac{1}{2} \|\theta' - \theta\|_1, \qquad \|f(\theta') - f(\theta)\|_2 \leq \sqrt{6}\, \|\theta' - \theta\|_1.$$

Proof. For a Pauli rotation the derivative has operator norm $1/2$. Integrating along the real angle interval gives $\|R(t') - R(t)\| \leq |t' - t|/2$. Subtract two products of unitary gates by replacing one factor at a time. The operator norm of each difference term is at most the norm difference of that one factor, because all other factors are unitary. The sum over the twelve independently varied gates proves the state bound. Lemma 18.6 then bounds the absolute change of each of the six features by $\|\theta' - \theta\|_1$. Summing their squared bounds proves the second inequality. □

Combining these propositions gives a sufficient margin condition for an *affine readout directly on the six features*. It does not give the same constant for the full transformer. The latter's nonlinear layers and learned weights contribute further Lipschitz factors; max-plus ties also require the appropriate nonsmooth interpretation. Moreover a polyhedron in expectation-feature space need not have a polyhedral preimage in angle space. Even the one-qubit feature $\cos\theta$ has periodic, nonlinear level sets. This qualification prevents a useful local sign guarantee from becoming an unsupported global geometry claim.

18.6 Three source corrections with complete small counterarguments

The reviewed ToricGT V text contains several elementary interfaces that are useful when their hypotheses are retained. The following localized qualifications concern its saturation example, a sign in its predictive transfer definition, and a proposed monotonicity interpretation of coverage homology. The original pinned sources are preserved. The review ledger records the exact commit, path, and line ranges; the corrections here are new mathematical explanations, not silent edits to that source. [144]

18.6.1 Saturation is relative to the generated lattice

Proposition 18.15 (The parity semigroup is normal in its own group)**.** *Let* $Q = \mathbb{N}(2,0) + \mathbb{N}(0,2) + \mathbb{N}(1,1) \subseteq \mathbb{Z}^2$*. Then*

$$\operatorname{gp}(Q) = \{(a,b) \in \mathbb{Z}^2 : a+b \text{ is even}\}, \qquad Q = \operatorname{gp}(Q) \cap \mathbb{R}^2_{\geq 0}.$$

In particular Q *is saturated in* $\operatorname{gp}(Q)$*, and* $(1,0)$ *is not a saturation hole in that group.*

Proof. Each generator has even coordinate sum, so every integer combination does too. Conversely if $a+b$ is even, then $a-b$ is even and $(a,b) = ((a-b)/2)(2,0) + b(1,1)$ is an integer combination of generators. This proves the group description.

For nonnegative (a,b) with even sum, the coordinates have the same parity. If both are even, use $(a/2)(2,0) + (b/2)(0,2)$. If both are odd, use $((a-1)/2)(2,0) + ((b-1)/2)(0,2) + (1,1)$. All coefficients are nonnegative integers, proving the intersection formula. If $m(a,b) \in Q$ with $m > 0$ and $(a,b) \in \operatorname{gp}(Q)$, then $a, b \geq 0$, so the formula implies $(a,b) \in Q$. Finally $(1,0)$ has odd coordinate sum and is outside the generated group. □

The fact that $2(1,0) \in Q$ does not contradict normality: its antecedent omits membership in $\operatorname{gp}(Q)$. If one instead studies the larger ambient lattice $\mathbb{Z}^2$, that is a different lattice specification. This is the same distinction that matters when proving normality of the semigroup degenerations in the standard-monomial chapters.

18.6.2 Inclusions do not make Betti numbers monotone

Proposition 18.16 (A point, a circle, and a filled triangle)**.** *There exist inclusions of finite simplicial complexes* $K_0 \subset K_1 \subset K_2$ *for which the first Betti numbers over* $\mathbb{Q}$ *are* $0, 1, 0$*, respectively. Thus inclusion alone implies neither nondecreasing nor nonincreasing first Betti numbers.*

Proof. Take K_0 to be vertex 0, K_1 the boundary of the triangle on vertices $0, 1, 2$, and K_2 the filled triangle. In K_1, orient the edges $0 \to 1$, $1 \to 2$, $2 \to 0$. The boundary matrix is

$$\partial_1 = \begin{pmatrix} -1 & 0 & 1 \\ 1 & -1 & 0 \\ 0 & 1 & -1 \end{pmatrix}.$$

Its first two columns are independent and its three columns sum to zero, so its rank is two and its kernel is spanned by $(1,1,1)^T$. There are no two-simplices in K_1, giving $\dim H_1(K_1;\mathbb{Q}) = 1$. In K_2, orient the added triangle so that $\partial_2(1) = (1,1,1)^T$. This map has rank one and its image equals $\ker \partial_1$; hence $H_1(K_2;\mathbb{Q}) = 0$. The point has no one-chains, so $H_1(K_0;\mathbb{Q}) = 0$ as well. □

A growing coverage complex can create a class and later kill it. A statement that a support construction grows under evidence addition may be correct while a monotone "hole penalty" conclusion is not justified. One must specify induced homology maps, persistence intervals, and what the complexes represent. Even a genuine persistent class is a property of that representation, not a theorem that a scientific assertion is false or that all relevant evidence has been gathered.

18.6.3 Use one sign for predictive gain

For normalized predictive distributions with positive probabilities on the observed symbols, define the unaugmented codelengths

$$C_0 = -\sum_t \log_2 q_0(b_t \mid b_{<t}), \qquad C_1 = -\sum_t \log_2 q_1(b_t \mid b_{<t}).$$

Their gain is

$$C_0 - C_1 = \sum_t \log_2 \frac{q_1(b_t \mid b_{<t})}{q_0(b_t \mid b_{<t})}.$$

If a deployed auxiliary costs B additional bits, the total codelength gain is $C_0 - C_1 - B$. Positive gain means improvement with this convention. For a single observed symbol with $q_0 = 1/4$ and $q_1 = 1/2$, the unaugmented gain is one bit. Reversing the ratio gives minus one bit. Artifact-adjusted gain and ordinary predictive bits per byte should be named separately, rather than changing the metric's definition mid-proof. These identities say nothing about the task accuracy of a neural model unless the probabilities and evaluation protocol have been supplied.

The companion's `algebra_checks.py` verifies the displayed boundary ranks and zero composition exactly over $\mathbb{Q}$, checks the parity decomposition on 545 bounded pairs, and records the rational probability example. The general semigroup claim follows from its proof above, not from finite enumeration. The finite checks provide reproducible illustrations and catch transcription mistakes.

18.7 Completed learning experiment and numerical evidence

18.7.1 Protocol and comparators

The fixed dataset contains 512 training graphs and 128 graphs in each of the validation, test, and length-shift sets. All underlying strings are distinct across the four sets. Training, validation, and ordinary test strings have lengths four through eight; the length-shift set has lengths nine through twelve. Token availability and integer cost parameters are sampled from the recorded generator. Each byte fallback costs sixteen integer units. In the following tables one unit of reported cost equals that byte-fallback cost, so integer teacher costs are divided by sixteen. The network itself predicts distances divided by thirty-two.

All three model variants use the same graph encoder layout, width twenty-four, two attention blocks, four experts per block, and scalar node readout. The trainable quantum and classical variants each contain 25,515 trainable parameters. The classical feature comparator uses the same six encoded angles and twelve phase parameters, with bounded features

$$g_j(a;p) = \tfrac{1}{2}\big(\sin(a_j + p_j) + \cos(a_{j-1} + p_{j+6})\big), \qquad j \pmod 6.$$

The other input and graph layers are identical. This is a matched parameter-count comparator, not a claim that it exhausts possible classical architectures. In the frozen-circuit ablation, both the 54-parameter angle encoder and twelve circuit offsets are frozen, while the remaining 25,449 parameters are trained.

Each variant is trained at seeds $17, 29, 43$ for eighty epochs, batch size thirty-two, AdamW learning rate 0.0015, weight decay 10^{-4}, and a cosine learning-rate schedule ending at one tenth of its initial value. The loss is mean squared error over valid node predictions, with padded and edge-token targets excluded. The same initialization seed and minibatch ordering are used for the corresponding variants. For each run, the checkpoint with lowest validation node loss is selected

before its test and length-shift metrics are computed. The representative browser checkpoint is the quantum run with lowest validation loss, not a checkpoint chosen for its test score. A preceding two-epoch smoke run checked execution and export; its results were not candidate checkpoints for the reported experiment.

The runs used PyTorch 2.8.0 with CUDA 12.8 on an NVIDIA GeForce RTX 4090. They were classical GPU statevector simulations with analytic expectations and automatic differentiation, not quantum-hardware experiments and not cuStateVec training. The existing environment was used without changing its packages. Process allocation was capped at one tenth of GPU memory. The largest measured PyTorch allocation during the nine runs was 183,426,048 bytes, about 175 MiB; this allocator measurement does not include every CUDA-context allocation or the memory of other applications. The full set of runs completed in approximately 109 seconds.[36]

18.7.2 Results, including the negative evidence

Table 18.1 reports means and sample standard deviations over the three seeds. "Optimal" refers to the original learned proposal. The initial certificate is stricter: the proposal must attain the lower bound produced from its rounded, repaired potential. A proposal may be optimal even when that particular potential does not certify it. All final paths pass exact certification after the separately reported dynamic-programming repair where required.

Set	Variant	Node RMSE	Optimal (%)	Initial cert. (%)
Test	Trainable quantum	0.265 ± 0.003	83.85 ± 2.39	24.74 ± 0.45
Test	Classical	0.266 ± 0.009	84.38 ± 0.78	30.47 ± 5.63
Test	Frozen quantum	0.284 ± 0.011	83.07 ± 1.97	25.78 ± 5.12
Longer	Trainable quantum	0.684 ± 0.150	54.69 ± 3.41	6.51 ± 3.85
Longer	Classical	0.790 ± 0.117	50.00 ± 4.75	5.47 ± 3.41
Longer	Frozen quantum	0.676 ± 0.069	52.86 ± 0.90	9.90 ± 2.26

Table 18.1: Measured learning outcomes on the fixed token-DAG dataset. Uncertainty entries are sample standard deviations over three seeds, not confidence intervals or statistical significance claims.

The mean proposal cost regret on the ordinary test set is 0.0386 for the trainable quantum model, 0.0356 for the classical comparator, and 0.0399 for frozen quantum features. On longer strings these means are 0.2246, 0.2625, and 0.2324, respectively. The mixed ordering is substantive evidence: the classical comparator has slightly better ordinary-test path accuracy and regret, while the trainable quantum model has somewhat better mean path accuracy and regret under this length shift. Frozen quantum features have a slightly lower mean length-shift node RMSE than the trainable circuit. Three seeds on one synthetic task do not support a general superiority conclusion for any of the variants.

The twelve quantum offsets change by Euclidean norms 0.146, 0.151, and 0.158 at the selected checkpoints, while the corresponding changes in the frozen ablation are exactly zero. Initial gradients for the circuit offsets and angle encoder are recorded. This establishes that the new circuit feature map participates in training. It does not show that every parameter is identifiable, that all gradients are equally useful, or that optimization found a global minimum.

A batch of thirty-two graph forwards takes approximately 2.56 ms for the trainable quantum variant and 1.81 ms for the classical comparator on the measured GPU, after warm-up and synchronization. These are small model-forward timings, excluding exact path repair and data

preparation. The exact dynamic program is an algorithmic baseline with zero regret; the model comparison does not establish an end-to-end speed advantage over it. No inference about a physical quantum device follows from these classical timings.

18.7.3 Independent checks and browser execution

The circuit is implemented three ways: indexed PyTorch tensor operations, independent dense NumPy Kronecker matrices, and a dependency-free JavaScript statevector reference. The GPU verification uses double precision and reports a maximum PyTorch–NumPy state difference of $3.73 \cdot 10^{-16}$, feature difference of $4.45 \cdot 10^{-16}$, and automatic-differentiation versus parameter-shift Jacobian difference of $3.34 \cdot 10^{-16}$ on the recorded cases. The full nonlinear-loss chain-rule check differs by $1.12 \cdot 10^{-16}$. These are observed numerical agreement bounds on finite tests, not formal floating-point error theorems.

The same verification includes forty-eight independent exhaustive path enumerations, exact rational LP-dual checks, token permutation and padding checks, and the two normalization axes of the expert router. Finite differences provide an additional numerical gradient check, but the analytic parameter-shift identity is the mathematical reference. All source and dataset hashes associated with the training runs are retained. Checkpoints use ordinary NumPy array archives, and per-example predictions and certificates are stored in JSON. The complete results of all nine runs are retained, including the runs that do not favor the quantum variant.

The browser runs the whole selected graph model. Its `runGraph` function accepts the graph, exported weights and optional circuit offsets. It recomputes input angles, circuit expectations, both attention blocks, expert routing, node predictions, and the candidate path. The optional offset tensor has shape $2 \times 3 \times 2$. It changes the circuit offsets for every token and propagates the resulting features through the full model. The interface then checks the candidate and, if necessary, displays the separate exact repair. Recorded predictions are not substituted for the live computation.

JavaScript uses double arithmetic and an approximation to the error function in GELU. Against saved GPU predictions on two ordinary-test and two length-shift graphs, the largest discrepancy in normalized node prediction is $6.81 \cdot 10^{-7}$. The four learned paths and final exact certificates agree. Increasing one circuit offset by 0.4 radians changes the full model's predictions by as much as 0.555 in the checked example; the interactive control therefore changes the actual model computation. This comparison is recorded with the hashes of the JavaScript modules and exported weights.[36]

18.8 Connections that follow, and research questions that remain

The finite Weyl algebra is the strongest direct connection to quantum algorithm theory in this experiment. Clock and shift matrices encode finite translation and character operations. Similar algebra is central to finite Fourier analysis. Here, however, their expectations are learned features of a very small circuit. No hidden subgroup is identified, no elliptic-curve discrete logarithm is extracted, and no improvement to Shor's complexity follows. The earlier cryptographic chapters identify the additional problem-specific structures required for those claims.

The binomial path kernel connects more directly to the standard-monomial and semigroup chapters. It is an actual quotient map with a proved kernel, and its simplest path comparison is an actual initial-term wall. The parity example shows why the lattice used for normalization must be specified. A hidden feature called a "toric coordinate" is not by itself a point on a chosen toric variety; one must supply the monomial map, relations, and any positivity or saturation hypotheses. Our finite path construction supplies those data in an elementary setting.

Graph rewrites connect to complexes through the condition $\partial_1\partial_2 = 0$ and to learning through proposal and acceptance. A trainable model may suggest a boundary-preserving replacement; exact checks enforce the interface and nonnegativity constraints. This is a practical division of labor for a graph grammar. Extending it to arbitrary typed hypergraphs requires the appropriate matching and gluing conditions, which are developed separately in the graph-rewrite chapter. It is not justified by a count-vector identity alone.

There is also a useful possible quantum extension whose theorem can be stated independently of its implementation. It makes precise what it would mean for a variational state to remain within a declared feasible space. It is not the circuit trained in the preceding experiment.

Proposition 18.17 (A constraint-preserving variational interface)**.** *Let Ω be a finite set, $F \subseteq \Omega$ nonempty, and Π_F the orthogonal projector onto the span of basis vectors $|x\rangle$ with $x \in F$. Let ψ_0 be a unit vector in this subspace. If every gate G_j is unitary and commutes with Π_F, the resulting state $\psi = G_m \cdots G_1\psi_0$ remains in that subspace. Measurement in the indicated basis returns an element of F with probability one. For a real cost function c on F, put $H_c = \sum_{x\in F} c(x)|x\rangle\langle x|$ on this subspace and $c_* = \min_{x\in F} c(x)$. Then $\langle H_c\rangle \geq c_*$. If every nonoptimal cost is at least $c_* + \Delta$ for some $\Delta > 0$, the probability of a nonoptimal outcome is at most $(\langle H_c\rangle - c_*)/\Delta$.*

Proof. Commutation gives $\Pi_F G_j\eta = G_j\Pi_F\eta = G_j\eta$ whenever $\Pi_F\eta = \eta$. Induction proves subspace preservation, and unitarity preserves norm. Writing $\psi = \sum_{x\in F} \alpha_x|x\rangle$ gives $\sum_x |\alpha_x|^2 = 1$ and $\langle H_c\rangle = \sum_x c(x)|\alpha_x|^2$. This is a convex combination of costs and is at least c_*. Subtracting c_* and retaining only nonoptimal terms yields

$$\langle H_c\rangle - c_* \geq \Delta \sum_{x:c(x)>c_*} |\alpha_x|^2,$$

which is the claimed probability bound. □

To make this a useful algorithm one must efficiently encode F, prepare ψ_0, construct expressive gates commuting with its projector, and estimate the cost. Those are substantial additional tasks. In particular, the three-qubit feature circuit in this chapter has no basis state for every tokenization path of an arbitrary graph. Its unitarity certifies normalization, while the separate graph checker certifies tokenization. Confusing these two invariants would erase the main proof boundary.

Finally, the finite graph and tensor operations do not by themselves establish a continuum gauge theory, an elliptic index theorem, a positive Yang–Mills mass gap, or a cosmological model. They can provide finite data structures and exactly checked algebraic interfaces for such research. A field-theoretic or cosmological conclusion additionally requires the physical interpretation, limiting procedure, analytic hypotheses, and observational tests appropriate to that problem. The cosmology and gauge-theory chapters retain those requirements. The achievement established here is narrower and concrete: a source-grounded trainable quantum graph feature model, explicit mathematical guarantees for its finite interfaces, measured comparator results, and a live reproducible implementation.

Part V

Quantum fields and a precise research synthesis

Chapter 19

Gauge fields, curvature, and the meaning of a complex

19.1 The precise connection to the preceding algebra

The theory of orders concerns modules carrying a multiplication action. The theory of complexes concerns maps whose consecutive compositions vanish. Gauge theory adds a third ingredient: the maps are differential operators depending on a connection, and the obstruction to their composing to zero is curvature. This observation supplies a concrete bridge between the subjects, but it also identifies the hypothesis that any proposed bridge must preserve. An arbitrary covariant exterior derivative is not a differential in a cochain complex. Its square is a curvature action. A flat connection removes that obstruction; an anti-self-dual connection removes the appropriate projected obstruction. Neither operation turns a classical deformation problem into a construction of the quantum vacuum.

The following derivations fix conventions sufficiently explicitly to permit comparison with both algebraic complexes and formal proofs. The foundational gauge theory dates to Yang and Mills; the geometric analysis of the anti-self-dual deformation complex is due to Atiyah, Hitchin, and Singer [15, 181]. The proofs below are elementary derivations of the identities used here. The analytic index formula and reconstruction theorems are identified separately as external results.

Let M be a smooth, oriented Riemannian manifold without boundary and let G be a compact Lie group. Let $P \to M$ be a principal G-bundle and write $\operatorname{ad} P = P \times_{\mathrm{Ad}} \mathfrak{g}$. Choose an Ad-invariant positive definite inner product on $\mathfrak{g}$. If G is realized as a compact matrix group in $U(m)$, a positive multiple of $-\operatorname{Re}\operatorname{Tr}(XY)$ is an example on its skew-Hermitian Lie algebra. The numerical multiple matters when comparing coupling constants, instanton charges, and short-distance coefficients. It does not change the differential identities.

The affine space of smooth connections has tangent space $\Omega^1(M; \operatorname{ad} P)$. In a local trivialization a connection is a Lie-algebra-valued one-form A, with

$$F_A = dA + \tfrac{1}{2}[A \wedge A], \qquad d_A\alpha = d\alpha + [A \wedge \alpha]. \tag{19.1}$$

For homogeneous forms the bracket includes exterior multiplication and the Lie bracket. In coordinates this means

$$(F_A)_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu], \qquad D_\mu X = \partial_\mu X + [A_\mu, X].$$

The bracket is graded antisymmetric: $[\alpha \wedge \beta] = -(-1)^{pq}[\beta \wedge \alpha]$ for degrees p, q. In particular it is symmetric for two one-forms. Keeping this sign is essential in the second variation.

Proposition 19.1 (Curvature is the failure of the complex identity)**.** *For every adjoint-valued form* α*,*

$$d_A^2\alpha = [F_A \wedge \alpha], \qquad d_A F_A = 0. \tag{19.2}$$

Consequently, if $F_A = 0$*, the sequence* $\Omega^0(\operatorname{ad} P) \xrightarrow{d_A} \Omega^1(\operatorname{ad} P) \xrightarrow{d_A} \cdots$ *is a complex.*

Proof. For a degree-zero section s, direct expansion gives

$$(D_\mu D_\nu - D_\nu D_\mu)s = [\partial_\mu A_\nu - \partial_\nu A_\mu, s] + [A_\mu, [A_\nu, s]] - [A_\nu, [A_\mu, s]] = [F_{\mu\nu}, s],$$

where the last equality is the Jacobi identity. Antisymmetrizing the same calculation in the additional form indices gives the first identity for every degree; differentiation of the scalar form coefficients contributes the ordinary identity $d^2 = 0$. Equivalently, the graded Leibniz rule and graded Jacobi identity give $d_A^2 = [dA + \frac{1}{2}[A \wedge A], \cdot]$ directly. For the second identity, the cyclic sum of $D_\lambda F_{\mu\nu}$ has cancelling second derivatives and cancelling first-derivative brackets. The remaining terms are $[A_\lambda, [A_\mu, A_\nu]]$ and their cyclic permutations, whose sum is zero by Jacobi. This is precisely $d_A F_A = 0$. If $F_A = 0$, the first identity says that every consecutive composition is zero. □

One must not reverse the last implication without a hypothesis on the representation. A central curvature can act trivially in the adjoint representation. For instance an abelian connection may have nonzero curvature although its induced adjoint differential is the ordinary exterior derivative. For a centerless Lie algebra the statement that $[F_A, \phi] = 0$ for every local adjoint section ϕ does force $F_A = 0$. This qualification is a small example of the semantic checks required when a formal statement is used as a bridge between different mathematical settings.

19.2 Gauge covariance and the variational equations

Use the convention $A^u = u^{-1}Au + u^{-1}du$ for a gauge transformation u. Then $d_{A^u}(u^{-1}\alpha u) = u^{-1}(d_A\alpha)u$ and $F_{A^u} = u^{-1}F_A u$. These formulas can be checked in a local matrix representation by differentiating $u^{-1}u = 1$. Since the chosen inner product is invariant, the Yang–Mills functional

$$S(A) = \frac{1}{2g^2}\int_M |F_A|^2 \, d\operatorname{vol} \tag{19.3}$$

is gauge invariant. Here $g > 0$ is fixed. On a noncompact manifold all variations used below have compact support, or sufficient decay to justify integration by parts. On a manifold with boundary a boundary condition must be supplied; omitting that condition changes the Euler–Lagrange problem.

Proposition 19.2 (First and second variations)**.** *For* $a, b \in \Omega^1(M; \operatorname{ad} P)$*,*

$$F_{A+ta} = F_A + t\, d_A a + \tfrac{1}{2}t^2[a \wedge a], \tag{19.4}$$

$$DS_A(a) = g^{-2}\langle d_A^* F_A, a\rangle_{L^2}, \tag{19.5}$$

$$D^2 S_A(a,b) = g^{-2}\left(\langle d_A a, d_A b\rangle_{L^2} + \langle F_A, [a \wedge b]\rangle_{L^2}\right). \tag{19.6}$$

Thus a smooth connection is stationary for all permitted variations if and only if $d_A^* F_A = 0$*.*

Proof. Substitute $A + ta$ into (19.1). The linear terms are $da + [A \wedge a] = d_A a$ because the bracket of one-forms is symmetric. The quadratic term is $\frac{1}{2}t^2[a \wedge a]$. Differentiating $\frac{1}{2}g^{-2}\langle F, F\rangle$ once gives $g^{-2}\langle F_A, d_A a\rangle$, which equals $g^{-2}\langle d_A^* F_A, a\rangle$ by the definition of the formal adjoint. Mixed differentiation of $F_{A+sa+tb}$ gives $\partial_s\partial_t F = [a \wedge b]$. Applying the product rule to the norm square proves the Hessian formula. If the first variation vanishes for every compactly supported a, the fundamental lemma of the calculus of variations gives the stated equation; its converse follows from the formula. □

The equation $d_A F_A = 0$ is an identity, while $d_A^* F_A = 0$ is a field equation. Their different logical roles are important. Imposing the first does not restrict the set of connections, whereas imposing the second selects the classical Yang–Mills critical points. In a Lean development the Bianchi identity should be proved from the definition of curvature, not added as an unexplained assumption that could conceal an inconsistent definition.

The infinitesimal gauge direction is $d_A\phi$, because $\frac{d}{dt}|_{t=0} A^{\exp(t\phi)} = d_A\phi$. At a critical point, the Hessian annihilates these directions in its associated bilinear form. One way to see this without a coordinate calculation is to differentiate the gauge invariance identity twice along a gauge orbit and a transverse variation. The term involving acceleration of the orbit is paired with DS_A, which is zero at a critical point. Thus the degeneracy is forced by symmetry, and the unmodified Hessian cannot be elliptic on all one-forms.

For a nonzero covector ξ, omit the harmless factor i in the symbol convention. The principal symbol of d_A is exterior multiplication by ξ; that of its adjoint is contraction by $\xi^\sharp$, up to the consistent sign for a first-order adjoint. The principal symbol of the nonnegative second-order operator $d_A^* d_A$ on one-forms is

$$a \longmapsto |\xi|^2 a - \xi\, \iota_{\xi^\sharp} a.$$

It vanishes on multiples of ξ. Adding the gauge-fixing term $d_A d_A^*$ changes this to $|\xi|^2 a$. Curvature contributes only lower-order terms to the Hessian. Accordingly the gauge-fixed Hessian is elliptic, although its lower-order part need not be positive at an arbitrary critical point. Ellipticity concerns invertibility of the principal symbol; positivity is a separate property of the complete operator.

19.3 Four dimensions and self-duality

Suppose now that $\dim M = 4$. On two-forms the Hodge star satisfies $*^2 = 1$, and hence

$$\Lambda^2 = \Lambda^2_+ \oplus \Lambda^2_-, \qquad F_A^\pm = \tfrac{1}{2}(F_A \pm *F_A).$$

The two summands are orthogonal. Define $Q(A) = \int_M \langle F_A \wedge F_A \rangle$, where the invariant pairing is extended to forms. Then

$$\|F_A\|_2^2 = \|F_A^+\|_2^2 + \|F_A^-\|_2^2, \qquad Q(A) = \|F_A^+\|_2^2 - \|F_A^-\|_2^2. \tag{19.7}$$

For a closed manifold, $Q(A)$ is constant as the connection varies on a fixed bundle. Its normalization as an integral characteristic number depends on the representation and the invariant form.

Proposition 19.3 (Energy bound and classical minimizers)**.** *With the conventions above, $S(A) \geq |Q(A)|/(2g^2)$. Equality holds exactly when one of F_A^+ or F_A^- vanishes. Every such connection satisfies the Yang–Mills equation.*

Proof. For nonnegative x, y, $x + y \geq |x - y|$, with equality if and only if $xy = 0$. Apply this to the squared norms in (19.7). A smooth form with zero L^2 norm is zero, so equality is the self-dual or anti-self-dual condition. If $*F_A = \pm F_A$, then $d_A * F_A = \pm d_A F_A = 0$ by Bianchi. The formula for the formal adjoint, $d_A^* = \pm * d_A *$ in the relevant degree, therefore gives $d_A^* F_A = 0$. □

For completeness, the connection independence of Q follows directly from the first variation. Differentiation gives $2\int_M \langle d_A a \wedge F_A \rangle$. Invariance of the pairing and Bianchi identify the integrand with $d\langle a \wedge F_A \rangle$. Stokes’ theorem makes the integral zero on a closed manifold. The statement requires an additional boundary analysis on noncompact spaces; finite action alone should not be silently substituted for all the boundary hypotheses needed in a particular theorem.

This is a useful first encounter with a topological quantity controlling an analytic variational problem. It is nevertheless a statement about classical connections. It gives neither an interacting quantum Hilbert space nor a positive lowest nonzero quantum energy. In particular a positive action of a nontrivial instanton is not the mass of a lowest particle. The action of a Euclidean field configuration and the spectrum of a reconstructed Hamiltonian are different objects.

19.4 The anti-self-dual deformation complex

At an anti-self-dual connection, $F_A^+ = 0$. Linearization of this equation gives $d_A^+ a = 0$, where $d_A^+ = \pi_+ d_A$. Quotienting by infinitesimal gauge transformations leads to

$$0 \longrightarrow \Omega^0(\operatorname{ad} P) \xrightarrow{d_A} \Omega^1(\operatorname{ad} P) \xrightarrow{d_A^+} \Omega^2_+(\operatorname{ad} P) \longrightarrow 0. \tag{19.8}$$

The following proof gives both parts of the assertion that this is an elliptic complex. Having the correct ranks is not by itself sufficient: one needs the composition identity and exactness of the symbol at every nonzero cotangent vector.

Theorem 19.4 (Exactness of the anti-self-dual symbol sequence)**.** *If $F_A^+ = 0$, then $d_A^+ d_A = 0$. For every nonzero covector ξ, the symbol sequence of* (19.8) *is exact.*

Proof. For an adjoint section ϕ, the curvature identity gives $d_A^+ d_A \phi = \pi_+[F_A, \phi] = [F_A^+, \phi] = 0$. For symbol exactness the Lie algebra factor is passive, so it suffices to work with ordinary real forms. Rescale and choose an oriented orthonormal coframe e^1, e^2, e^3, e^4 with $\xi = |\xi| e^1$. The first symbol sends c to $c\xi$ and is injective. If $a = \sum_j a_j e^j$, the second sends it to

$$\frac{|\xi|}{2} \sum_{j=2}^{4} a_j \left(e^1 \wedge e^j + *(e^1 \wedge e^j) \right).$$

The three parenthesized forms are mutually orthogonal and nonzero, and they span the three-dimensional self-dual space. The map is therefore surjective and its kernel consists precisely of multiples of e^1. This kernel is the image of the first symbol. Tensoring the exact finite-dimensional sequence with $\mathfrak{g}$ preserves exactness. □

The proof displays a direct algebraic connection to varieties of complexes: each nonzero cotangent vector determines a pair of linear maps satisfying a zero-product equation and a particular exact rank condition. The geometric content is the existence of a smooth family of these maps, with specified homogeneity in ξ, over the punctured cotangent bundle. A single algebraic complex cannot replace that family. Its global symbol class is the input to index theory.

Let H_A^0, H_A^1, H_A^2 denote the cohomology of (19.8). On a closed manifold elliptic theory makes these spaces finite dimensional and gives harmonic representatives after choosing Sobolev completions. Their meanings are distinct: H_A^0 contains infinitesimal stabilizers, H_A^1 infinitesimal deformations modulo gauge, and H_A^2 obstructions. The operator

$$L_A = d_A^* \oplus d_A^+ : \Omega^1(\operatorname{ad} P) \longrightarrow \Omega^0(\operatorname{ad} P) \oplus \Omega^2_+(\operatorname{ad} P)$$

has index $\dim H_A^1 - \dim H_A^0 - \dim H_A^2$. At an irreducible, unobstructed connection this is the local dimension of the moduli space. Without those conditions the same integer is a virtual dimension, not an assertion that the quotient is a manifold.

Theorem 19.5 (Atiyah–Hitchin–Singer index formula, external)**.** *For an $SU(2)$ bundle on a closed, connected, oriented Riemannian four-manifold, write $k = \langle c_2(P \times_{SU(2)} \mathbb{C}^2), [M]\rangle$ with the convention in which anti-self-dual instantons have positive k. Then*

$$\operatorname{ind} L_A = 8k - 3(1 - b_1(M) + b_2^+(M)). \tag{19.9}$$

This external theorem is an application of the analytic index theorem to the symbol just computed [15]. The proof of the symbol calculation above is not a proof of the index theorem. On S^4 the formula becomes $8k - 3$ for the unframed irreducible problem. Framing at a point removes the three-dimensional constant gauge quotient in the usual regular situation and gives $8k$. Comparing these numbers without recording framing conventions creates a spurious disagreement.

19.5 A finite-dimensional algebraic model: monads and ADHM

There is a stronger bridge than a resemblance between equations. The Atiyah–Drinfeld–Hitchin–Manin construction relates appropriate finite-dimensional algebraic data to framed instantons [14]. We use only its elementary complex identity here, and state the correspondence as an external theorem rather than re-proving its analytic and algebro-geometric foundations.

Let V, W be finite-dimensional complex vector spaces. Choose $B_1, B_2 \in \operatorname{End}(V)$, $I \in \operatorname{Hom}(W, V)$, and $J \in \operatorname{Hom}(V, W)$. Set

$$\mu_{\mathbb{C}} = [B_1, B_2] + IJ.$$

For $z = (z_1, z_2) \in \mathbb{C}^2$, define

$$\alpha_z(v) = ((B_1 - z_1)v, (B_2 - z_2)v, Jv), \quad \beta_z(v_1, v_2, w) = -(B_2 - z_2)v_1 + (B_1 - z_1)v_2 + Iw.$$

Their types are $V \to V \oplus V \oplus W \to V$. Direct multiplication gives

$$\beta_z \alpha_z = [B_1, B_2] + IJ = \mu_{\mathbb{C}}. \tag{19.10}$$

Thus the complex moment-map equation is exactly a zero-composition condition, valid simultaneously for all z.

If α_z is injective and β_z surjective for every relevant z, the middle cohomology has dimension $\dim W$. Indeed $\dim \ker \beta_z = 2 \dim V + \dim W - \dim V$, and quotienting by the $\dim V$-dimensional image of α_z leaves $\dim W$. These constant-rank conditions give a vector bundle locally. They are essential: the equation $\mu_{\mathbb{C}} = 0$ alone permits rank drops and singular objects.

For unitary instantons the real moment-map equation also enters:

$$[B_1, B_1^*] + [B_2, B_2^*] + II^* - J^*J = 0.$$

Together with the appropriate stability and regularity conditions, the data modulo the unitary action on V yield the ADHM correspondence. The complex equation explains why varieties of complexes are relevant; the real equation, positivity, and rank conditions explain why arbitrary points of a variety of complexes are insufficient. In particular this construction describes classical instantons. Its finite-dimensional algebra does not prove a quantum mass gap or an algorithm for elliptic-curve discrete logarithms.

The monad identity is also a good candidate for machine checking: matrix multiplication, commutators, scalar identities, and dimensions have clear finite specifications. Formalizing the full ADHM correspondence would require extensive differential geometry and analysis beyond that elementary identity. Reporting the identity as “the ADHM theorem verified” would discard almost all of the theorem’s content.

19.6 Why an index cannot establish a quantitative gap

An index records a difference of dimensions. A spectral gap is an inequality with a positive numerical constant. The distinction can be proved by an example that is already a complex of finite-dimensional vector spaces.

Proposition 19.6 (Exactness does not control the smallest singular value)**.** *For every $\varepsilon > 0$, the complex $0 \to \mathbb{R} \xrightarrow{\varepsilon} \mathbb{R} \to 0$ is exact, but the smallest eigenvalue of either associated nonzero Laplacian is ε^2. There is no positive lower bound valid for all $\varepsilon > 0$.*

Proof. Multiplication by a nonzero real number is a bijection, so both cohomology groups vanish. Relative to the standard inner products, the adjoint is also multiplication by ε, and $d^*d = dd^* = \varepsilon^2$. For any proposed lower bound $c > 0$, choosing $0 < \varepsilon < \sqrt{c}$ contradicts the bound. □

The same phenomenon appears geometrically. On a flat torus of side length L, the scalar Laplacian has eigenvalues $(2\pi/L)^2|n|^2$ for $n \in \mathbb{Z}^d$. The first positive one is $4\pi^2/L^2$, which tends to zero as $L \to \infty$. Each fixed torus has a compact elliptic problem and a discrete spectrum, yet no volume-independent positive gap follows. The quantum Yang–Mills problem contains a much more difficult version of precisely this issue: a bound must survive both the ultraviolet limit and unbounded physical volume.

This gives a disciplined place for orders and homological algebra. They can organize integral models, modules, deformation complexes, characteristic classes, and exact relations among observables. To conclude a mass gap one still needs coercive estimates or spectral information with constants that remain positive in the required limits. The absence of such an estimate is a mathematical missing implication, not a matter of choosing more persuasive terminology.

Chapter 20

From finite gauge lattices to a quantum spectral gap

20.1 Three different uses of the word lattice

A lattice over an order is a module with a finiteness and torsion-freeness condition over the ground ring. A Euclidean lattice is a discrete subgroup of a real inner-product space, with a metric that makes shortest-vector and decoding questions meaningful. A lattice in a lattice gauge theory is a discrete approximation to spacetime, equipped with vertices, links, and plaquettes. The first two notions can be related by embeddings of number fields and choices of positive forms. The third need not carry any order module structure at all. Identifying all three because they share a word would erase their defining mathematical data.

There are productive interactions nevertheless. A gauge theory on a finite cell complex has incidence maps and chain groups. A choice of coefficient ring can make those chain groups modules over an order. Discrete Hodge operators add positive inner products. Fourier analysis on finite abelian groups, central in quantum algorithms, also diagonalizes many free lattice operators. Each such connection is a construction that can be specified and proved. None entails that the computational hardness of a Euclidean lattice problem controls the mass spectrum of an interacting gauge theory.

20.2 A finite Wilson model exists by elementary integration

Let a finite oriented cell complex have a set E of positively oriented edges and a set $\mathcal{P}$ of oriented plaquettes. Assign $U_e \in G$ to each $e \in E$, and use $U_{\bar{e}} = U_e^{-1}$ for the reverse orientation. The ordered product around a plaquette is U_p. For a fixed unitary representation $G \subset U(m)$ and $\beta \geq 0$, define

$$E_W(U) = \sum_{p \in \mathcal{P}} \left(1 - \frac{1}{m} \operatorname{Re} \operatorname{Tr} U_p\right), \quad Z_\beta = \int_{G^E} e^{-\beta E_W(U)} \prod_{e \in E} dh(U_e), \tag{20.1}$$

where h is normalized Haar measure. Wilson's lattice construction preserves gauge invariance before taking any limit [177]. The following modest result explains exactly how much existence is obtained from compactness alone.

Proposition 20.1 (Finite-volume normalization and gauge invariance)**.** *The quantity Z_β satisfies $e^{-2\beta|\mathcal{P}|} \leq Z_\beta \leq 1$, and $d\mu_\beta = Z_\beta^{-1} e^{-\beta E_W} \prod_e dh(U_e)$ is a probability measure. It is invariant under the vertex gauge action $U_{xy} \mapsto t_x U_{xy} t_y^{-1}$.*

Proof. Every eigenvalue of a unitary matrix has modulus one, so $-m \le \operatorname{Re}\operatorname{Tr} U_p \le m$. Thus $0 \le E_W \le 2|\mathcal{P}|$ and $e^{-2\beta|\mathcal{P}|} \le e^{-\beta E_W} \le 1$. Integration against a probability measure gives the bounds. In particular the normalization is finite and strictly positive. Under a vertex transformation the successive $t_y^{-1} t_y$ factors in a plaquette cancel, leaving conjugation by the transformation at its initial vertex. Trace is invariant under conjugation, so the action is unchanged. Haar measure on a compact group is invariant under both left and right translations. Applying this separately to each edge preserves the product measure and proves the last assertion. □

The proof works for every finite complex and every compact matrix group. It contains no control as the number of plaquettes tends to infinity or as the lattice spacing tends to zero. The lower bound on Z_β itself degenerates exponentially with $|\mathcal{P}|$. Dividing by Z_β produces a probability measure, but does not remove the need for uniform bounds on the observables whose limits one wants to take.

An observable such as a product of finitely many matrix entries is bounded at fixed lattice spacing. A renormalized curvature observable typically contains negative powers of the spacing and cutoff-dependent mixing with other observables. Its bound can diverge as the spacing tends to zero. Consequently, compactness of G is compatible with difficult and highly singular continuum limits. The relevant estimates must concern the renormalized observables and their correlations, not merely the unscaled link variables.

20.3 Reflection positivity is additional structure

Let ϑ reflect Euclidean time and let $\mathcal{A}_+$ denote a space of observables supported at positive times. Write $\Theta F = \overline{F \circ \vartheta}$, with any required adjoint on field labels. Reflection positivity asserts

$$\mathbb{E}[(\Theta F) F] \ge 0 \qquad (F \in \mathcal{A}_+). \tag{20.2}$$

It is not a consequence of the ordinary positivity of a probability measure. For example, take a symmetric random variable Z with $0 < \mathbb{E}[Z^2] < \infty$ and set $X_+ = Z$, $X_- = -Z$. The joint law is a positive probability measure and is invariant under exchanging the two time labels. Nevertheless, for $F = X_+$, $\mathbb{E}[(\Theta F)F] = \mathbb{E}[X_- X_+] = -\mathbb{E}[Z^2] < 0$. Both positivity and reflection symmetry can hold while reflection positivity fails.

The Wilson action has the special structure needed for appropriate lattice reflection-positivity and transfer-matrix theorems [92, 118]. A useful finite-dimensional ingredient can be proved directly.

Proposition 20.2 (Positive-type Wilson link weight)**.** *Let $\rho : G \to U(m)$ be a unitary representation, $\chi(g) = \operatorname{Tr}\rho(g)$, and $c \ge 0$. The function $w(g) = \exp(c \operatorname{Re}\chi(g))$ is of positive type: for all $g_1, \dots, g_n \in G$ and $z_1, \dots, z_n \in \mathbb{C}$,*

$$\sum_{i,j} \overline{z_i} z_j w(g_i^{-1} g_j) \ge 0.$$

Proof. The character χ is of positive type because

$$\sum_{i,j} \overline{z_i} z_j \chi(g_i^{-1} g_j) = \operatorname{Tr}\left(\Big(\sum_i z_i \rho(g_i)\Big)^* \Big(\sum_j z_j \rho(g_j)\Big) \right) \ge 0.$$

The conjugate character is the character of the conjugate unitary representation. Products of characters are characters of tensor products. Expanding

$$w(g) = \sum_{a,b \ge 0} \frac{(c/2)^{a+b}}{a!\,b!} \chi(g)^a \overline{\chi(g)}^b$$

therefore expresses w as a limit of nonnegative linear combinations of positive-type functions. Since $|\chi(g)| \le m$, the double series converges uniformly. For any fixed finite set of points and coefficients, the associated nonnegative quadratic forms converge to the displayed quadratic form, which remains nonnegative. □

This lemma proves a kernel positivity statement. The full lattice theorem also needs a factorization compatible with the reflection plane, the orientation of links, and the integration of gauge variables. Those additional steps are not omitted by calling the lemma a proof of the full Osterwalder–Schrader reconstruction theorem. The separation of claims is particularly useful in a formal development: a locally verified kernel identity can be reused without overstating its global consequence.

20.4 The Hilbert-space quotient behind reconstruction

Reflection positivity supplies the semidefinite sesquilinear form $(F, G)_{\mathrm{OS}} = \mathbb{E}[(\Theta F)G]$. Assume that the correlation functional is defined on a complex vector space $\mathcal{V}_+$ of finite positive-time field words, that this form is Hermitian, and that it is nonnegative on the diagonal. The following quotient step is purely algebraic.

Proposition 20.3 (Null vectors and positive quotient)**.** *For a positive semidefinite Hermitian form B on a complex vector space, the set $N = \{v : B(v, v) = 0\}$ is a linear subspace and $B(v, w) = 0$ for every $v \in N$ and w. The induced form on $\mathcal{V}_+/N$ is positive definite.*

Proof. If $B(w, w) > 0$, positivity of $B(v + \lambda w, v + \lambda w)$ for every $\lambda \in \mathbb{C}$ and minimization of this real quadratic polynomial gives $|B(v, w)|^2 \le B(v, v)B(w, w)$. If $B(w, w) = 0$, apply the same argument to $w + \varepsilon u$ with a positive-norm u and pass to the limit, or examine the linear term of the polynomial directly. Thus Cauchy–Schwarz holds in all cases. If $v \in N$ it forces $B(v, w) = 0$. This orthogonality implies that sums and scalar multiples of null vectors are null, and that changing a representative by a null vector does not change any inner product. Finally an equivalence class with zero norm has a representative in N and is therefore the zero class. □

Completing this quotient gives a Hilbert space. Euclidean time translation, under the additional continuity and boundedness estimates required in reconstruction, induces a strongly continuous self-adjoint contraction semigroup. Its generator is a nonnegative self-adjoint operator H and the semigroup is e^{-tH}. A complete reconstruction must also supply field domains, distributional regularity, locality, analytic continuation, and the Poincare action. The original Osterwalder–Schrader papers state precise regularity and growth hypotheses for these conclusions [116, 117]. The quotient lemma alone supplies none of them.

One should track test-function spaces carefully. A correlation kernel defined only where $x_i \ne x_j$ is not automatically a tempered distribution on all configurations. Terms supported on the coincidence diagonals are invisible on separated tests. A simple illustration is T and $T + c\delta_0$ on $\mathbb{R}^d$: they agree on every test function supported away from the origin, but disagree on a test with $f(0) \ne 0$. Even all derivatives of a test vanishing on a diagonal make every derivative of the corresponding delta distribution invisible. This does not invalidate a separated-point construction. It specifies the extra work necessary to extend it, establish positivity of an extension, and prove uniqueness under a chosen normalization.

20.5 What the mass-gap statement actually requires

In the Hamiltonian formulation, fix a complex Hilbert space $\mathcal{H}$, a normalized vacuum Ω, and a nonnegative self-adjoint Hamiltonian H with $\ker H = \mathbb{C}\Omega$. Let P_0 be orthogonal projection onto the vacuum. A positive gap means that there exists $\Delta > 0$ with

$$\operatorname{spec}(H) \subset \{0\} \cup [\Delta, \infty). \tag{20.3}$$

For the Yang–Mills Millennium problem this spectral assertion accompanies a nontrivial four-dimensional quantum field theory with the required Yang–Mills local-field and short-distance structure. It is not an independent name for any positive matrix eigenvalue [68]. Vacuum uniqueness, positive energy, existence of a nonzero local observable, and a mass gap are four different conditions.

Theorem 20.4 (Spectral gap and Euclidean-time decay)**.** *Under the hypotheses just stated, the following are equivalent for a fixed $\Delta > 0$:*

1. (20.3) *holds.*
2. $\|e^{-tH}(1 - P_0)\| \le e^{-t\Delta}$ *for every* $t \ge 0$.
3. *For every* $v \perp \Omega$ *and every* $t \ge 0$, $0 \le \langle v, e^{-tH} v\rangle \le e^{-t\Delta}\|v\|^2$.

Proof. Let E_H be the spectral measure of H. Under the first condition, the restriction to $\Omega^\perp$ has spectral support in $[\Delta, \infty)$. Functional calculus bounds $e^{-t\lambda}$ by $e^{-t\Delta}$ on that support, proving the second condition. Positivity of e^{-tH} and Cauchy–Schwarz then give the third.

Conversely suppose the third condition and suppose the spectral projection $E_H((0, \Delta))$ is nonzero. Since $(0, \Delta)$ is a countable union of closed intervals $[a, b]$ with $0 < a \le b < \Delta$, one such projection is nonzero. Choose a nonzero v in its range. It is orthogonal to the zero-energy eigenspace, and

$$\langle v, e^{-tH} v\rangle = \int_{[a,b]} e^{-t\lambda}\, d\langle v, E_H(\lambda) v\rangle \ge e^{-tb}\|v\|^2.$$

For every $t > 0$, $e^{-tb} > e^{-t\Delta}$, contradicting the third condition. Thus the projection is zero. The spectral theorem then gives the first condition. The case $t = 0$ is the equality bound and causes no exception. □

For a bounded observable A, set $a = (A - \langle \Omega, A\Omega\rangle)\Omega$. Then $a \perp \Omega$, and the associated Euclidean-time connected two-point function is

$$C_A(t) = \langle a, e^{-tH} a\rangle.$$

The theorem bounds it by $e^{-t\Delta}\|a\|^2$. For unbounded smeared fields, the same statement applies once the vacuum is in the relevant operator domain and a is a Hilbert-space vector. Writing products of unsmeared fields at the same point would evade this domain issue, not solve it. Spatial clustering in a relativistic theory requires the locality and spectrum properties that relate time and space; the elementary theorem here concerns Euclidean time only.

A useful converse needs control on a total collection of vectors. If the third inequality holds with one common Δ on a linear dense subspace of $\Omega^\perp$, continuity of the bounded operator e^{-tH} extends it to every vector, and the theorem proves the gap. Control of a single observable cannot do this. For example take $\mathcal{H} = \mathbb{C}\Omega \oplus \mathbb{C}a \oplus L^2(0, 1)$, let $H\Omega = 0$, $Ha = a$, and let H act by multiplication by x on $L^2(0, 1)$. The a-correlation is exactly $e^{-t}\|a\|^2$, but the third summand has spectrum down to zero. The complete Hamiltonian has no positive mass gap.

20.6 A rigorous countermodel to ultraviolet-to-infrared inference

Short-distance asymptotics probe high spectral energies, whereas a gap concerns low energies. This distinction can be shown with two explicit positive spectral measures. Define for $t > 0$

$$C_{\mathrm{g}}(t) = \int_1^\infty \frac{E^7}{\log^2(e+E)} e^{-tE}\, dE, \tag{20.4}$$

$$C_{\mathrm{ng}}(t) = C_{\mathrm{g}}(t) + \int_0^1 e^{-tE}\, dE. \tag{20.5}$$

The first measure has support in $[1, \infty)$; the second has support arbitrarily close to zero. Both Laplace transforms are finite for every $t > 0$ and have the same leading short-distance behavior.

Proposition 20.5 (Identical leading ultraviolet behavior, different gaps)**.** *For either function in* (20.5)*,*

$$\lim_{t \downarrow 0} t^8 \log^2(1/t) C(t) = 7!.$$

The corresponding positive spectral measures have, respectively, a positive lower support bound and no positive lower support bound.

Proof. Write $L = \log(1/t)$ and substitute $u = tE$ in the first integral. After multiplication by $t^8 L^2$, the result is

$$\int_t^\infty u^7 e^{-u} \left(\frac{L}{\log(e + u/t)} \right)^2 du.$$

For $u \geq \sqrt{t}$ and $t < 1$, $\log(e + u/t) \geq L/2$, so the integrand is at most $4u^7 e^{-u}$. For fixed $u > 0$ the parenthesized ratio tends to one. Dominated convergence on the region $u \geq \sqrt{t}$, extended by zero elsewhere, therefore gives $\int_0^\infty u^7 e^{-u} du = 7!$. The contribution from $t \leq u \leq \sqrt{t}$ is at most $L^2 \int_0^{\sqrt{t}} u^7 du = L^2 t^4/8$, which tends to zero. This proves the first limit. The additional integral is bounded by one, and $t^8 \log^2(1/t) \to 0$, so it does not change the limit. The support assertions follow immediately from the strictly positive densities on their respective intervals. □

These examples are positive spectral models, not constructions of local four-dimensional gauge theories. Their role is logical: a proposed inference using only positivity and this ultraviolet asymptotic cannot force a gap, because the two examples satisfy those inputs and have different infrared support. Additional Yang–Mills-specific estimates may well force a gap. They must be stated and proved rather than inferred from the leading asymptotic alone.

20.7 Uniform estimates that would survive a limit

Here is a precise version of one sufficient missing step. Work on one Hilbert space so that convergence has an unambiguous meaning. Let H_n and H be nonnegative self-adjoint operators, and let P_n, P be their zero-energy spectral projections. Assume $P_n \to P$ strongly and $e^{-tH_n} \to e^{-tH}$ strongly for each $t \geq 0$.

Proposition 20.6 (Passage of a uniform gap bound)**.** *If for one $\Delta > 0$ and all $n, t \geq 0$,*

$$\|e^{-tH_n}(1 - P_n)\| \leq e^{-t\Delta},$$

then $\|e^{-tH}(1 - P)\| \leq e^{-t\Delta}$ for all $t \geq 0$.

Proof. For fixed v and t, write

$$e^{-tH_n}(1-P_n)v - e^{-tH}(1-P)v = e^{-tH_n}(P-P_n)v + (e^{-tH_n}-e^{-tH})(1-P)v.$$

The first term tends to zero since the semigroups are contractions and $P_n \to P$ strongly. The second tends to zero by the assumed strong convergence. Taking norms in the uniform bound and passing to the limit gives $\|e^{-tH}(1-P)v\| \le e^{-t\Delta}\|v\|$. Taking the supremum over unit vectors proves the claim. □

For actual lattice-to-continuum constructions the Hilbert spaces may vary. One must first supply compatible embeddings or formulate convergence in terms of a common observable algebra and its states. The proposition is a conditional theorem with clear hypotheses; it does not assert that Yang–Mills satisfies those hypotheses. It is useful because it localizes the hard work in the uniform bound and the comparison maps, rather than hiding it inside the word "limit."

Physical units are equally important. If a transfer matrix T_a is normalized to have largest eigenvalue one and corresponds to time step a, its Hamiltonian is $H_a = -a^{-1}\log T_a$ on the relevant spectral subspace. A bound for a second eigenvalue must therefore be interpreted through $-a^{-1}\log\lambda_2(a)$. Merely proving $\lambda_2(a) < 1$ separately for each a does not imply a common positive physical gap. For example $\lambda_2(a) = e^{-a^2}$ is strictly less than one but gives physical gap $a \to 0$. This elementary counterexample rules out a frequent confusion between finite-cutoff positivity and continuum spectral control.

20.8 Renormalization and what a running coupling can prove

The passage from the finite Wilson law to continuum fields requires more than refinement of a mesh. The bare coupling and the observable normalizations must change with the cutoff. The four-dimensional renormalization-group program of Bałaban constructs effective actions and controls transformations between scales; the 1987 paper cited here explicitly works in its small-field approximation [19]. The limitation matters: a small-field expansion is not automatically control of the full probability measure, and estimates for all fields require additional arguments.

The original perturbative discovery of asymptotic freedom determines a negative leading beta function for the appropriate nonabelian theories [57]. An exact scalar differential equation captures the leading behavior without claiming to be the full quantum theory. Let $b > 0$ and suppose, on an interval of scales $\mu > 0$, that

$$\mu\frac{dg}{d\mu} = -bg^3, \qquad g(\mu) > 0. \tag{20.6}$$

Then

$$\frac{1}{g(\mu)^2} = \frac{1}{g(\mu_0)^2} + 2b\log\frac{\mu}{\mu_0}, \qquad g(\mu)^2 = \frac{1}{2b\log(\mu/\Lambda)}, \tag{20.7}$$

where $\Lambda = \mu_0\exp(-1/(2bg(\mu_0)^2))$ and the displayed positive solution is considered only for $\mu > \Lambda$. The proof is simply $\frac{d}{d\log\mu}(g^{-2}) = -2g^{-3}\frac{dg}{d\log\mu} = 2b$, followed by integration. In particular $g(\mu) \to 0$ at large μ.

The parameter Λ illustrates dimensional transmutation: an initial dimensionless coupling together with a reference scale determines a dimensional scale. Equation (20.7) does not prove a particle of mass Λ, a mass gap proportional to Λ, or a controlled continuation through $\mu = \Lambda$. Its denominator vanishes there because the leading differential model leaves its useful weak-coupling range. Deducing

the infrared spectrum from that singularity would replace a missing theorem with an artifact of an approximation.

The same distinction appears in the control of irrelevant interactions. For a local operator of engineering degree d in four dimensions, the classical rescaling factor of its coupling under a length enlargement by $b_{\mathrm{s}} > 1$ is b_{s}^{4-d}. The exponent explains the words relevant, marginal, and irrelevant. It does not bound the nonlinear errors or the infinite collection of generated operators. A simple inequality isolates one of the required analytic mechanisms.

Proposition 20.7 (A contractive recurrence with controlled errors)**.** *Let $0 \le \rho < 1$, $a_n \ge 0$, $\varepsilon_n \ge 0$, and $a_{n+1} \le \rho a_n + \varepsilon_n$. Then*

$$a_n \le \rho^n a_0 + \sum_{j=0}^{n-1} \rho^{n-1-j} \varepsilon_j.$$

If $\varepsilon_j \le \varepsilon$ for all j, this is at most $\rho^n a_0 + \varepsilon/(1-\rho)$. If instead $\varepsilon_j \to 0$, then $a_n \to 0$.

Proof. The finite-sum bound follows by induction. Multiplying the bound for a_n by ρ and adding ε_n gives the bound for a_{n+1}. The uniform-error assertion follows from summing a geometric series. For the final assertion fix $\eta > 0$ and choose J with $\varepsilon_j < (1-\rho)\eta/2$ for $j \ge J$. The sum over these indices is at most $\eta/2$. The finitely many earlier terms, together with $\rho^n a_0$, tend to zero. For sufficiently large n they contribute less than $\eta/2$. Hence $a_n < \eta$ eventually. □

A renormalization proof must identify the norms represented by a_n, show that the nonlinear transformation satisfies the recurrence, and bound the errors with the right uniformity in volume and insertion degree. The elementary lemma makes none of those analytic inputs disappear. What it supplies is a reusable logical step once those inputs exist. This is also a sensible formalization boundary: verify the recurrence lemma independently, then connect it to the actual effective-action estimates through a theorem whose hypotheses remain visible.

20.9 Pure phases, clustering, and the order of quantifiers

Suppose a translation-invariant state is a mixture $\omega = \int_X \omega_\xi \, d\nu(\xi)$ of component states. Even if every component clusters, the mixture need not. For a self-adjoint observable A, assume enough boundedness to pass the large-separation limit through the integral. Component clustering gives

$$\lim_{|x|\to\infty} \omega(A\,\tau_x(A)) = \int_X \omega_\xi(A)^2 \, d\nu(\xi).$$

Subtracting $\omega(A)^2$ leaves the variance of $\xi \mapsto \omega_\xi(A)$, which can be strictly positive. The argument is an elementary proof that a phase decomposition and clustering of each phase do not imply clustering of the original mixture.

Likewise, the statement that every diverging volume sequence has a convergent subsequence does not assert that all subsequential limits agree. The scalar sequence $(-1)^n$ already distinguishes sequential compactness from convergence. In a field theory the extracted object must include all correlation functions coherently. Choosing a different subsequence for each observable is weaker than finding one subsequence working for all of them. Conversely, finding one common subsequence remains weaker than proving independence of the subsequence. These distinctions are precisely the ones needed to read the recent formal construction correctly.

There is also a separate distinction between uniqueness of a vacuum inside one representation and uniqueness of the theory selected by a continuum limit. Two different pure phases can each

have a unique vacuum in their own Hilbert spaces. Neither uniqueness statement identifies the phases with one another. A formal verification plan should represent these questions with different predicates so that a theorem about one cannot be mistakenly reported as a theorem about the other.

Chapter 21

The 2026 OpenAI Yang–Mills construction: statement and verification

21.1 A dated research result, not the older amplitude result

This chapter records the public evidence available on 13 November 2026. OpenAI's announcement dated 26 October 2026 includes a claimed construction of four-dimensional pure Yang–Mills theory, as part of broader announced work on gauge theories with fermions. The announcement identifies a natural-language Yang–Mills existence proof dated 9 September and a completed formalization dated 16 October. It describes an independent group organized around David Gross, with Michael Douglas, Arthur Jaffe, Igor Klebanov, and Nikita Nekrasov, to examine and explain the work [112]. These are dates and claims of the announcing organization; the subsequent released submission and the independent assessment provide the mathematical objects examined here.

The independent group's assessment, dated 10 November 2026, addresses the submission of 28 October. Its stated judgment is that the submitted construction yields a Wightman quantum field theory from a four-dimensional Wilson lattice scaling limit for the specified simple gauge groups. It explicitly withholds a mass-gap claim and uniqueness of the infinite-volume limit. The assessors also say they do not yet fully understand the proof and distinguish their review from a Clay Mathematics Institute decision [45]. It is therefore inaccurate both to dismiss the announcement as merely a perturbative amplitude calculation and to present it as a completed solution of the whole Millennium problem.

The earlier paper on single-minus gluon amplitudes is a different result. Its arXiv submission was posted on 12 February 2026, with a second version on 9 March. It concerns special half-collinear kinematics in Klein space or complexified momenta, rather than a four-dimensional nonperturbative vacuum construction [58]. Its AI-assisted derivation may be relevant to research methodology, but it does not supply the infrared estimate required in (20.3). Conflating the two projects would misstate both the scope and the chronology of the recent work.

21.2 Immutable source identities and review coverage

The original public source was inspected at commit `3456e5e75cf8d4cd558955a5c7638fba2d68f464` of `yang-mills-group/OpenAI-YM-20261028`. The group's review material was inspected at commit `55ba9fc36293b3f4d52c6b1f813169a0be76ee98` of `yang-mills-group/YM`. The accompanying

evidence manifest records URLs and hashes of the files used in this review. The 35-page assessment was read in full, together with the headline theorem, its public specifications, the stronger construction theorem, the combined acceptance specification and proof, and selected explanatory verification modules. This is a review of those materials, not a claim to have read every proof in the approximately million-line upstream development.

The distinction matters operationally. A source snapshot, a successful local compilation, a kernel replay, and an independent semantic review are different deliverables. The present project can pin and analyze an external formal result without pretending that a historical upstream build was performed again here. Reproduction instructions must identify the toolchain, dependencies, target declarations, and exact evidence actually produced by a new run. A link to an upstream pass is valuable evidence of that upstream check; it is not a substitute for a fresh local pass when a fresh pass is being claimed.

21.3 The headline theorem and its quantifiers

At the frozen revision, the file `Solution.lean` proves the declaration `YangMills.jointExistence`. Its target is defined in `ChallengeSpec.lean`. In ordinary mathematical notation, its quantifier pattern is

$$\forall G\ \big(\mathrm{Simple}(G) \Longrightarrow \exists C\,[\mathrm{Reg}(C)\ \wedge\ \forall p \in \mathrm{Inc}(\mathbb{N})\ \exists q \in \mathrm{Inc}(\mathbb{N})\ \exists T\ \ \mathrm{Properties}(C, T, p \circ q)]\big). \quad (21.1)$$

Here the gauge-group type is a closed matrix subgroup of a unitary group, and simplicity concerns its nonabelian real Lie algebra. The object C is one Wilson family, p requests an increasing sequence of physical volume indices, q refines it, and T is a continuum hierarchy associated to the same family [111, 113].

Several consequences follow from this syntax before any deep proof is examined. The family is chosen before the volume sequence, so this is stronger than choosing a new finite model for each requested sequence. The hierarchy is selected after refinement, so uniqueness across refinements is not a conclusion. All four fields of Properties concern the same T, so the reconstruction, operator-product estimates, and trace test cannot be supplied by unrelated models. These are elementary consequences of the order of quantifiers. They are exactly the kind of details lost in a slogan such as "a limit exists."

The public regularity predicate requires a positive logarithmic slope of the bare action coefficient, uniformly in volume, continuous and gauge-invariant lattice insertions, and one finite lattice stencil for each token that works at every cutoff and volume. The public Properties predicate then bundles a Wilson-limit statement, reconstructed fields, a separated-point operator-product expansion, and a specified trace-based exclusion of free-field profiles [111]. The strength of the result rests on the definitions and proofs of these structures, not on the short final tactic script that assembles them.

21.4 A hierarchy with finite-degree local fields

The formal hierarchy includes a Hilbert space, field labels, a Wightman theory, finite-dimensional complex spaces F_d of fields of bounded engineering degree, maps from F_d to the full label space, and a surjective representation of those fields by lattice tokens. It also includes a compact space X of phases with a probability measure, a selected phase, correlation functions for the mixture and phases, analytic continuations, OPE channels and coefficients, and a centered degree-four trace field [111].

This is a structural resemblance to filtered algebra that can be used carefully. If F_d is finite dimensional and increasing with d, then one can choose successive complements and index finitely many new channels at each degree. Such choices support finite truncations of an expansion. They do not imply that multiplication of quantum fields at the same point is an everywhere-defined map $F_d \otimes F_e \to F_{d+e}$. Coincidence limits need renormalization, and the OPE is an asymptotic statement about separated fields. Replacing that analytic statement by an exact multiplication table would change the theorem.

Similarly, a quotient by correlation-null labels is a mathematically natural way to obtain faithful fields. If a preliminary label space V maps linearly to correlation functions by R, the quotient $V/\ker R$ injects into the image of R by the first isomorphism theorem. This elementary fact explains why faithfulness of a quotient is plausible; it does not prove that the chosen preliminary labels have the desired Yang–Mills content. That latter assertion requires comparison with curvature polynomials and control of their null relations. The independent review consequently includes a separate operator-content check instead of treating quotient faithfulness as a complete answer [182].

21.5 The two limits and the physical normalization

Let $a_N \downarrow 0$ denote the chosen spacings and $L_p \to \infty$ the physical side lengths. A schematic smeared lattice correlation is

$$C_{N,p}(v; f) = a_N^{4n} \sum_{x_1,\dots,x_n} f(a_N x_1, \dots, a_N x_n) \int \prod_{j=1}^{n} I_{N,p}(v_j, x_j; U)\, d\mu_{N,p}(U). \tag{21.2}$$

The original Wilson specification fixes the finite law, positions, insertions, and separated-test convergence; the proof retains a full ultraviolet limit for the chosen spacing sequence at fixed volume followed by the asserted refined volume limit [115]. The ordered limit is $\lim_{k\to\infty} \lim_{N\to\infty} C_{N,p(q(k))}$. No interchange of these limits is implicit.

The distinction can be understood without field theory. The numbers $f_{N,p} = N/(N+p)$ satisfy $\lim_p \lim_N f_{N,p} = 1$ and $\lim_N \lim_p f_{N,p} = 0$. Therefore an ordered limit theorem cannot be turned into a joint limit theorem without uniformity of an appropriate kind. In the gauge problem, locality, renormalization, and positivity make the situation more structured, but the logical requirement remains.

The stronger acceptance specification also fixes a reference spacing and a common reference coupling after renormalization-group evolution, and requires the insertions to read out the same reference labels across cutoffs and volumes. Its comparison with another finite lattice formalization includes the conversion of coupling conventions by the matrix representation dimension [182]. These clauses prevent a change in units or a rescaling of observables from masquerading as a physically consistent limit.

In particular, the declaration named `TwoPointScale.MassLimit` in the combined specification concerns a positive limiting two-point calibration. Its name must not be read as the spectral condition (20.3). A nonzero correlation provides a scale-sensitive observable; the spectrum could still have arbitrarily small positive energies. The countermodels in the preceding chapter demonstrate the logical difference explicitly.

21.6 Reconstruction, phases, and test spaces

The reconstructed theory includes a strongly continuous unitary translation representation with spectrum in the forward cone, a dense common domain for smeared fields, a normalized cyclic

vacuum, adjoint and regularity properties, Poincare covariance, and spacelike commutativity. The selected phase has vacuum uniqueness. The formal public definitions specify these properties in the field system, rather than leaving the meaning of "Wightman theory" as an informal comment [114]. In this project these are statements of an external formal development; our elementary semigroup arguments use their consequences conditionally.

There are three separate classes of tests to keep apart. First, lattice-limit statements use smooth compactly supported tests away from coincidences. Second, the review's independent scalar OS comparison uses Schwartz tests flat to every order on coincidence diagonals. Third, the reconstructed Wightman distributions act on unrestricted Schwartz tests. The review compares the selected centered trace hierarchy to a separately specified OS interface and growth condition [188]. These domains have different powers to detect contact terms, as the delta-distribution example above shows.

The mixed infinite-volume correlation table is related to phase correlations by integration over X. Each phase's reconstructed vacuum uniqueness is compatible with nonuniqueness of the mixture's vacuum and with disagreement between different infinite-volume subsequences. The combined acceptance theorem keeps the selected scalar field, phase, trace, and translation tensor attached to one construction witness [184]. This coherence is mathematically significant. It does not remove the remaining uniqueness questions by definition.

21.7 The trace law, non-freeness, and OPE estimates

For the normalized centered stress trace Θ in the selected phase, the public theorem includes a positive leading short-distance constant:

$$\lim_{r\downarrow 0} r^8 \log^2(1/r) S_{\Theta,\Theta}(-re_0/2, re_0/2) = \frac{3\dim_{\mathbb{R}} \mathfrak{g}}{\pi^4}. \tag{21.3}$$

The independent trace adapter expresses the same limit in all spatial directions for the same centered field and normalization [186]. The detailed coefficient is convention dependent; copying it while changing the normalization of the stress tensor would be incorrect.

The public non-freeness test excludes finite combinations of specified mass-shell Wick profiles, including finite products of kernels with nonnegative masses and homogeneous numerators. This is a precise exclusion class, and the formal target describes it explicitly [111]. It should not be expanded into an assertion excluding every possible Gaussian representation, every generalized free field with a continuous mass spectrum, or every conceivable equivalence of theories. These stronger statements would need new definitions and proofs. Nor does (21.3) establish the gap: multiplying the examples in (20.5) by the appropriate positive constant gives exactly the same leading coefficient with different infrared supports.

The separated OPE has an input cluster of fields v_i of total degree $w = \sum_i d_i$, a cluster radius r, and spectators kept a fixed positive distance away. At cutoff degree D it has the form

$$S^s(v(rx), u(y)) = \sum_{d_\alpha \le D} c_\alpha(v; rx) S^s(o_\alpha(0), u(y)) + R_{D,s}(r), \tag{21.4}$$

with bounds of the type $|c_\alpha| \le r^{d_\alpha - w - \varepsilon}$ and $|R_{D,s}(r)| \le r^{D+1-w-\varepsilon}$ for all sufficiently small r, uniformly over the prescribed separated configurations and phases. The same coefficients serve all spectators, phases, and truncation degrees. The independent OPE check restates such bounds directly and derives them from the source theorem [187].

Here "all orders" means arbitrarily high *field-degree truncation* in this expansion. It is not a claim that every perturbative loop coefficient has been compared with a particular renormalization scheme. The difference is visible in the variables: D indexes the degrees of local fields, whereas loop order indexes powers of a coupling in a perturbative series. An equality between the two kinds of information would be another theorem. Maintaining this distinction makes the OPE result more useful, because it specifies the exact remainder one can legitimately use in further arguments.

21.8 A stronger acceptance target and the scope of checking

The original short public target did not retain every connection between the lattice construction and the continuum witness. The review's `CombinedAcceptance` specification explicitly keeps one family, one extracted table, and one final construction object with those certificates. Its proof applies the stronger source construction theorem once and assembles the additional consequences for that same object [182, 184]. This is a concrete example of why verifying a theorem's proof term and verifying that its type expresses the intended scientific claim are distinct activities.

The recorded independent check reports a clean build of 9,470 project modules using Lean 4.32.0, with 9,477 supplied source files matched to a manifest. It reports no project axioms, no direct placeholder dependency, and no unsafe declarations in the inventoried project declarations. A reference challenge contains a placeholder, but the solution does not import that reference proof. The reported principal proof dependencies use the standard axioms of propositional extensionality, classical choice, and quotient soundness [183].

The group's separate Comparator records cover the public endpoint, the combined acceptance endpoint, and the OPE endpoint. They include controls that reject a changed target, an added forbidden axiom, and a changed definition. The historical macOS launcher did not provide Linux process isolation, and the check used Comparator's built-in Lean kernel rather than another independently implemented kernel [185]. These qualifications describe the actual experiment. They do not negate a kernel check, and they should not disappear when the check is summarized.

This project distinguishes four levels of confidence. A well-typed declaration verifies a particular proposition relative to its dependencies. An axiom audit identifies the assumptions used. A specification comparison checks that a fixed target was not replaced. Semantic review assesses whether the fixed target captures the mathematical and physical claim. The last step cannot be reduced to searching for the word `sorry`: a theorem with a vacuous hypothesis can have a completely valid proof and still express the wrong scientific question. Conversely, a separate unfinished reference stub is not an assumption of a proof that never imports it. Both observations follow from the dependency structure of formal logic.

21.9 What can be built on the result now

The most useful next contributions are statements whose inputs can be bound to the already constructed objects. Examples include a precise uniqueness criterion for the phase mixture, a criterion that identifies different volume subsequences, or a volume-uniform decay estimate on a total observable domain. Each should specify the family, normalization, test spaces, and phase before stating a conclusion. The spectral proposition in the preceding chapter explains exactly how a successful uniform estimate would yield a gap.

An algebraic improvement could be valuable without resolving the infrared problem. A clearer description of the filtered local-field spaces might simplify the verification of null relations. A finite presentation of low-degree relations could make the operator-content interface easier to audit. A

normalized family of finite complexes could organize Ward identities and cohomological constraints. Such work would need to show that its algebraic relations are respected by the same renormalized correlation hierarchy, rather than merely by a classically related polynomial ring.

The appropriate positive conclusion is therefore substantial and precise: the current public development supplies a formally verified external construction with documented independent checks, and it supplies concrete objects to which further theorems can refer. The remaining gap is not an absence of numerical plausibility. It is the absence, in the reviewed claim, of a theorem giving a strictly positive lower bound on the nonvacuum energy spectrum in the infinite-volume continuum theory. The present manuscript does not supply that missing theorem and does not label its classical or finite-dimensional propositions as substitutes for it.

Chapter 22

An explicit arithmetic–geometric bridge through elliptic tori

22.1 Orders acting on a complex elliptic curve

A concrete example can connect orders, elliptic curves, unitary gauge fields, and elliptic operators without making a conjectural identification. Let $\Lambda = \mathbb{Z}\omega_1 + \mathbb{Z}\omega_2 \subset \mathbb{C}$ with $\mathrm{Im}(\omega_2/\omega_1) > 0$, and form the complex torus $E = \mathbb{C}/\Lambda$. Its algebraic realization as a complex elliptic curve is the usual uniformization theorem. The elementary constructions below begin with this displayed torus, so they do not need uniformization as an unproved step.

For $\alpha \in \mathbb{C}$, multiplication on $\mathbb{C}$ descends to a map $[z] \mapsto [\alpha z]$ on E exactly when $\alpha\Lambda \subseteq \Lambda$. Necessity follows because each $\lambda \in \Lambda$ represents the same point as zero. Sufficiency follows because $z - z' \in \Lambda$ then implies $\alpha z - \alpha z' \in \Lambda$. Thus

$$\mathrm{End}_{\mathrm{mult}}(E) = \{\alpha \in \mathbb{C} : \alpha\Lambda \subseteq \Lambda\}$$

is a ring. If it contains a nonreal α, it lies in an imaginary quadratic number field and is an order in that field. To see the essential algebra, represent multiplication by α on the chosen lattice basis by a matrix $M_\alpha \in M_2(\mathbb{Z})$. Its characteristic polynomial is monic of degree two, annihilates α, and has negative discriminant when α is nonreal. The finite-index assertion for the endomorphism ring is the standard complex-multiplication structure theorem; the matrix calculation identifies its integrality mechanism. One can make the order assertion explicit. Rescale so that $\Lambda = \mathbb{Z}+\mathbb{Z}\tau$. Write $\alpha = a+c\tau$ and $\alpha\tau = b + d\tau$ with integers a, b, c, d. Since α is nonreal, $c \neq 0$, and $c\tau^2 + (a - d)\tau - b = 0$. Thus $K = \mathbb{Q}(\tau)$ is imaginary quadratic. Every other multiplier β has $\beta \cdot 1 \in \Lambda$, so $\beta \in K$. Its integral matrix on Λ makes it an algebraic integer, hence the multiplier ring lies in $\mathcal{O}_K$. It contains $\mathbb{Z}[\alpha]$, a rank-two sublattice of $\mathcal{O}_K$. Therefore it too is a rank-two subring of finite index in $\mathcal{O}_K$, which is exactly an order. Here the usual fact that $\mathcal{O}_K$ is a rank-two lattice is the number-field theorem established or cited in the arithmetic chapters.

This is an actual use of an order acting on a lattice. It concerns endomorphisms of a complex torus and should be compared with the arithmetic endomorphism-order results in the earlier chapters. It does not replace the separate problem of computing endomorphisms or discrete logarithms over a finite field.

22.2 The same lattice controls flat unitary line bundles

A unitary character $\chi : \Lambda \to U(1)$ determines a line bundle

$$L_\chi = (\mathbb{C} \times \mathbb{C})/\sim, \qquad (z, v) \sim (z + \lambda, \chi(\lambda)v).$$

Here χ is the automorphy character: with this quotient convention, parallel-transport holonomy is χ^{-1}. The trivial connection on the product descends because the multipliers are constant. It is flat and unitary. Conversely, parallel transport of a flat unitary line bundle around the two lattice generators gives its holonomy character. In rank one, conjugacy of holonomy characters is trivial. The general relation between stable bundles and unitary representations is a deep theorem of Narasimhan and Seshadri; Atiyah's classification gives the special structure of bundles on elliptic curves [9, 103].

For this explicit rank-one model, all characters are

$$\chi_{s,t}(m\omega_1 + n\omega_2) = \exp(2\pi i(ms + nt)), \qquad (s, t) \in (\mathbb{R}/\mathbb{Z})^2. \tag{22.1}$$

Indeed a homomorphism out of a free abelian group is uniquely determined by the images of its basis, and each unit complex number has an argument modulo 2π. Tensor product of line bundles corresponds to multiplying characters, hence adding (s, t). The characters with $\chi^n = 1$ are exactly the n^2 points with $s, t \in \frac{1}{n}\mathbb{Z}/\mathbb{Z}$. This explains how finite torsion and finite Fourier characters occur inside a moduli space of flat connections.

Proposition 22.1 (Order action and dual action)**.** *If $\alpha\Lambda \subseteq \Lambda$, pullback by $[z] \mapsto [\alpha z]$ sends L_χ to the flat line bundle with character $\chi \circ \alpha$. If M_α represents multiplication on column lattice coordinates, the corresponding character parameters transform by M_α^{T} modulo $\mathbb{Z}^2$.*

Proof. Under pullback, the constant multiplier for a lattice translation λ is the multiplier for $\alpha\lambda$, namely $\chi(\alpha\lambda)$. If $x = (m, n)^{\mathsf{T}}$ denotes lattice coordinates and $u = (s, t)^{\mathsf{T}}$ denotes character coordinates, then (22.1) is $\exp(2\pi i u^{\mathsf{T}} x)$. Precomposition replaces x by $M_\alpha x$, giving $u^{\mathsf{T}} M_\alpha x = (M_\alpha^{\mathsf{T}} u)^{\mathsf{T}} x$. □

The transpose is the same algebraic operation that appears in dual lattices and Fourier transforms. What unifies these examples is the evaluation pairing between a module and its characters. That pairing can be formalized directly. It should not be enlarged into a claim that arithmetic elliptic curves, all gauge moduli spaces, and quantum hidden subgroup problems are equivalent categories.

22.3 An elliptic operator with a completely explicit spectrum

Take a rectangular torus with side lengths $L_1, L_2 > 0$ and automorphy parameters (s, t). Sections of its flat line bundle can be represented by smooth functions satisfying

$$f(x + L_1, y) = e^{2\pi i s} f(x, y), \qquad f(x, y + L_2) = e^{2\pi i t} f(x, y).$$

The nonnegative covariant Laplacian is represented by $-\partial_x^2 - \partial_y^2$ with these boundary conditions. Its Fourier modes are

$$f_{m,n}(x, y) = \exp\left(2\pi i \left(\frac{m + s}{L_1} x + \frac{n + t}{L_2} y\right)\right), \qquad (m, n) \in \mathbb{Z}^2,$$

with eigenvalues

$$\lambda_{m,n}(s, t) = 4\pi^2 \left(\frac{(m + s)^2}{L_1^2} + \frac{(n + t)^2}{L_2^2}\right). \tag{22.2}$$

The boundary conditions and eigenvalue equation follow by direct substitution. Completeness follows by multiplying by $\exp(-2\pi i(sx/L_1 + ty/L_2))$, which turns such sections into periodic functions, and applying the usual complete Fourier basis of the rectangular torus.

There is a zero mode exactly for the trivial character: a zero eigenvalue requires $m+s = n+t = 0$ for integers m, n. For every nontrivial character the lowest eigenvalue is positive, but it tends to zero as the character approaches the trivial one. This example simultaneously exhibits a moduli parameter, a jump in a kernel dimension, and failure of a uniform spectral bound over moduli. It can be plotted exactly, without numerically solving a differential equation. A diagram should label its two displayed automorphy coordinates and should describe a four-dimensional parameter view as a projection or slice if side lengths are varied too.

The corresponding Dolbeault operator $\overline{\partial}_{L_\chi} : \Omega^{0,0}(L_\chi) \to \Omega^{0,1}(L_\chi)$ is an elliptic complex on the complex curve. Riemann–Roch gives index zero for a degree-zero line bundle on a genus-one curve. Thus the index remains constant while individual kernel dimensions and positive eigenvalues vary. The explicit flat spectrum already demonstrates the crucial distinction. Atiyah–Singer index theory organizes stable topological information; it does not fix every eigenvalue in a family.

22.4 How this bridge relates to cryptography

The finite n-torsion characters above form $(\mathbb{Z}/n\mathbb{Z})^2$, and their evaluation pairing is the familiar finite Fourier pairing. A quantum Fourier transform can sample characters of a finite abelian group when that group is presented with efficient reversible operations. The mathematical resemblance is exact at the level of the pairing. The complexity statement needs more data: an encoding of elements, circuits for the group law, access to the hidden function, and a cost model. A geometric interpretation of characters does not automatically provide those resources.

For cryptographic elliptic curves over $\mathbb{F}_q$, the public computational problem usually concerns a specified finite subgroup of $E(\mathbb{F}_q)$ and scalar multiplication in that subgroup. The complex torus model instead uses a lattice in $\mathbb{C}$. Relating the two requires arithmetic reduction or lifting with its own hypotheses, and a proposed algorithm must preserve the relevant unknown and control the cost of every conversion. No such efficient conversion is supplied by the flat-line-bundle example. Its contribution is a rigorous conceptual connection, useful for organizing structures and formal interfaces, rather than a claimed cryptanalytic improvement.

There is a parallel lesson for lattice-based cryptography. A module over an order can be equipped with a positive form and embedded as a Euclidean lattice. Its dual then represents characters, and Fourier methods can be expressed geometrically. Hardness, however, concerns quantitative approximation or average-case distributions. The purely algebraic existence of a dual or a character basis does not bound the complexity of finding a shortest vector, recovering a noisy secret, or implementing the needed state preparation. Those are additional statements and belong in the algorithmic part of the research program.

22.5 A precise synthesis and a testable research program

The preceding constructions establish a network of actual mathematical maps. Orders act on integral lattices; a complex multiplication order acts on an elliptic torus; pullback acts contragrediently on its unitary characters; characters specify flat line bundles; their covariant differentials form complexes; elliptic symbols yield index classes; and finite gauge discretizations provide a setting for quantum reconstruction when the necessary analytic limits are proved. Each arrow has a defined

domain and codomain. A useful new result must strengthen one of those arrows, or supply a missing quantitative estimate, rather than merely draw the same objects closer together in a diagram.

Three research directions are consequently concrete. First, formalize the elementary matrix, character, and deformation identities and connect them to existing verified libraries through explicit equivalences. Second, develop finite algebraic presentations of selected gauge-observable relations and check their compatibility with one pinned continuum construction. Third, investigate quantitative inequalities on a total observable domain, retaining dependence on cutoff, volume, normalization, and phase. Success in the first direction provides reliable infrastructure; success in the second may improve exposition and auditing of the construction; a sufficiently strong success in the third could address the remaining mass-gap question.

These directions have different completion criteria. A matrix identity is complete when its stated assumptions yield its conclusion in the kernel. A compatibility theorem is complete when the diagram of maps commutes for the actual objects used by the external construction. A mass-gap argument is complete only when its uniform estimates and limit passages prove (20.3) for the required nontrivial quantum Yang–Mills theory. Keeping those criteria explicit allows substantial intermediate progress to be recognized without relabeling it as a solution of an unresolved problem.

Chapter 23

A character family inside a variety of complexes

There is a concrete bridge from a complex multiplication order to a parameterized variety of complexes. The bridge uses characters of the underlying rank-two lattice, and its elementary algebra can be written down completely. It supplements the continuous flat-torus spectrum with a finite cochain model, an explicit pullback map, and a quantitative contracting homotopy. It also supplies a useful example of cohomology jumping under specialization. No cryptographic complexity improvement or interacting quantum-field spectral estimate is asserted.

23.1 A universal complex over the character torus

Put $S = \mathbb{C}[X^{\pm 1}, Y^{\pm 1}]$, the Laurent polynomial ring, and write $A = X - 1$, $B = Y - 1$. Consider the cochain complex

$$0 \longrightarrow S \xrightarrow{D_0} S^2 \xrightarrow{D_1} S \longrightarrow 0, \qquad D_0 = \begin{pmatrix} A \\ B \end{pmatrix}, \quad D_1 = \begin{pmatrix} -B & A \end{pmatrix}. \tag{23.1}$$

Its degrees are $0, 1, 2$. The identity $D_1 D_0 = 0$ follows from commutativity. This is a Koszul complex for two regular parameters, unlike the periodic resolution in the worked nodal-order example.

Proposition 23.1 (Universal exactness and the exceptional fiber)**.** *The universal complex has $H^0 = H^1 = 0$ and $H^2 = S/(X-1, Y-1) \cong \mathbb{C}$. After specialization at $(u, v) \in (\mathbb{C}^\times)^2$, its differentials are*

$$d_0 = \begin{pmatrix} u - 1 \\ v - 1 \end{pmatrix}, \qquad d_1 = \begin{pmatrix} -(v-1) & u - 1 \end{pmatrix}.$$

The specialized complex is exact if $(u, v) \neq (1, 1)$. At $(1, 1)$ its cohomology dimensions are $(1, 2, 1)$.

Proof. The ring S is a domain, so D_0 is injective. If $-Bp + Aq = 0$, reduction modulo A gives $B\overline{p} = 0$ in $S/(A) \cong \mathbb{C}[Y^{\pm 1}]$. The latter ring is a domain and B is nonzero in it. Thus $p = Ar$ for some $r \in S$; substitution and cancellation of A give $q = Br$. This proves $\ker D_1 = \operatorname{im} D_0$. The image of D_1 is the ideal (A, B), proving the cokernel assertion.

In a specialized fiber, put $a = u - 1$, $b = v - 1$. If $(a, b) \neq (0, 0)$, the first map has rank one, the second is surjective, and their composite vanishes. The kernel of the second therefore has dimension one and equals the image of the first. If $a = b = 0$, both maps vanish and the cohomology is the entire graded vector space. □

The assignment $(u, v) \mapsto (d_0, d_1)$ is a regular map from an algebraic torus to the variety of complexes with dimension vector $(1, 2, 1)$. Its image lies in the equation $d_1 d_0 = 0$. The generic rank pair is $(1, 1)$; the identity character maps to rank pair $(0, 0)$. These are the explicit rank and determinantal data that connect the example to the general theory of varieties of complexes [27, 40, 41].

The distinction between the universal cohomology and the fiber cohomology is essential. Evaluation at $(1, 1)$ is a nonflat base change. Indeed, appending $S \to \mathbb{C}$ to (23.1) and reversing the degree indexing gives a free resolution of the augmentation module $\mathbb{C}$. Tensoring it with that module makes every differential zero, so

$$\operatorname{Tor}_0^S(\mathbb{C}, \mathbb{C}) = \mathbb{C}, \qquad \operatorname{Tor}_1^S(\mathbb{C}, \mathbb{C}) = \mathbb{C}^2, \qquad \operatorname{Tor}_2^S(\mathbb{C}, \mathbb{C}) = \mathbb{C},$$

and higher Tor groups vanish. The exceptional fiber records derived information, rather than contradicting the universal exactness calculation.

23.2 Characters, local coefficients and sign conventions

Let $\Lambda = \mathbb{Z}\omega_1 + \mathbb{Z}\omega_2 \subset \mathbb{C}$, and $E = \mathbb{C}/\Lambda$. A unitary representation $\rho : \Lambda \to U(1)$ is determined by $u = \rho(\omega_1)$ and $v = \rho(\omega_2)$. Here ρ denotes parallel-transport holonomy. A bundle with this convention is

$$(\mathbb{C} \times \mathbb{C})/((z, w) \sim (z + \lambda, \rho(\lambda)^{-1} w)).$$

Consequently its sections have automorphy multiplier ρ^{-1}. The flat-torus chapter uses the automorphy character χ; the conventions agree by setting $\rho = \chi^{-1}$.

A torus has a cell presentation with one vertex, two oriented edges and one oriented face. With a consistent choice of lifted cells, its cochains with these rank-one coefficients have the displayed differences $u - 1$ and $v - 1$, and the oriented face supplies the row $(-(v - 1), u - 1)$. Changing all character conventions to their inverses gives an equivalent model and leaves the conclusions below unchanged.

The cohomology dimensions also follow directly in the smooth model. Write $u = e^{2\pi i s}$, $v = e^{2\pi i t}$, and use real coordinates $z = x_1\omega_1 + x_2\omega_2$. Coefficient functions have Fourier frequencies $k = (m - s, n - t)$, with $(m, n) \in \mathbb{Z}^2$. On each mode the flat exterior derivative is exterior multiplication by $2\pi i(k_1 dx_1 + k_2 dx_2)$. When $k \neq 0$, contraction by $k_1\partial_{x_1} + k_2\partial_{x_2}$, divided by $2\pi i|k|^2$, is a contracting homotopy: the exterior multiplication and contraction anticommutator equals $|k|^2$ times the identity. For a fixed nontrivial character every mode is nonzero, with frequencies bounded away from zero. The resulting multipliers preserve rapidly decreasing Fourier coefficients and therefore smooth forms. All cohomology vanishes. For the trivial character, only the constant mode survives and has dimensions $(1, 2, 1)$. Thus the finite and differential complexes exhibit precisely the same cohomology jump.

23.3 A complete Hodge calculation

Equip the three finite cochain spaces with their standard Hermitian inner products. For $a = u - 1$, $b = v - 1$, put

$$\sigma(u, v) = |a|^2 + |b|^2.$$

Proposition 23.2 (Finite Laplacians and a quantitative homotopy)**.** *The degree-zero, degree-one and degree-two Hodge Laplacians are*

$$\Delta_0 = \sigma I_1, \qquad \Delta_1 = \sigma I_2, \qquad \Delta_2 = \sigma I_1.$$

For $\sigma > 0$, the maps $h^1 = d_0^*/\sigma$ and $h^2 = d_1^*/\sigma$ satisfy $dh + hd = I$. Their operator norms are $\sigma^{-1/2}$.

Proof. The endpoint products $d_0^* d_0$ and $d_1 d_1^*$ equal σ. At the middle term,

$$d_0 d_0^* + d_1^* d_1 = \begin{pmatrix} |a|^2 & a\bar{b} \\ b\bar{a} & |b|^2 \end{pmatrix} + \begin{pmatrix} |b|^2 & -\bar{b}a \\ -\bar{a}b & |a|^2 \end{pmatrix} = \sigma I_2.$$

Division by σ gives the homotopy identities in every degree. Each nonzero row or column differential has operator norm $\sqrt{\sigma}$, proving the norm assertion for its adjoint divided by σ. □

For unitary characters,

$$\sigma = 4\sin^2(\pi s) + 4\sin^2(\pi t).$$

Choose $s, t \in [-1/2, 1/2]$ and let $\delta^2 = s^2 + t^2$. Concavity of sine on $[0, \pi/2]$ and the inequality $\sin x \le x$ give

$$16\delta^2 \le \sigma \le 4\pi^2 \delta^2.$$

This supplies explicit estimates for the failure of a uniformly bounded contracting homotopy as the character approaches the identity. The Euler characteristic remains $1 - 2 + 1 = 0$ throughout. Exactness holds at every nontrivial character, but its quantitative inverse becomes unbounded. The number σ belongs to this chosen finite cellular metric; it is not the continuum Laplacian eigenvalue formula or a Yang–Mills mass.

23.4 An order action with an explicit cochain map

Suppose $\alpha\Lambda \subseteq \Lambda$ and write its matrix on column lattice coordinates as $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$. Pullback by $[z] \mapsto [\alpha z]$ sends the character pair to

$$(u', v') = (u^a v^c, u^b v^d). \tag{23.2}$$

This assertion concerns the multiplier ring of a complex lattice; when that ring has rank two it is an imaginary quadratic order, as proved in the flat-torus chapter.

For every integer r, define the Laurent polynomial $S_r(z) = (z^r - 1)/(z - 1)$, with its polynomial value at $z = 1$. For $r > 0$ it is $1 + \cdots + z^{r-1}$, for $r = 0$ it is zero, and for $r < 0$ it is $-\sum_{j=r}^{-1} z^j$. In all cases $S_r(1) = r$. Set

$$F_M(u, v) = \begin{pmatrix} v^c S_a(u) & S_c(v) \\ v^d S_b(u) & S_d(v) \end{pmatrix}.$$

Theorem 23.3 (A character pullback bridge)**.** *The maps 1, $F_M(u, v)$ and $\det F_M(u, v)$ in degrees $0, 1, 2$ form a cochain map from the complex at (u, v) to the complex at (u', v'). At the trivial character the induced maps are 1, M^{T} and $\det M$, respectively. If $\alpha \neq 0$, precisely $\det M = |\alpha|^2$ unitary characters pull back to the trivial character.*

Proof. The identity $u^a v^c - 1 = v^c S_a(u)(u - 1) + S_c(v)(v - 1)$, and its counterpart for (b, d), show that $d_0' = F_M d_0$. For any two-by-two matrix F and column x, the alternating row associated to Fx, multiplied by F, equals $\det(F)$ times the alternating row associated to x. This follows by expanding the two coordinates, or from $F^{\mathsf{T}} J F = (\det F) J$ with $J = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$. Hence $d_1' F_M = (\det F_M) d_1$, proving the second cochain identity. At $(1, 1)$ the differentials vanish and $F_M(1, 1) = \left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right) = M^{\mathsf{T}}$, proving the stated cohomology maps.

Multiplication by a nonzero complex number has real determinant $|\alpha|^2$. Changing from a real orthonormal basis to a lattice basis does not change its determinant, so $[\Lambda : \alpha\Lambda] = \det M = |\alpha|^2$. A

character pulls back trivially exactly when it is trivial on $\alpha\Lambda$, equivalently when it factors through the finite group $\Lambda/\alpha\Lambda$. A finite abelian group has as many unitary characters as elements: decompose it into cyclic factors and choose one root of unity for each generator. This proves the count. □

The theorem makes an order action, a character pairing, a variety of complexes, and a cohomology jump coexist in a single explicit diagram. Its cochain map is an algebraic representative obtained from the displayed factorization; no choice of a hidden scalar is required. For example, $\alpha = n$ kills exactly the n^2 torsion characters. Several initially acyclic character complexes consequently map to a fiber with nonzero cohomology. The map is not asserted to be a cohomology isomorphism.

These finite torsion characters also carry the evaluation pairings used in finite Fourier analysis. The present theorem supplies their algebraic structure and functorial maps. An efficient encoding, a hidden-function oracle, and a resource analysis remain separate requirements for a quantum algorithm. Similarly, the explicit estimate for σ is a bound in a specified finite complex. Connecting it to a physical Hamiltonian would require a proved operator comparison and estimates uniform in the relevant limits. The bridge is substantial precisely because its conclusions and its remaining obligations can both be stated exactly.

Part VI

Information reduction, fluid dynamics and the released singularity proof

Chapter 24

Microscopic information, kinetic limits, and fluid closure

Evidence status. The 249 timed English transcript segments of the requested braintruffle video were reviewed and independently reread in the public YouTube transcript interface. The first entry starts at 0:00 and the last at 32:36; the latter is an entry start, not the video's duration. An exact timestamped UTF-8 extraction has a recorded SHA-256 digest and a verified private source-cache copy. The public audit retains provenance without republishing the transcript. The timestamped source audit records precisely which accompanying papers, lecture decks, and book sections were accessible. The full main texts of Stodolna, Vozzi, Itatani, Allen, the abridged English Jönsson article, and all six requested Volkov decks were reviewed. Several other requested originals were accessible only as abstracts or metadata; these are not represented as full-text reviews. The arguments proved below are ordinary mathematical proofs. Their presence here does not assert a Lean formalization of the continuum theories.

24.1 What the microscopic perspective establishes

The video organizes fluid simulation around successive reductions in the information retained about a physical system: a quantum state, molecular positions and velocities, a velocity distribution, and a few macroscopic fields [25]. This is a productive organizing principle provided that each reduction is supplied with its own assumptions. An exact change of coordinates, a statistical marginal, an asymptotic limit, and an empirical constitutive law are mathematically different operations. The distinction determines which conclusions can be carried from one level to another.

For example, expanding a wave function in a complete eigenbasis retains all information. Keeping only finitely many modes discards information. Taking the partial trace of a many-body density operator discards correlations. Replacing a quantum nuclear state by a classical trajectory requires a localization or semiclassical approximation. Replacing an N-particle classical distribution by a one-particle distribution requires a closure of its correlation hierarchy. Replacing that distribution by density, velocity, and temperature requires another closure. None of these steps follows merely from a count of variables.

The central mathematical lesson is therefore a sequence of conditional statements. The Born–Huang expansion identifies the terms discarded by an adiabatic approximation. Liouville's equation identifies the correlations discarded by a kinetic approximation. Collision invariants identify the moments retained by a fluid approximation. An entropy inequality restricts admissible closures but does not determine all their coefficients. A hydrodynamic-limit theorem supplies compactness and

convergence information that a formal moment calculation does not supply. The structural results in Chapter 26 and the quantum algorithms in Chapter 27 fit into these particular stages.

24.1.1 Five distinct mathematical state spaces

For N distinguishable spinless particles in three dimensions, a pure quantum state is a unit vector in $L^2(\mathbb{R}^{3N})$, up to an overall phase. An N-particle classical microstate is a point of $\mathbb{R}^{6N}$. An uncertain classical state is instead a probability measure on that same phase space; when it has a density, that density is a function of $6N$ variables. Thus the video's contrast between a wave function of $3N$ coordinates and a classical list of $6N$ numbers compares a general quantum state with a single classical microstate. A comparison of statistical descriptions must also include classical ensemble distributions.

The one-particle kinetic density $f(t, x, v)$ is a function on $\mathbb{R}^3_x \times \mathbb{R}^3_v$. In the unit-mass convention used below, its basic moments are

$$\rho = \int f \, dv, \qquad \rho u = \int v f \, dv, \qquad 3\rho\theta = \int |v - u|^2 f \, dv. \tag{24.1}$$

Here $\theta = k_B T/m$ in physical units. The fluid variables (ρ, u, θ) contain five scalar fields. They generally do not determine f, its pressure tensor, or its heat flux.

In the quantum setting, mixed states are positive trace-one operators ϱ. A reduced one-particle density matrix and a classical kinetic density are different kinds of objects. A Wigner transform can relate them in suitable semiclassical regimes, but a generic Wigner function is not a nonnegative probability density. A classical probabilistic interpretation therefore cannot be inserted without a positivity or limiting argument.

24.2 Measurement, phase, and the limits of tomography

The experiments cited by the video support wave interference and model-dependent reconstruction of electronic structure. Jönsson's multi-slit experiment compares electron diffraction with interference predictions [72]. Stodolna and collaborators image nodal structures in an outgoing electron distribution associated with hydrogen Stark states; the detector distribution is related to a propagated wave by the apparatus and its model [155]. Itatani's molecular-orbital tomography uses high-harmonic emission, alignment information, a reference atom, and a recollision approximation [67]. Vozzi's generalized tomography improves phase recovery and the treatment of Coulomb distortion while retaining finite-bandwidth and reconstruction assumptions [176].

These are not measurements of a particle trajectory or unrestricted tomography of an arbitrary many-electron wave function. In particular, a single orbital reconstruction does not by itself determine every many-particle correlation. The distinction can be expressed without depending on an interpretation of quantum mechanics.

Proposition 24.1 (Finite-dimensional informational completeness)**.** *Let $E_1, \dots, E_s$ be positive operators on $\mathbb{C}^d$ satisfying $\sum_i E_i = I$. The probabilities $p_i = \operatorname{Tr}(E_i \varrho)$ uniquely determine every density operator ϱ if and only if the real linear span of the E_i is the space of Hermitian $d \times d$ matrices.*

Proof. If the span is the full Hermitian space and two states give the same probabilities, their difference X is orthogonal to every Hermitian matrix in the inner product $\operatorname{Tr}(XY)$. Taking $Y = X$ gives $\operatorname{Tr}(X^2) = 0$, hence $X = 0$.

Conversely, if the span is proper, choose a nonzero Hermitian matrix X orthogonal to it. Since $I = \sum_i E_i$ belongs to the span, $\operatorname{Tr} X = 0$. For sufficiently small $a > 0$, both $\varrho_\pm = I/d \pm aX$ are positive trace-one matrices. They are distinct and satisfy $\operatorname{Tr}(E_i\varrho_+) = \operatorname{Tr}(E_i\varrho_-)$ for every i. □

Position probabilities alone fail this criterion. The qubit vectors $(|0\rangle + |1\rangle)/\sqrt{2}$ and $(|0\rangle - |1\rangle)/\sqrt{2}$ have identical computational-basis probabilities and different relative phase. Similarly, $\psi(x)$ and $e^{i\alpha(x)}\psi(x)$ have the same $|\psi(x)|^2$ where both define admissible states, but may have different currents and future dynamics. The Born density for a complex wave function is $|\psi|^2$, not its algebraic square.

Proposition 24.2 (Local operations do not transmit a signal through a marginal)**.** *Let ϱ_{AB} be a density operator on a finite tensor product. For operators K_j on B satisfying $\sum_j K_j^* K_j = I_B$, define*

$$\widetilde{\varrho}_{AB} = \sum_j (I_A \otimes K_j)\varrho_{AB}(I_A \otimes K_j^*).$$

Then $\operatorname{Tr}_B \widetilde{\varrho}_{AB} = \operatorname{Tr}_B \varrho_{AB}$.

Proof. For every observable A on the first factor, cyclicity of the trace gives

$$\operatorname{Tr}((A \otimes I)\widetilde{\varrho}) = \sum_j \operatorname{Tr}((A \otimes K_j^* K_j)\varrho) = \operatorname{Tr}((A \otimes I)\varrho).$$

Equality against every A determines the reduced operator. □

Conditioning on a recorded outcome can change a conditional state. The outcome record must then be communicated before the other observer can condition on it. Bell correlations consequently do not justify an assertion of controllable instantaneous communication. This qualification is needed when using the introductory discussion in van Dommelen's book [168].

24.2.1 What the multi-slit formula proves

In the Fraunhofer approximation, suppose s equally illuminated identical slits contribute a common amplitude $a(\vartheta)$ and a phase increment $\delta = 2\pi d \sin\vartheta/\lambda$ between neighboring slits. Linearity gives

$$A(\vartheta) = a(\vartheta) \sum_{j=0}^{s-1} e^{ij\delta}.$$

The geometric-series identity and continuity at removable singularities give

$$|A(\vartheta)|^2 = |a(\vartheta)|^2 \frac{\sin^2(s\delta/2)}{\sin^2(\delta/2)}. \tag{24.2}$$

For a uniformly illuminated slit of width b, the Fourier integral $\int_{-b/2}^{b/2} e^{2\pi i x \sin\vartheta/\lambda}\, dx$ supplies a squared sinc envelope. Equation (24.2) therefore proves the interference factor under stated coherence, illumination, and far-field assumptions. It does not remove the need to model the source, finite detector resolution, or partial coherence. The reviewed 1974 article explicitly identifies itself as an abridged translation; it cannot stand in for a complete review of the longer 1961 German original [71, 72].

24.3 Exterior powers are a physical state constraint

The project's exterior algebra is directly relevant to identical fermions. Let V be a one-particle space including spin. The antisymmetric N-particle subspace is $\bigwedge^N V$, identified with the antisymmetric tensors inside $V^{\otimes N}$. Exchanging two complete spin-space labels changes a fermionic state by a sign. Equality of position probabilities under an exchange, by itself, does not specify that sign or the spin dependence.

Proposition 24.3 (Slater normalization and dimension)**.** *If $e_1, \dots, e_M$ are an orthonormal basis of the finite-dimensional one-particle space V, then for $1 \le i_1 < \cdots < i_N \le M$ the tensors*

$$\Psi_{i_1 \cdots i_N} = \frac{1}{\sqrt{N!}} \sum_{\sigma \in S_N} \operatorname{sgn}(\sigma) e_{i_{\sigma(1)}} \otimes \cdots \otimes e_{i_{\sigma(N)}}$$

form an orthonormal basis of $\bigwedge^N V$. In particular, $\dim \bigwedge^N \mathbb{C}^M = \binom{M}{N}$*, with the dimension zero for $N > M$.*

Proof. For fixed distinct indices the $N!$ elementary tensors in the sum are mutually orthonormal, so the squared norm is $N!/N! = 1$. Different index sets have disjoint elementary-tensor supports. Antisymmetry forces the coefficient of any tensor with a repeated index to vanish, and relates all coefficients belonging to permutations of a distinct index set by the permutation sign. Thus every antisymmetric tensor is a linear combination of the displayed vectors. □

In position representation this is the normalized Slater determinant. The sign representation of the symmetric group selects the Young diagram (1^N), while symmetric bosonic tensors select (N). More general Schur functors organize symmetry types of tensor features and multi-index responses. This is a precise use of the algebraic material developed earlier in the monograph: a representation specifies allowed states or tensor components. It does not, on its own, furnish a bound for a nonlinear fluid equation.

24.3.1 Number states and coherent states

Mathur's accessible introductory modules are useful for interference, tensor products, and the distinction between classical amplitude and measurement probabilities [95]. One teaching identification needs correction: a state with a fixed number of identical bosons in one mode is a number state, and is not a coherent state. Let $a|n\rangle = \sqrt{n}\,|n-1\rangle$. A coherent state is

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle. \tag{24.3}$$

The exponential series shows that its squared norm is one. Shifting the summation index shows $a|\alpha\rangle = \alpha|\alpha\rangle$. Its number probabilities are $e^{-|\alpha|^2}|\alpha|^{2n}/n!$, whose mean and variance are both $|\alpha|^2$. A fixed $|n\rangle$ has number variance zero and $\langle a \rangle = 0$. Quantization of the number observable therefore does not discretize every complex amplitude or every field expectation. This matters when deciding what a quantum neural model's input amplitudes and measured outputs represent.

24.4 Born–Huang expansion before Born–Oppenheimer truncation

The video's move from electronic wave functions to molecular potential energy surfaces is best made precise by first writing the coupling terms that are later neglected. The treatment in van

Dommelen's Section 9.2 and derivation D.51 distinguishes electronic and nuclear quantum dynamics; Allen's molecular-dynamics notes then describe classical motion on a chosen potential [4, 168].

For an exact, domain-transparent derivation, take a finite electronic space $\mathbb{C}^r$, nuclear coordinate $R \in \mathbb{R}^d$, and a C^2 family of Hermitian matrices $H_e(R)$. Suppose an orthonormal eigenbasis $\phi_n(R)$ can be chosen C^2 on the open region being considered. For C^2 nuclear coefficient functions χ_n, set

$$H = -\frac{\hbar^2}{2M}\Delta_R + H_e(R), \qquad \Psi(R) = \sum_{n=1}^{r} \chi_n(R)\phi_n(R).$$

Scalar nuclear repulsion can be included in H_e. The corresponding continuum-electronic formulas require common operator domains and the regularity needed to justify these derivatives; the finite statement does not assert those analytic hypotheses automatically.

Proposition 24.4 (Exact coupled nuclear equations)**.** *Set*

$$d_{mn} = \langle \phi_m, \nabla_R \phi_n \rangle,$$
$$\tau_{mn} = \langle \phi_m, \Delta_R \phi_n \rangle.$$

Then $H\Psi = E\Psi$ is equivalent to

$$-\frac{\hbar^2}{2M}\left(\Delta\chi_m + 2\sum_n d_{mn}\cdot\nabla\chi_n + \sum_n \tau_{mn}\chi_n\right) + E_m(R)\chi_m = E\chi_m \tag{24.4}$$

for each m. Moreover, $d_{mn} = -\overline{d_{nm}}$ componentwise.

Proof. Apply the product rule to each $\Delta(\chi_n\phi_n)$ and take its electronic inner product with ϕ_m. Orthonormality gives the first term and diagonalizes H_e. Differentiating $\langle \phi_m, \phi_n \rangle = \delta_{mn}$ proves the final identity. □

The usual single-surface approximation retains one electronic sector and controls or omits its couplings to others. Which diagonal geometric terms are retained must be specified. Rephasing $\phi_n(R) \mapsto e^{i\gamma_n(R)}\phi_n(R)$ also rephases χ_n and changes diagonal derivative couplings; a claim that they vanish everywhere needs a gauge and topology argument.

Proposition 24.5 (Energy gradients and the gap denominator)**.** *Let $E_n(R)$ be a simple normalized eigenpair of the finite matrix family. For a coordinate derivative ∂_j,*

$$\partial_j E_n = \langle \phi_n, (\partial_j H_e)\phi_n \rangle, \qquad \langle \phi_m, \partial_j \phi_n \rangle = \frac{\langle \phi_m, (\partial_j H_e)\phi_n \rangle}{E_n - E_m} \quad (m \neq n).$$

If $|E_n - E_m| \geq \gamma > 0$ for all $m \neq n$, then

$$\|(I - |\phi_n\rangle\langle\phi_n|)\partial_j\phi_n\| \leq \gamma^{-1}\|\partial_j H_e\|.$$

Proof. Differentiate $H_e\phi_n = E_n\phi_n$ and take inner products with ϕ_n and with ϕ_m. The terms containing $\partial_j\phi_n$ cancel in the first case. In the second they give the difference $E_n - E_m$. Squaring the off-diagonal coefficients and summing over the orthonormal eigenbasis gives the stated norm bound. □

Large nuclear mass reduces the prefactor in (24.4); it does not prevent a small electronic gap from amplifying the derivative coupling. Near crossings and conical intersections, a single-surface argument requires additional work. Likewise, an excited electronic surface may have an adiabatic nuclear description. The Born–Oppenheimer approximation is not a theorem that electrons automatically dissipate into their ground state.

Proposition 24.6 (Unitary evolution does not supply ground-state relaxation)**.** *For a time-independent finite Hermitian matrix H and a spectral projector P, the population*

$$\mathrm{Tr}(Pe^{-itH/\hbar}\varrho_0 e^{itH/\hbar})$$

is constant.

Proof. The projector commutes with both exponentials. Cyclicity of the trace reduces the expression to $\mathrm{Tr}(P\varrho_0)$. □

A bath, emitted radiation, collisions with other degrees of freedom, or a specified open-system approximation can generate relaxation. Adiabatic following and relaxation are distinct physical mechanisms. The video at 21:37–21:53 compresses them into a single explanation; equations (24.4) and Proposition 24.6 state what must be separated.

24.5 When a quantum packet behaves approximately as a particle

For $H = P^2/(2M) + V(X)$, the formal Ehrenfest identities are

$$\frac{d}{dt}\langle X\rangle = \frac{\langle P\rangle}{M}, \qquad \frac{d}{dt}\langle P\rangle = -\langle \nabla V(X)\rangle. \tag{24.5}$$

They follow from $d\langle A\rangle/dt = (i/\hbar)\langle [H, A]\rangle$ on a common invariant domain. For instance, they hold for sufficiently regular rapidly decaying solutions for which all integrations by parts and differentiated moments are justified. They are not yet Newton's equations for the mean, because $\langle \nabla V(X)\rangle$ need not equal $\nabla V(\langle X\rangle)$.

Proposition 24.7 (Localization controls the mean-force error)**.** *If ∇V is L-Lipschitz and the position distribution has a finite second moment, then*

$$|\langle \nabla V(X)\rangle - \nabla V(\langle X\rangle)| \le L\left\langle |X - \langle X\rangle|^2\right\rangle^{1/2}.$$

Proof. Apply the Lipschitz inequality inside the position integral, then the triangle inequality and Cauchy–Schwarz for its probability measure. □

This gives a quantitative condition rather than a universal particle-size threshold. To turn it into a trajectory approximation on a time interval, one must control the packet width throughout that interval and estimate the resulting perturbation of the classical flow.

For a free particle in one dimension with $\sigma > 0$ and

$$\psi_0(x) = (2\pi\sigma^2)^{-1/4} e^{-x^2/(4\sigma^2)},$$

the exact Heisenberg identity $X(t) = X + tP/M$ gives

$$\mathrm{Var}(X(t)) = \sigma^2 + \frac{\hbar^2 t^2}{4M^2\sigma^2}. \tag{24.6}$$

Indeed $\mathrm{Var}\, X = \sigma^2$, $\mathrm{Var}\, P = \hbar^2/(4\sigma^2)$, and the initial symmetrized covariance vanishes. Increasing M slows free spreading for a fixed packet, but (24.6) still grows with time. Heavy particles can also exhibit interference. A suitable scale comparison involves action, wavelength, temperature, environmental coupling, and the observable being approximated.

24.5.1 One-dimensional Coulomb models require a domain

The video explicitly uses a softened one-dimensional potential to explain electronic forces. Such a model is pedagogically useful, but it is a different operator from the singular expression

$$-\frac{\hbar^2}{2m}\frac{d^2}{dx^2} - \frac{\kappa}{|x|}.$$

A differential expression does not specify a self-adjoint Hamiltonian until its domain and behavior at the singular point are given. The reviewed operator-domain analysis by de Oliveira and Verri finds a family of self-adjoint realizations for the one-dimensional operator initially defined away from the origin [42]. We use that domain distinction, not an unqualified assertion about the spectrum of every regularization. The requested Loudon papers are separately registered as unavailable full texts in this review [88, 89].

The limit of a family of softened potentials must specify the topology of operator convergence and any renormalization of energy. It cannot be inferred merely by drawing potentials that approach $-\kappa/|x|$ away from zero. This is the same structural issue that recurs in continuum limits of graph Laplacians: a formula for coefficients must be accompanied by domains, boundary conditions, and convergence estimates.

24.6 Molecular dynamics and conserved microscopic information

A classical molecular Hamiltonian has the form

$$H(q,p) = \sum_{i=1}^{N} \frac{|p_i|^2}{2m_i} + U(q), \qquad \dot{q}_i = \frac{p_i}{m_i}, \quad \dot{p}_i = -\nabla_{q_i} U. \tag{24.7}$$

The force is the negative gradient of potential energy. General molecular force fields include angular, torsional, many-body, constrained, or reactive terms; they need not be sums of pairwise distance functions. Allen's notes explain these modeling choices, symplectic integrators, constraints, thermostats, and statistical errors [4]. Their scope is much wider than the two-particle illustration in the video.

Proposition 24.8 (Liouville volume and fine-grained entropy)**.** *Let $H \in C^2(\mathbb{R}^{2n})$ have a smooth Hamiltonian flow Φ_t on the interval considered. Then* $\det D\Phi_t = 1$*. If a smooth probability density evolves by $f_t = f_0 \circ \Phi_{-t}$ and the entropy integral is absolutely integrable, then*

$$\int f_t \log f_t \, dq\, dp = \int f_0 \log f_0 \, dq\, dp.$$

Proof. For the Hamiltonian vector field $X_H = (\partial_p H, -\partial_q H)$, equality of mixed derivatives gives $\operatorname{div} X_H = \sum_i (\partial_{q_i}\partial_{p_i} H - \partial_{p_i}\partial_{q_i} H) = 0$. Jacobi's determinant identity for the flow derivative yields $\frac{d}{dt} \det D\Phi_t = (\operatorname{div} X_H) \circ \Phi_t \det D\Phi_t = 0$. Its initial value is one. A change of variables in the entropy integral proves the second assertion. □

Thus a Hamiltonian description preserves fine-grained information. Entropy growth in a reduced description is not the destruction of information by an invertible microscopic flow. It involves a marginalization, an environment, a limiting operation, or a restriction on what is resolved. Quantum unitary conjugation similarly preserves the spectrum of a density operator and hence its von Neumann entropy when that entropy is defined. A reduced subsystem can have changing entropy because its correlations with the omitted subsystem change.

24.6.1 A useful exact statement about velocity Verlet

For separable $H = T(p) + U(q)$, the kick $(q, p) \mapsto (q, p - h\nabla U(q))$ and drift $(q, p) \mapsto (q + h\nabla T(p), p)$ are exact Hamiltonian subflows. Their symmetric composition, kick by $h/2$, drift by h, kick by $h/2$, is symplectic. This proves preservation of the symplectic form; it does not prove exact preservation of the original H.

For the scalar oscillator $H = (p^2 + \omega^2 q^2)/2$, the update is

$$q' = \left(1 - \frac{h^2\omega^2}{2}\right) q + hp, \quad p' = \left(1 - \frac{h^2\omega^2}{2}\right) p - h\omega^2 \left(1 - \frac{h^2\omega^2}{4}\right) q.$$

Direct expansion proves that it preserves

$$\widetilde{H}_h(q, p) = \frac{1}{2}p^2 + \frac{1}{2}\omega^2 \left(1 - \frac{h^2\omega^2}{4}\right) q^2. \tag{24.8}$$

This quadratic form is positive definite when $0 < h\omega < 2$. It is a valid special-case modified invariant. In a general nonlinear system, a formal or asymptotic modified Hamiltonian from backward error analysis must not be promoted to an exact global invariant without the additional theorem needed for that claim.

24.7 Marginals, correlations, and the kinetic approximation

For a smooth symmetric N-particle density F_N, let $F_N^{(s)}$ denote its s-particle marginal. With unit masses and a pair force $K(x) = -\nabla\Phi(x)$, integrating the Liouville equation over all but particle 1 gives

$$\partial_t F_N^{(1)} + v \cdot \nabla_x F_N^{(1)} + (N-1)\nabla_v \cdot \int K(x-y) F_N^{(2)}(x, v, y, w)\, dy\, dw = 0, \tag{24.9}$$

assuming the boundary terms vanish. The two-particle marginal appears because the force on a particle depends on others. Its own equation involves the three-particle marginal. This is the BBGKY hierarchy.

Example 24.9 (A first marginal cannot determine a pair correlation)**.** On $\{0, 1\}^2$, let P give mass $1/2$ each to $(0, 0)$ and $(1, 1)$, and let Q give mass $1/2$ each to $(0, 1)$ and $(1, 0)$. Both first and second one-coordinate marginals are uniform. Nevertheless, $P(X = Y) = 1$ and $Q(X = Y) = 0$. Thus no function of the one-coordinate marginals alone can reconstruct all pair observables.

Replacing $F_N^{(2)}$ by a product is a substantive assumption. For dilute collisions it is the incoming, appropriately scaled correlations that matter. Hard-sphere exclusion already prevents exact independence at fixed particle diameter. The word "chaos" in propagation of chaos denotes an asymptotic factorization property of fixed-order marginals; it is not interchangeable with deterministic sensitivity to initial conditions.

24.7.1 A rigorous classical-to-kinetic bridge and its time interval

In the hard-sphere Boltzmann–Grad scaling in dimension $d \geq 2$,

$$N\varepsilon^{d-1} = 1, \qquad N\varepsilon^d = \varepsilon \longrightarrow 0.$$

The first relation retains an order-one collision scale, while the second sends the occupied volume fraction to zero. Gallagher, Saint-Raymond, and Texier give a detailed derivation for hard spheres

and suitable short-range repulsive potentials [53]. Their hard-sphere theorem propagates admissible asymptotic initial factorization on a positive short time interval, with convergence of velocity-averaged observables locally away from positional diagonals. It uses weighted uniform bounds on the hierarchy and quantitative control of recollisions.

The reviewed source's introductory statement assumes a continuous probability density f_0 with a Gaussian velocity bound

$$\sup_{x,v} f_0(x,v)e^{\beta|v|^2/2} < \infty \quad \text{for some } \beta > 0.$$

Its precise Theorem 8 includes the admissibility and hierarchy-space hypotheses. The proof does not assert convergence for arbitrary finite-N initial data, arbitrarily long times, every potential, or every boundary geometry. This chapter reviewed the introduction, main convergence statement, and concluding scope discussion; it does not claim to reproduce or independently verify the monograph's complete recollision proof.

This time restriction has a practical logical consequence. A short-time particle-to-Boltzmann theorem and a long macroscopic-time Boltzmann-to-fluid theorem cannot simply be concatenated. The former must hold uniformly on the time scales required by the latter, or a separate joint-limit theorem must be supplied. No amount of formal diagram composition replaces that analytic requirement.

24.8 The collision algebra and its five invariants

We now prove the algebraic facts behind the simplest elastic collision operator. They are useful both for conservative simulation and for checking learned collision models. Set $g = v - v_*$, let $n \in S^{d-1}$, and define

$$v' = v - (g \cdot n)n, \qquad v'_* = v_* + (g \cdot n)n. \tag{24.10}$$

Proposition 24.10 (Elastic reflection)**.** *The map* (24.10) *is an involution with absolute Jacobian determinant one. It preserves total momentum and kinetic energy:*

$$v' + v'_* = v + v_*, \qquad |v'|^2 + |v'_*|^2 = |v|^2 + |v_*|^2.$$

Proof. The center velocity $V = (v + v_*)/2$ is unchanged and $g' = g - 2(g \cdot n)n$. This is an orthogonal reflection, so $|g'| = |g|$, its square is the identity, and its determinant has absolute value one. Since $|v|^2 + |v_*|^2 = 2|V|^2 + |g|^2/2$, all assertions follow after the invertible linear change from (v, v_*) to (V, g). □

A scalar function ϕ is a collision invariant when $\phi(v) + \phi(v_*) = \phi(v') + \phi(v'_*)$ for every elastic collision. The functions $1, v_1, \dots, v_d, |v|^2$ are immediately invariant. The following proof makes the converse precise under a simple regularity assumption.

Proposition 24.11 (Classification of smooth collision invariants)**.** *For $d \geq 2$, every C^2 collision invariant has the form*

$$\phi(v) = a + b \cdot v + c|v|^2.$$

Consequently the scalar invariant space has dimension $d + 2$, hence five in three dimensions.

Proof. For fixed V and radius r, collisions can send g to any vector of the same length: reflection in a suitable hyperplane does so for two given vectors. Hence $\phi(V + r\omega) + \phi(V - r\omega)$ is independent

of the unit vector ω. Its second-order expansion at $r = 0$ implies that $\omega^T D^2\phi(V)\omega$ is constant on the unit sphere. Taking coordinate vectors and their pairwise sums shows $D^2\phi(V) = \lambda(V)I$.

The mixed second derivatives are zero, so $\partial_i\phi(v)$ depends only on v_i. The equality of all diagonal second derivatives then says that a function of v_i alone equals a function of v_j alone for independently variable v_i, v_j. Since $d \geq 2$, each is the same constant. Integrating twice yields the displayed quadratic polynomial. The converse follows from Proposition 24.10. □

Some introductory notes group the invariants into three physical categories: particle number, momentum, and energy. Momentum is a vector, so those categories contain five independent scalar invariants in three dimensions, not three scalar equations.

24.8.1 A collision kernel contains more information than a diameter

For a central classical interaction, the differential cross section is obtained from the map between impact parameter b and scattering angle χ. On a branch where the map is invertible,

$$\frac{d\sigma}{d\Omega} = \frac{b}{\sin\chi}\left|\frac{db}{d\chi}\right|.$$

If multiple branches produce the same angle, their contributions must be added. Angular singularities and grazing collisions require further analysis. The hard-sphere model and the variable hard-sphere model used in Volkov's notes do not reproduce every angular distribution of a general potential [171]. Matching a viscosity law or a mixing statistic is a calibration of particular observables, not an equality of the underlying scattering laws.

Internal excitation and chemical reactions add variables, transitions, and conservation relations. Parker's requested paper concerns rotational and vibrational relaxation, while Valentini and collaborators' requested study concerns nitrogen dissociation on an ab initio potential surface [119, 167]. Only their accessible abstracts and metadata were reviewed here. Their reported models must not be treated as fully audited derivations or as evidence that every gas can be represented by an elastic monatomic kernel.

24.9 The weak collision identity and entropy production

Let $B(g, n) \geq 0$ be a collision kernel invariant under exchange of the particles and under the pre/post-collision reflection. Define

$$Q(f, f)(v) = \int_{\mathbb{R}^d}\int_{S^{d-1}} B(v - v_*, n)(f'f'_* - ff_*)\, dn\, dv_*. \tag{24.11}$$

Here all four values of f have the same (t, x), when those variables are present. The next statements assume absolute integrability of the displayed integrals, so that Fubini's theorem and the changes of variables are legitimate.

Proposition 24.12 (Weak collision identity)**.** *For an admissible test function ϕ,*

$$\begin{aligned}\int \phi Q(f, f)\, dv &= \frac{1}{4}\iiint B(f'f'_* - ff_*)(\phi + \phi_* - \phi' - \phi'_*)\, dn\, dv_*\, dv \\ &= \frac{1}{2}\iiint Bff_*(\phi' + \phi'_* - \phi - \phi_*)\, dn\, dv_*\, dv.\end{aligned} \tag{24.12}$$

Proof. Exchange v and v_* in the original integral and average the two expressions. In a second copy, change variables from (v, v_*) to (v', v'_*). The measure and B are unchanged, the gain-minus-loss factor changes sign, and the test functions acquire primes. Averaging once more gives the first identity. Applying the same pre/post change to the gain term gives the second. □

Testing with collision invariants proves the mass, momentum, and energy collision balances. Testing with the logarithm supplies a stronger constraint, subject to the additional integrability needed for that unbounded test.

Theorem 24.13 (Classical homogeneous H theorem under explicit regularity)**.** *Let $f(t, v) > 0$ be a differentiable solution of $\partial_t f = Q(f, f)$ for which $\int f \log f$ is finite, differentiation under the integral is valid, and all terms in (24.12) with $\phi = \log f$ are integrable. Then*

$$\frac{d}{dt}\int f \log f \, dv = -D(f) \le 0, \quad D(f) = \frac{1}{4}\iiint B(f'f'_* - ff_*) \log \frac{f'f'_*}{ff_*} \, dn \, dv_* \, dv. \tag{24.13}$$

Proof. The derivative equals $\int (1+\log f) Q(f, f) \, dv$. The constant part vanishes by mass conservation. Use (24.12) for the remaining part. Since $(a-b)(\log a - \log b) \ge 0$ for $a, b > 0$, the displayed entropy production is nonnegative. □

For spatially varying smooth solutions, transport adds the entropy flux $\nabla_x \cdot \int v f \log f \, dv$. Boundary entropy fluxes must then be included. For weak or renormalized solutions, an entropy inequality is established by approximation and lower semicontinuity; one cannot simply insert an unjustified logarithmic test into a weak equation.

Proposition 24.14 (Finite energy does not imply finite positive entropy)**.** *There is a probability density supported in $[0, 1]^3$ with finite second moment and $\int f \log f = +\infty$.*

Proof. Let $c = 6/\pi^2$ and choose disjoint measurable sets A_n in $[0, 1]^3$ with volumes $|A_n| = ce^{-n^2}/n^2$. Such sets exist because the sum of their volumes is less than one. Define $f = e^{n^2}$ on A_n and $f = 0$ elsewhere, with $0 \log 0 = 0$. Then

$$\int f = c\sum_{n\ge 1} n^{-2} = 1, \qquad \int |v|^2 f \, dv \le 3, \qquad \int f \log f = c\sum_{n\ge 1} 1 = +\infty.$$

□

This corrects a gap in the elementary entropy discussion in the reviewed Volkov slides [172]. Similarly, integrability of a continuous density alone does not force pointwise decay to zero at infinity: separated continuous spikes of height one and summable widths provide a counterexample. Decay assumptions used in integration by parts must therefore be stated, not inferred solely from normalization.

24.9.1 Entropy minimization is not a convergence proof

Proposition 24.15 (Maxwellian minimizes H at fixed first two moments)**.** *Let f be a probability density with mean u and $\int |v-u|^2 f \, dv = d\theta$, where $\theta > 0$. Set*

$$M_{u,\theta}(v) = (2\pi\theta)^{-d/2} e^{-|v-u|^2/(2\theta)}.$$

Whenever the entropy expressions are defined,

$$\int f \log f - \int M_{u,\theta} \log M_{u,\theta} = \int f \log \frac{f}{M_{u,\theta}} \ge 0.$$

Equality holds exactly when $f = M_{u,\theta}$ almost everywhere.

Proof. The function $\log M_{u,\theta}$ is a quadratic polynomial in v, so its integral against f agrees with its integral against $M_{u,\theta}$. This proves the identity. With $r = f/M_{u,\theta}$, use $r \log r - r + 1 \geq 0$ and integrate against $M_{u,\theta}\, dv$. Equality in this strictly convex inequality requires $r = 1$ almost everywhere. □

Under a positive kernel, sufficient continuity, and $\log f \in C^2$, vanishing $D(f)$ forces detailed balance $f'f'_* = ff_*$. Proposition 24.11 applied to $\log f$ then gives the Maxwellian form, with a negative quadratic coefficient required by integrability.

These facts characterize entropy minimizers and stationary collision equilibria. They do not prove that every trajectory converges to one. If a scalar function $H(t)$ is nonincreasing and bounded below, it has a limit, but its derivative need not tend to zero. For example, take a nonnegative smooth function b supported in $(-1/4, 1/4)$ with $b(0) = 1$ and define

$$h(t) = \sum_{n \geq 1} b(2^n(t - n)), \qquad H(t) = \int_t^\infty h(s)\, ds.$$

The bumps have disjoint supports and summable areas. Hence H is smooth, nonnegative, and decreasing to zero, while $H'(n) = -1$. Even $H'(t) \to 0$ alone would not supply compactness of the state or identify a unique limit. Convergence-to-equilibrium theorems require the appropriate coercivity, compactness, integrability, and control of invariants in the stated regime.

24.10 Exact moment balances and the closure problem

For a smooth unit-mass solution of

$$\partial_t f + v \cdot \nabla_x f = Q(f, f),$$

assume sufficient velocity integrability to justify the following moments. Let $c = v - u$ and define

$$P = \int c \otimes c\, f\, dv, \quad e = \frac{1}{2\rho} \int |c|^2 f\, dv, \quad q = \frac{1}{2} \int |c|^2 c f\, dv, \quad E = \frac{1}{2}\rho|u|^2 + \rho e.$$

Here P is a positive semidefinite pressure tensor. The Cauchy stress uses the opposite sign for the kinetic contribution; keeping this convention explicit avoids sign errors in viscosity formulas.

Proposition 24.16 (Moment equations before closure)**.** *Under these hypotheses,*

$$\partial_t \rho + \nabla \cdot (\rho u) = 0, \tag{24.14}$$

$$\partial_t (\rho u) + \nabla \cdot (\rho u \otimes u + P) = 0, \tag{24.15}$$

$$\partial_t E + \nabla \cdot (Eu + Pu + q) = 0. \tag{24.16}$$

Proof. Test the kinetic equation with $1, v, |v|^2/2$ and use the collision invariants. Since $\int cf\, dv = 0$, expansion of $v = u + c$ gives $\int v \otimes v f\, dv = \rho u \otimes u + P$. Likewise,

$$\frac{1}{2} \int |v|^2 v f\, dv = \left(\frac{1}{2}\rho|u|^2 + \rho e\right) u + Pu + q.$$

These expansions give all three equations. □

Example 24.17 (Equal five moments, different stress)**.** Let f_A and f_B be centered normalized Gaussian densities on $\mathbb{R}^3$ with covariance matrices

$$A = I, \qquad B = \operatorname{diag}(1+\delta, 1-\delta, 1), \quad 0 < \delta < 1.$$

They have the same mass, mean, and scalar temperature because $\operatorname{tr} A = \operatorname{tr} B = 3$. Their pressure tensors are A and B, respectively. Thus the five fluid moments do not determine the momentum flux.

A local Maxwellian imposes $P = \rho\theta I$ and $q = 0$, giving the Euler fluxes. An arbitrary local Maxwellian with space- and time-dependent parameters is generally not a solution of the full Boltzmann equation: its collision term vanishes, but its transport term need not. Euler equations arise as moment equations under that closure or as a justified limit, rather than as a proof that every Euler solution lifts to an exact local-Maxwellian kinetic solution. This qualification is stated explicitly in the rigorous hydrodynamic-limit literature [54].

24.10.1 First-order constitutive laws and scaling

For a monatomic gas near local equilibrium, a formal Chapman–Enskog expansion produces the constitutive form

$$P = pI - 2\mu S_0 - \zeta(\nabla\cdot u)I, \qquad q = -\kappa\nabla T, \quad S_0 = \tfrac{1}{2}(\nabla u + \nabla u^T) - \tfrac{1}{3}(\nabla\cdot u)I. \tag{24.17}$$

The bulk coefficient is zero in the simplest monatomic elastic model. Internal modes and other models change this conclusion. The values of the transport coefficients require solving equations for the linearized collision operator; conservation and symmetry alone do not fix them.

The Knudsen number is the ratio of a microscopic mean free path to the macroscopic length scale. Small Knudsen number, small Mach number, and fixed Reynolds number specify distinct relationships among scales. A formal expansion in one does not silently impose the others. A global small value also does not rule out local breakdown near a shock, boundary layer, or strongly varying rarefied region. Volkov's first and third decks make the physical scale separation and the failure of elementary closure in rarefied regimes explicit [170, 172].

24.10.2 A precise historical hydrodynamic-limit theorem

The result of Golse and Saint-Raymond that is useful here concerns the scaled equation on $\mathbb{R}^3_x \times \mathbb{R}^3_v$

$$\varepsilon^2\partial_t F_\varepsilon + \varepsilon v\cdot\nabla_x F_\varepsilon = Q(F_\varepsilon, F_\varepsilon), \qquad M(v) = (2\pi)^{-3/2}e^{-|v|^2/2}. \tag{24.18}$$

Their collision kernels obey the exchange and reflection symmetries and hard-cutoff bounds of the form

$$0 < B(z,n) \le C(1+|z|)^\beta |\widehat{z}\cdot n|, \qquad \int_{S^2} B(z,n)\,dn \ge C^{-1}\frac{|z|}{1+|z|}, \quad 0 \le \beta \le 1,$$

with the inequalities understood almost everywhere. Initial data satisfy

$$\sup_{\varepsilon>0}\varepsilon^{-2}H(F_\varepsilon^{\mathrm{in}} \mid M) < \infty, \quad H(F\mid M) = \iint \left(F\log\frac{F}{M} - F + M\right)dx\,dv. \tag{24.19}$$

Their Theorem 2.4 describes limit points of the fluctuations $g_\varepsilon = (F_\varepsilon/M - 1)/\varepsilon$ for the corresponding renormalized solutions. In the stated weak local topology the limit has the form

$$g(t,x,v) = u(t,x)\cdot v + \tfrac{1}{2}\theta(t,x)(|v|^2 - 5),$$

where (u, θ) solves an incompressible Navier–Stokes–Fourier system and satisfies its energy inequality [54, Theorem 2.4]. The initial velocity is determined by the weak limit of the Leray projection of the scaled momentum; a corresponding thermal moment determines the initial temperature fluctuation. For appropriate well-prepared initial entropy, the velocity is a Leray solution with the corresponding initial energy bound.

This is a theorem about a specified collision class, spatial domain, scaling, solution concept, and convergence topology. Its proof uses entropy bounds, coercivity of the linearized collision operator, velocity averaging, vanishing conservation defects, and control of nonlinear fluxes. It yields weak fluid solutions; it does not by itself upgrade them to every stronger regularity or uniqueness statement. Statements about the separate 2026 proof release are analyzed in Chapter 25. That chapter's precise theorem must be used with its own forcing and solution hypotheses, rather than inferred from this historical limit theorem.

24.11 DSMC as a stochastic approximation, with auditable assumptions

The requested sixth Volkov deck explains direct simulation Monte Carlo as a particle approximation to kinetic dynamics, with statistical weights, free-flight and collision substeps, cell-based collision selection, inflow and wall kernels, and sampling of moments [175]. These simulated particles represent weighted samples. Their paths are not claimed to be the paths of every physical molecule in the gas.

Three error categories must be distinguished: the physical model (collision law and wall interaction), numerical discretization (time step, cells, particle count, cutoffs), and statistical estimation. Increasing a GPU sample count primarily addresses the latter two categories. It cannot validate a physical collision law by itself.

Proposition 24.18 (Flux sampling at an equilibrium wall)**.** *Let a stationary wall emit an equilibrium gas with thermal variance $\theta > 0$ and inward unit normal n. The normalized velocity density of molecules crossing the wall into the domain is*

$$p_{\mathrm{flux}}(v) = \frac{(v \cdot n)_+}{2\pi\theta^2} e^{-|v|^2/(2\theta)}.$$

Its positive normal component s has density $(s/\theta)e^{-s^2/(2\theta)}$, and its two tangential components are independent centered Gaussians of variance θ.

Proof. A bulk density $M_\theta(v)$ contributes a crossing flux $(v \cdot n)_+ M_\theta(v)$. The integral of this flux is $\sqrt{\theta/(2\pi)}$. Dividing by that value gives the displayed density. In orthonormal normal/tangential coordinates the exponential and its prefactor factor into the three stated one-dimensional densities. □

Sampling the normal velocity from an unweighted half-Gaussian would therefore give the wrong boundary ensemble. A finite reservoir implementation also requires an estimate of the missing high-speed tail; the Maxwellian has no finite maximum speed.

Proposition 24.19 (The exact frozen-rate majorant identity)**.** *Consider $K \geq 1$ possible events with fixed rates $r_j \geq 0$, and choose $R > 0$ with $R \geq \max_j r_j$. Generate a Poisson proposal process of rate KR, choose a label uniformly from $\{1, \ldots, K\}$ at each proposal, and accept label j with probability r_j/R. The accepted events of label j form a Poisson process of rate r_j; distinct accepted labels have independent counts.*

Proof. The probability that a proposal is accepted with label j is $r_j/(KR)$. For an interval of length t, condition on its Poisson number of proposals and take the joint probability generating function of the marked counts. Summing the Poisson series gives

$$\exp\left(t\sum_{j=1}^{K} r_j(z_j-1)\right)=\prod_{j=1}^{K}\exp(tr_j(z_j-1)).$$

This is the product of the generating functions of independent Poisson variables with the claimed means. The same argument on disjoint time intervals proves the process statement. □

This proves the sampling principle behind majorant collision schemes. It does not alone prove convergence of a full DSMC method whose rates change after collisions and whose transport is split in time. In that case the bound must remain valid for the current state and the discretization error must be controlled. If an alleged acceptance probability exceeds one, silently clipping it changes the rate. A temperature-based constant in illustrative code is not a proof of a global bound on an unbounded velocity distribution.

Proposition 24.20 (Sampling variance with correlations)**.** *If $Y_1,\dots,Y_n$ have finite second moments, then*

$$\operatorname{Var}\left(\frac{1}{n}\sum_i Y_i\right)=\frac{1}{n^2}\sum_{i,j}\operatorname{Cov}(Y_i,Y_j).$$

For identically distributed independent samples this equals σ^2/n. For a stationary sequence with variance $\sigma^2>0$ and lag correlations r_k, it equals

$$\frac{\sigma^2}{n}\left(1+2\sum_{k=1}^{n-1}(1-k/n)r_k\right).$$

Proof. Expand the square of the centered sum and take expectations. For a stationary sequence, group off-diagonal pairs by their lag; there are $n-k$ pairs at lag k in each direction. □

Thinning a time series can reduce some correlations but does not prove independence. A central-limit error bar requires its own assumptions; a fixed seed provides reproducibility, not randomness certification. For independent samples, Chebyshev's inequality gives the explicit bound $\Pr(|\overline{Y}-\mathbb{E}Y|\ge a)\le\sigma^2/(na^2)$. The statistics and Monte Carlo decks are useful introductions [173, 174]; implementations should additionally handle finite precision, zero arguments of logarithms, array bounds, correct random-number-generator state, and reproducible restart state. The period of a generator is constrained by its internal state space, not merely by the number of possible values of a single output word.

24.12 Reduced models and learned closures

An orthogonal reduced basis minimizes a chosen projection loss on its training snapshots. That statement does not imply that the projected trajectory is stable, that the discarded modes exert no force, or that the basis will predict new parameter regimes. The currently available replacement Farhat course decks explicitly distinguish projection error from model-dynamics error [50]. Their complete version-specific review is recorded separately in the source audit.

For an orthonormal column matrix V, write a full state $z=Va+w$ with $V^Tw=0$. If $\dot z=F(z)$, the exact retained coordinate satisfies $\dot a=V^TF(Va+w)$. Replacing the right side by $V^TF(Va)$ is a closure. The exact memory identity in Chapter 27 explains why eliminated coordinates can leave both memory and an initial-condition contribution.

Proposition 24.21 (What weighted POD actually minimizes)**.** *Let $X = [x_1 \ \cdots \ x_s] \in \mathbb{R}^{n\times s}$, let $W = \operatorname{diag}(w_1, \dots, w_s)$ with $w_j > 0$, and let $\lambda_1 \geq \cdots \geq \lambda_n \geq 0$ be the eigenvalues of $C = XWX^T$. Fix $1 \leq r \leq n$. Among all $V \in \mathbb{R}^{n\times r}$ with $V^TV = I$,*

$$\min_V \sum_j w_j \|x_j - VV^T x_j\|^2 = \sum_{i>r} \lambda_i.$$

A minimizing range is spanned by eigenvectors for the largest r eigenvalues.

Proof. The objective is $\operatorname{tr} C - \operatorname{tr}(V^TCV)$. In an orthonormal eigenbasis e_i of C, put $a_i = \|V^Te_i\|^2$. Then $0 \leq a_i \leq 1$, $\sum_i a_i = r$, and $\operatorname{tr}(V^TCV) = \sum_i \lambda_i a_i$. Transferring any weight from a smaller eigenvalue to an unfilled larger one cannot decrease this sum. Its maximum is therefore $\sum_{i\leq r} \lambda_i$, and the indicated eigenspace attains it. □

The proposition concerns the chosen snapshots and weights. It contains no assertion about an unsampled parameter or a differential equation. Indeed the stable matrix

$$A = \begin{pmatrix} -1 & 10 \\ 0 & -1 \end{pmatrix}, \qquad V = 2^{-1/2} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

has both eigenvalues equal to -1, while its orthogonal Galerkin reduction is $V^TAV = 4$. The reduced scalar solution grows exponentially. Nonsymmetric stable dynamics are therefore not automatically preserved by orthogonal projection.

For Petrov–Galerkin projection with full-column-rank trial and test matrices V, W, the matrix W^TV must be invertible to use $\Pi = V(W^TV)^{-1}W^T$. Direct multiplication gives $\Pi^2 = \Pi$, $\operatorname{range}\Pi = \operatorname{range} V$, and $\ker\Pi = \ker W^T = (\operatorname{range} W)^\perp$. The direction of this oblique projection is the latter kernel.

24.12.1 Hyperreduction adds another controlled approximation

State reduction need not make a nonlinear residual cheap to evaluate. Let $U \in \mathbb{R}^{n\times r}$ have orthonormal columns and let $P \in \mathbb{R}^{n\times r}$ select r coordinates. If P^TU is invertible, the discrete empirical interpolation map is

$$D = U(P^TU)^{-1}P^T.$$

Proposition 24.22 (Interpolation exactness and amplification)**.** *For every $y \in \mathbb{R}^n$, $DU = U$, $P^TDy = P^Ty$, and*

$$\|y - Dy\| \leq \left(1 + \|(P^TU)^{-1}\|\right) \|(I - UU^T)y\|.$$

Proof. The first two identities follow by cancellation of P^TU. Because D is the identity on the range of U, $(I - D)y = (I - D)(I - UU^T)y$. Use $\|I - D\| \leq 1 + \|D\|$ and the unit operator norms of U and P^T to obtain the estimate. □

The bound is deliberately elementary; sharper oblique-projection bounds are possible. Even this version shows why a small discarded snapshot norm is insufficient if the sampled basis is ill-conditioned. For oversampled least squares, full column rank of P^TU is the corresponding uniqueness and exact-reproduction requirement. An online cost independent of n additionally requires that sampled nonlinear entries and the restricted reconstruction can be evaluated without assembling all n entries. A dense nonlocal nonlinear map can violate that requirement.

Other reduction methods have different hypotheses. Classical balanced truncation concerns stable finite-dimensional linear time-invariant systems, with reachability and observability conditions,

and its usual input–output error estimate assumes zero initial state. An empirical balanced-POD approximation does not inherit every exact balancing guarantee. Moment matching requires the relevant full and reduced resolvents to exist at its interpolation nodes. Grassmann-chart interpolation requires the chart to be valid, fixed subspace rank, and appropriate alignment before interpolating reduced operators. The detailed Farhat review documents these distinctions [50].

Likewise, a weighted finite-element energy $\widetilde{U}(q) = \sum_e \xi_e U_e(q)$ with fixed weights has gradient $\sum_e \xi_e \nabla U_e(q)$. If a neural model makes $\xi_e = \xi_e(q)$, the gradient instead includes $\sum_e U_e(q) \nabla \xi_e(q)$. Omitting those terms loses the claimed potential-force relation. A preservation theorem for a fixed-weight mechanical hyperreduction must therefore be re-proved when the weights become state dependent or the equation changes to a different fluid formulation.

24.12.2 A provable constraint for a neural constitutive model

A classical or quantum neural model can be used to predict constitutive coefficients from invariant local features. For example, a model can output nonnegative scalars μ, ζ, κ in (24.17). This architecture has a limited but useful exact guarantee.

Proposition 24.23 (Nonnegative constitutive entropy production)**.** *For $T > 0$, let the viscous stress be $\tau = 2\mu S_0 + \zeta(\nabla \cdot u)I$ and the heat flux $q = -\kappa \nabla T$, with $\mu, \zeta, \kappa \geq 0$. Then*

$$\frac{\tau : \nabla u}{T} - \frac{q \cdot \nabla T}{T^2} = \frac{2\mu |S_0|^2}{T} + \frac{\zeta(\nabla \cdot u)^2}{T} + \frac{\kappa |\nabla T|^2}{T^2} \geq 0.$$

Proof. The symmetric stress is orthogonal in the Frobenius inner product to the antisymmetric part of ∇u. The traceless symmetric part S_0 is orthogonal to I. Decompose the symmetric part of ∇u into $S_0 + (\nabla \cdot u)I/3$ and substitute the heat-flux formula. □

If the scalar inputs are rotational invariants, these constitutive tensors transform correctly under rotations. Dependence on gradients and thermodynamic invariants also permits Galilean covariance. Neither property proves accurate transport coefficients. Moreover, nonnegative entropy production does not by itself make $P = pI - \tau$ positive semidefinite far from equilibrium. A kinetic moment model that requires realizable covariances can instead parameterize $P = LL^T$ or another positive-definite form and impose the remaining energy and flux constraints explicitly. These are different requirements and must be checked together.

Tensor decomposition makes these constraints transparent. In three dimensions a rank-two tensor splits into scalar trace, antisymmetric, and symmetric traceless pieces, of dimensions $1, 3, 5$. The antisymmetric piece is an exterior-square representation; the symmetric pieces arise from symmetric tensor powers and their trace decomposition. Schur functors and invariant contractions therefore help build features and outputs with known transformation laws. Determinants test nonsingularity or volume change in specified finite-dimensional maps. They do not determine a turbulent closure or imply the analytic estimates required for a continuum theorem.

For a quantum version of ToricGT, the surface or dessin specifies connectivity and admissible local operations, while circuit parameters specify a family of quantum channels and observables. A conservation law can be enforced by restricting channels or observables to a specified invariant sector. A measured expectation then estimates a well-defined feature or coefficient. State preparation, finite-shot variance, circuit approximation, and the relation between that feature and the physical kinetic state remain separate parts of the model. The quantum-fluid chapter discusses these algorithmic costs and encodings. Learning on a surface does not identify an arbitrary three-dimensional fluid with a two-dimensional toric code.

24.13 What passes from complexes to a continuum theorem

On a finite cell complex, the boundary-of-boundary identity supplies exact incidence relations. Its adjoint cochain identities, Hodge decompositions, and positive semidefinite Laplacians can enforce discrete conservation and isolate harmonic components. A gentle surface algebra or a dessin provides additional combinatorial organization of local operators. A variety of complexes organizes matrices with prescribed compositions and rank conditions. These constructions can certify the algebraic structure of a discretization or constrain a neural architecture.

To pass from such a model to a continuum claim requires a separate chain of estimates: consistency of the geometric approximation, stability uniform in mesh size, compactness in a topology suitable for nonlinear products, identification of the limiting equation, and preservation of the required initial and boundary data. Spectral gaps on individual finite graphs need not have a uniform positive lower bound as the graphs are refined. Rank constraints on finitely many matrices need not control high-frequency concentration. A learned closure with positive entropy production need not supply the required continuum regularity.

The constructive role of the project's algebra is to make some requirements exact and testable, then to identify which analytic requirements are still independent. Chapter 26 proves the finite structural statements, and Chapter 25 identifies the actual theorem and proof obligations of the recent release. The microscopic derivations in this chapter provide the physical interpretation and the correct hierarchy of approximations without changing those theorems' hypotheses.

Chapter 25

The Navier–Stokes release: exact scope and proof interfaces

Evidence status. The 166-page released manuscript has been read in full across the two mathematical review passes recorded in the source ledger. The repository revision, final theorem statements, dependency graph, and selected implementation modules have been inspected. Reading a mathematical proof, rebuilding a Lean target, and independently checking an exported kernel certificate are distinct operations; their actual outcomes are recorded separately. The propositions proved in this chapter explain mathematical interfaces of the construction. They do not replace the analytic construction imported from the release.

25.1 What the released theorem says

OpenAI's release of 8 September 2026 concerns a smooth, externally forced, three-dimensional incompressible flow. It supplies a manuscript and Lean formalization. The Clay Mathematics Institute's response of 11 September acknowledges the apparent resolution while describing an ongoing evaluation process. These facts must be distinguished from an award decision [35, 110].

Let $\nu > 0$ and define the momentum residual

$$\mathcal{R}_\nu(u,p) = \partial_t u + (u\cdot\nabla)u - \nu\Delta u + \nabla p. \tag{25.1}$$

The source's main conclusion can be recorded without changing its quantifiers:

$$\begin{aligned} &f \in C_c^\infty(\mathbb{R}^3\times(0,\infty);\mathbb{R}^3), \qquad u,p \in C^\infty(\mathbb{R}^3\times[0,1)),\\ &\mathcal{R}_\nu(u,p) = f, \quad \nabla\cdot u = 0, \quad u(\cdot,0)=0, \qquad (25.2)\\ &\operatorname{supp} u(\cdot,t)\cup\operatorname{supp} p(\cdot,t) \subset K \quad (0\le t<1),\\ &\sup_{t<1}\|u(t)\|_2 < \infty, \qquad \limsup_{t\uparrow 1}\|u(t)\|_\infty = \infty. \qquad (25.3) \end{aligned}$$

Here K is compact. With the same datum and force there is no global smooth whole-space solution with uniformly bounded kinetic energy. A corresponding periodic result is also established [109, Theorem 1.1 and Corollary 10.6].

In Fefferman's formulation, alternatives (A) and (B) impose $f = 0$, whereas (C) and (D) permit a smooth force satisfying the specified decay conditions. The released statements address (C) and (D). They do not logically negate (A) or (B): an existential assertion with $f \neq 0$ does not supply an example with $f = 0$. Pressure periodicity is part of the periodic formulation, including its erratum

[51]. The separate Euler result concerns a different equation, obtained by setting viscosity to zero, and cannot be transferred to positive viscosity by deleting the viscous estimates.

This distinction changes the appropriate goal of the present extension. A more informative presentation can expose the cancellation, support, analytic norms, and formal dependency interfaces. A stronger theorem would require a new mathematical conclusion with all its additional hypotheses discharged. The material below provides the former and identifies precisely where the source's construction supplies the latter's indispensable analytic inputs.

25.2 An exact algebra for residual correction

The same distinction between a complex and an exact complex encountered earlier reappears here. A differential identity constrains what a correction can do. It does not prove that a correction of the necessary size, support, and regularity exists.

Proposition 25.1 (Quadratic residual identity)**.** *For smooth u, w, p, π, with $\nabla \cdot w = 0$,*

$$\mathcal{R}_\nu(u + w, p + \pi) = \mathcal{R}_\nu(u, p) + L_{u,\nu}(w, \pi) + \nabla \cdot (w \otimes w),$$
$$L_{u,\nu}(w, \pi) = \partial_t w + (u \cdot \nabla)w + (w \cdot \nabla)u - \nu \Delta w + \nabla \pi.$$

The tensor convention is $(w \otimes w)_{ij} = w_i w_j$ and $(\nabla \cdot S)_i = \sum_j \partial_j S_{ij}$.

Proof. All terms other than advection are linear. Expanding the latter gives $(u + w)_j \partial_j (u_i + w_i)$, with the four products obtained by choosing one of u, w in each factor. Moreover, $\sum_j \partial_j (w_i w_j) = \sum_j w_j \partial_j w_i + w_i \sum_j \partial_j w_j = (w \cdot \nabla) w_i$. This proves every component of the identity. No averaging, limiting argument, or neglected term occurs. □

In particular, solving $L_{u,\nu}(w, \pi) = -\mathcal{R}_\nu(u, p)$ leaves the quadratic residual $\nabla \cdot (w \otimes w)$. A numerical decrease in the linear residual alone is therefore an incomplete verification. Conversely, if a background satisfies $\mathcal{R}_\nu(u, p) = -\nabla \cdot S + E$, a fluctuation with averaged covariance S can cancel the selected mean stress while creating other linear and oscillatory terms. Each of those terms needs a place in the next residual estimate.

Proposition 25.2 (A transverse wave and its covariance)**.** *Let $a, k \in \mathbb{R}^3$ with $a \cdot k = 0$ and put $w(x, \phi) = a \cos(k \cdot x + \phi)$. Then*

$$\nabla \cdot w = 0, \qquad (w \cdot \nabla) w = 0, \qquad \frac{1}{2\pi} \int_0^{2\pi} w \otimes w \, d\phi = \frac{1}{2} a \otimes a.$$

The vanishing self-advection statement requires constant a and k.

Proof. Differentiating the cosine gives $\nabla \cdot w = -(a \cdot k) \sin(k \cdot x + \phi)$ and $(w \cdot \nabla) w = -(a \cdot k) a \cos(k \cdot x + \phi) \sin(k \cdot x + \phi)$. Both vanish. Integration of $\cos^2 s = (1 + \cos 2s)/2$ proves the tensor identity. If the amplitude or phase gradient varies, the product rule introduces additional derivatives, so this proof gives no cancellation for them. □

The covariance is an element of $\mathrm{Sym}^2(\mathbb{R}^3)$, with nonnegative quadratic form $v \mapsto (a \cdot v)^2/2$. The positivity is a concrete use of a symmetric tensor power. A rank-one tensor representation does not by itself solve the vector-valued stress-divergence equation; the allowed polarization directions and the differential constraints still matter. The exterior-power test for independence of two stress directions is developed in Chapter 26.

Proposition 25.3 (Averaging and nonlinear closure)**.** *Let P be a linear averaging operator on a function algebra with $P1 = 1$. Assume $P^2 = P$ and the conditional-expectation identity $P(ag) = aP(g)$ whenever $Pa = a$. If $v = Pu$ and $w = u - v$, then*

$$P(u \otimes u) = v \otimes v + P(w \otimes w).$$

If P is positive, the final term is positive semidefinite pointwise in the averaged variables.

Proof. Idempotence gives $Pw = 0$. Expand $(v + w) \otimes (v + w)$. The two mixed terms have zero average by the assumed multiplication identity, and $P(v_i v_j) = v_i v_j$. For a constant test vector b, $b^T P(w \otimes w)b = P((b \cdot w)^2) \geq 0$. □

An orthogonal projection onto an arbitrary collection of modes need not have the multiplication property in this proposition. Likewise, averaging an extended function $F(x, Y)$ over independent Y and then substituting a phase $Y = \Phi(x)$ is different from first restricting and averaging the resulting physical function. If $F(x, Y) = \cos(2\pi(Y - x))$ on one-dimensional tori, then $\int F(x, Y)dY = 0$, whereas $F(x, x) = 1$. This elementary example prevents an unjustified replacement of a pointwise physical residual by an auxiliary mean. The phase derivatives and support geometry in a complete construction are precisely what allow these two kinds of expression to be related.

25.3 Concentration, energy, and the information lost by norms

Proposition 25.4 (Energy estimate from rest)**.** *Let u, p be a smooth solution on $[0, T]$ of $\mathcal{R}_\nu(u, p) = f$, $\nabla \cdot u = 0$, $u(0) = 0$. Assume periodicity, or sufficient decay for the integrations by parts below; fixed compact spatial support of u, p is sufficient. If $F(t) = \int_0^t \|f(s)\|_2 \, ds$, then*

$$\|u(t)\|_2^2 + 2\nu \int_0^t \|\nabla u(s)\|_2^2 \, ds \leq F(t)^2.$$

Proof. The pressure pairing is $\int u \cdot \nabla p = -\int p \nabla \cdot u = 0$. Advection contributes $\int u \cdot (u \cdot \nabla)u = \frac{1}{2} \int \nabla \cdot (u|u|^2) = 0$. Thus

$$\frac{1}{2} \frac{d}{dt} \|u\|_2^2 + \nu \|\nabla u\|_2^2 = \langle f, u \rangle.$$

Write $y_\delta = (\|u\|_2^2 + \delta^2)^{1/2}$, $\delta > 0$. Then $y_\delta' \leq \|f\|_2$, so $y_\delta(t) \leq \delta + F(t)$. Passing to $\delta \downarrow 0$ gives $\|u(t)\|_2 \leq F(t)$. Integrating the energy identity once more,

$$\|u(t)\|_2^2 + 2\nu \int_0^t \|\nabla u\|_2^2 \leq 2 \int_0^t F(s) F'(s) \, ds = F(t)^2.$$

The argument on each $T < 1$, followed by monotone convergence of the nonnegative dissipation integral, also applies up to a singular time. □

This estimate controls an integral, not the maximum of the integrand. Information reduction becomes mathematically consequential at this exact point: a small spatial set can carry a large amplitude while making a small contribution to L^2 energy.

Proposition 25.5 (Anisotropic concentration test)**.** *Let $0 < h < 1/6$, $\tau > 0$, and $W \in C_c^\infty(\mathbb{R}^3; \mathbb{R}^m)$ be nonzero. Define the comparison family*

$$V_\tau(x_1, x_2, x_3) = \tau^{-1/2-h} W\left(\frac{x_1}{\tau^{1/2}}, \frac{x_2}{\tau^{1/2}}, \frac{x_3}{\tau^{1/2-h}}\right).$$

For $1 \le p < \infty$,

$$\|V_\tau\|_p^p = \tau^{3/2-h-p(1/2+h)}\|W\|_p^p, \qquad \|V_\tau\|_\infty = \tau^{-1/2-h}\|W\|_\infty.$$

Its energy tends to zero, while its maximal amplitude diverges. Moreover, $\int_0^{\tau_0} \|\nabla V_\tau\|_2^2\, d\tau < \infty$ for every finite $\tau_0 > 0$. This comparison family is not asserted to solve Navier–Stokes.

Proof. The change of variables has determinant $\tau^{3/2-h}$. This proves the norm identities. At $p = 2$ the exponent is $1/2 - 3h > 0$. Each of the first two spatial derivatives contributes $\tau^{-1/2}$, so its squared norm has exponent $-1/2 - 3h > -1$. The axial derivative contributes $\tau^{-1/2+h}$ and has squared-norm exponent $-1/2-h > -1$. Both powers are integrable at zero. The supremum formula follows because the coordinate dilation is onto. □

For this family the critical spatial integrability exponent is

$$p_*(h) = \frac{3/2 - h}{1/2 + h} < 3.$$

At $p = 3$, $\|V_\tau\|_3^3 = \tau^{-4h}\|W\|_3^3$, which diverges. Thus the scaling test does not contradict a regularity theorem requiring a bounded critical L^3 norm. Such a theorem is a substantially stronger statement than an energy bound; the endpoint theorem for the unforced Cauchy problem is due to Escauriaza, Seregin, and Šverák [49]. Our comparison family proves only the logical insufficiency of the weaker norm information, independently of any existence construction.

25.4 Similarity coordinates as an exact change of variables

The geometry relevant to the release can be organized using the positive variables

$$\tau = 1 - t, \qquad D = \tfrac{1}{2} - h, \qquad \tau = q(1-\eta^2), \quad z = q^D\eta, \quad X = \frac{r^2}{2q}, \quad 0 < h < \tfrac{1}{2}.$$

The next calculation verifies the differential interface rather than assuming that symbolic substitution preserves all derivatives.

Proposition 25.6 (Coordinate existence and derivative weights)**.** *For $\tau > 0$ and $z \in \mathbb{R}$ there is a unique $q > |z|^{1/D}$ with $q - z^2q^{2h} = \tau$. Set $\eta = zq^{-D}$, $d = 1 - \eta^2$, and $L = 1 - 2h\eta^2$. Then $|\eta| < 1$, $L \ge 1 - 2h > 0$, and for a smooth profile $G(X, \eta)$ and any real b,*

$$\partial_t(q^bG) = q^{b-1}L^{-1}(-bG + D\eta\partial_\eta G + X\partial_X G),$$
$$\partial_z(q^bG) = q^{b-D}L^{-1}(2b\eta G + d\partial_\eta G - 2\eta X\partial_X G).$$

Proof. On $q > |z|^{1/D}$, the function $q - z^2q^{2h}$ begins at zero, tends to infinity, and has derivative $1 - 2hz^2q^{2h-1} \ge 1 - 2h > 0$. The intermediate-value theorem and strict monotonicity give existence and uniqueness. The implicit-function theorem gives smoothness, and differentiation yields

$$q_t = -L^{-1}, \quad q_z = 2\eta q^{1-D}/L, \quad \eta_t = D\eta/(qL), \quad \eta_z = d/(q^DL),$$

$$X_t = X/(qL), \qquad X_z = -2\eta X/(q^DL).$$

The two formulas follow by the product and chain rules. The endpoint $z = 0$ also satisfies them, because the implicit derivative remains nonzero. □

On bounded profile ranges, a transverse derivative has weight $q^{-1/2}$, an axial derivative has weight q^{-D}, and a time derivative has weight q^{-1}. The denominators remain bounded because $L \geq 1 - 2h$. Constants can depend on the derivative order and fixed profile parameters. They cannot be allowed to hide an extra loss growing with the approximation stage if one later intends to sum an infinite correction sequence.

The apparent cylindrical singularity at $r = 0$ must also be resolved. For instance, if the azimuthal velocity is $r\,a(r^2, z, t)$, then

$$r\,a(r^2, z, t)e_\theta = a(x_1^2 + x_2^2, z, t)(-x_2, x_1, 0)$$

is smooth in Cartesian coordinates. Boundedness of a cylindrical component alone would not establish this. This is a local regularity condition on a representation, comparable to requiring that coordinate expressions of a section glue smoothly; it is not supplied by the topology of the domain.

25.5 A quadratic relation that controls analytic derivative loss

An especially useful connection to the project's algebraic methods occurs in the analytic linear recursion. The relevant feature is a forbidden composition of derivative-carrying blocks. We isolate its mechanism in a statement independent of fluid terminology.

Lemma 25.7 (Words with no adjacent derivative letters)**.** *A word of length k in the letters N, D with no consecutive D has at most $\lceil k/2 \rceil$ occurrences of D. The number with exactly j occurrences is $\binom{k-j+1}{j}$, interpreted as zero when $j > k - j + 1$. The total number a_k satisfies $a_0 = 1$, $a_1 = 2$, and $a_k = a_{k-1} + a_{k-2}$ for $k \geq 2$.*

Proof. Between j derivative letters there are at least $j - 1$ other letters, so $2j - 1 \leq k$. If their positions are $1 \leq i_1 < \cdots < i_j \leq k$ with $i_{r+1} \geq i_r + 2$, the shifted positions $i_r - (r - 1)$ form an ordinary j-element subset of $\{1, \ldots, k - j + 1\}$. This is a bijection, proving the binomial formula. For the recurrence, an admissible word ending in N is an arbitrary admissible word of length $k - 1$ followed by N. One ending in D, when $k \geq 2$, ends in ND and corresponds to a word of length $k - 2$. These two classes are disjoint and exhaustive. □

Proposition 25.8 (Sparse Volterra systems)**.** *Let $E = E_0 \oplus E_1$ be a finite-dimensional normed complex vector space. Fix a compact real interval I and let $S_\rho = \{z \in \mathbb{C} : \operatorname{dist}(z, I) < \rho\}$. Consider an integral equation on $0 \leq s \leq a$ and this complex parameter neighborhood,*

$$W = Gf + G(A_0 W + A_1 \partial_\eta W), \qquad (Gg)(s) = \int_0^s K(s, r)g(r)\,dr.$$

Assume $\|K(s, r)\| \leq 1$, K preserves the two summands and is independent of η; the coefficients and kernel are continuous in their radial variables; A_0, A_1, f are uniformly bounded and holomorphic in η; and every value and every η-derivative of A_1 maps E_0 into E_1 and annihilates E_1. On each smaller neighborhood $S_{\rho'}$ with $0 < \rho' < \rho$, the Picard series converges for every finite a. It gives the unique solution in the class of solutions bounded and holomorphic on a parameter neighborhood strictly larger than the one on which they are compared.

Proof. Expand the k-fold iterate of $G(A_0 + A_1\partial_\eta)$ into words. The kernel between adjacent coefficients preserves E_0, E_1. Therefore

$$A_1(s)K(s, r)A_1(r) = 0, \qquad A_1(s)K(s, r)\partial_\eta A_1(r) = 0.$$

The second identity is necessary: in a composition of two derivative operators the first derivative can strike the second coefficient. It vanishes for the same block reason. Thus a nonzero word cannot contain two adjacent derivative operators, even after expanding the product rule. By Lemma 25.7, at most $p_k = \lceil k/2 \rceil$ Cauchy radius losses are required.

Put $\Delta = \rho - \rho' > 0$. Distribute that loss in p_k equal parts. The Cauchy estimate for one parameter derivative costs at most p_k/Δ; when $p_k = 0$ use cost one. Operator norms absorb the coefficient bounds, and the nested integration domain has volume $a^{k+1}/(k+1)!$. Bounding the number of words by 2^k gives, after enlarging a constant C,

$$\left\| [G(A_0 + A_1\partial_\eta)]^k Gf \right\|_{S_{\rho'}} \le \frac{C^{k+1} a^{k+1}}{(k+1)!} \max\{1, p_k/\Delta\}^{p_k}.$$

Stirling's bound shows that the kth root of the right-hand side is $O(k^{-1/2})$, tending to zero. Hence the series converges uniformly on the smaller neighborhood. Applying the same estimate after one additional radius restriction permits differentiation in η and termwise substitution into the integral equation. If $W_1 - W_2 = G(A_0 + A_1\partial_\eta)(W_1 - W_2)$, iteration and the same estimate make the difference zero on the smaller neighborhood. Analytic continuation gives the stated uniqueness. □

The source's six-component linear recursion has the splitting $E_0 = \mathbb{C}^4$, $E_1 = \mathbb{C}^2$, with diagonal radial kernels. Its use of the above relation is traceable to Lemma 5.1, rather than inferred from a superficial resemblance to a quiver [109]. The relation has three precise interpretations. It eliminates forbidden operator words; it limits derivative loss; and, with the integration simplex, it makes a convergent analytic solve possible on a common radial interval.

The language of admissible words is regular: a finite automaton needs only to remember whether the preceding letter was D, plus a rejecting state. The analytic conclusion nevertheless requires unbounded iteration and quantitative estimates. A regular syntax does not make the represented analytic task finite-dimensional.

Nor does this calculation establish a gentle algebra. The abstract one-vertex algebra $k\langle N, D\rangle/(D^2)$ is infinite-dimensional, and the arrow N has two allowed successors. It therefore fails a usual gentle-algebra continuation condition. The honest bridge is to a representation with a quadratic zero relation and a controlled path language. Any claim of a particular gentle presentation would require a separate equivalence or construction.

25.6 From formal orders to actual smooth fields

The order of a term in a formal expansion must not be confused with a convergence estimate for the coefficients. Suppose coefficients of order j have constants growing faster than every exponential. Positive powers of a small parameter do not automatically make their unrestricted sum converge. A suitable realization uses cutoffs chosen after the coefficients are known.

Lemma 25.9 (A diagonal realization estimate)**.** *Let $0 < q < q_0 \le 1$ on an open domain, with q smooth and bounded below on every compact subset. Suppose smooth increments Z_j, $j \ge 1$, satisfy*

$$|Z_j|_m \le C_{j,m} q^{g_j - \ell_m} (1 + |\log q|)^{P_{j,m}}, \qquad 0 < g_1 \le g_2 \le \cdots \longrightarrow \infty,$$

where $|\cdot|_m$ is the pointwise maximum of coordinate derivatives through order m, and ℓ_m is independent of j. Assume that differentiating a cutoff $\chi(aq)$ costs a power of q^{-1} depending only on derivative

order, uniformly in a on its transition region. Then one can choose $a_{j+1} \geq 2a_j$ so that the cutoff sum $Z = \sum_{j\geq 1} \chi(a_j q) Z_j$ is locally finite and smooth, and its tail after a fixed initial block satisfies

$$\left| Z - \sum_{j=1}^{J} Z_j \right|_m \leq 2^{-J} q^{g_{J+1}/2 - L_m} \quad \left(J \geq \max\{1, m\},\ q < \frac{1}{2a_J} \right).$$

Here $\chi = 1$ on $[0, 1/2]$, $\chi = 0$ on $[1, \infty)$, and L_m is independent of J.

Proof. The product rule and the cutoff hypothesis bound the mth seminorm of the jth summand by $\widehat{C}_{j,m} q^{g_j - L_m} (1 + |\log q|)^{\widehat{P}_{j,m}}$. At stage j impose this finite list of requirements for $0 \leq m \leq j$:

$$\widehat{C}_{j,m} q^{g_j/2} (1 + |\log q|)^{\widehat{P}_{j,m}} \leq 2^{-j} \quad (0 < q \leq a_j^{-1}).$$

It is possible because every positive power of q dominates every fixed power of $|\log q|$ at zero. Increase a_j further to preserve its geometric growth. On a compact subset where $q \geq \delta > 0$, terms with $a_j > \delta^{-1}$ vanish, proving local finiteness and smoothness. If $q < (2a_J)^{-1}$, the first J cutoffs equal one. For $j > J \geq m$, the bound just imposed and $g_j \geq g_{J+1}$ give

$$\sum_{j>J} 2^{-j} q^{g_j/2 - L_m} \leq 2^{-J} q^{g_{J+1}/2 - L_m},$$

as required. □

To preserve a differential constraint, apply the cutoffs to suitable potentials. In three-dimensional Euclidean space,

$$\nabla \times (\chi(a_j q) A_j) = \chi(a_j q) \nabla \times A_j + a_j \chi'(a_j q) \nabla q \times A_j.$$

The second term is a real part of the field and must be retained. Taking divergence of the full curl yields zero. Cutting off an already divergence-free velocity directly would instead give $\nabla \cdot (\chi u) = \nabla \chi \cdot u$, which need not vanish.

Proposition 25.10 (All-order residual realization)**.** *Let $U^{[J]}$ be finite approximations and U their cutoff realization. Assume their derivatives have polynomial growth in q^{-1} with exponent losses independent of J, and that the tail estimate of Lemma 25.9 applies through every fixed order. Let F be a fixed differential polynomial of finite order and degree whose coefficient derivatives have polynomial growth in q^{-1}. If*

$$|F(U^{[J]})|_m \leq C_{J,m} q^{\rho_J - K_m} (1 + |\log q|)^{P_{J,m}} + E_{J,m}, \quad \rho_J \to \infty,$$

with K_m independent of J and each fixed $E_{J,m} = O(q^N)$ for every N, then $|F(U)|_m = O(q^N)$ for every m, N.

Proof. Fix m, N. Let s, d bound the differential order and polynomial degree of F. For products use the exact identity

$$\prod_{i=1}^{d} a_i - \prod_{i=1}^{d} b_i = \sum_{r=1}^{d} \left(\prod_{i<r} a_i \right) (a_r - b_r) \left(\prod_{i>r} b_i \right).$$

The product rule then bounds $F(U) - F(U^{[J]})$ in seminorm m by a fixed power of q^{-1} times the tail seminorm $m + s$ and a polynomial in the two field seminorms $m + s$. All exponent losses depend on m, s, d, not J. Choose J large enough that both ρ_J and $g_{J+1}/2$ exceed N plus these losses and one extra unit. Keep that J fixed. Shrink the q-neighborhood so all its finitely many constants and logarithms are absorbed by that extra unit, and the first J cutoffs equal one. The finite residual, its fixed flat remainder, and the residual difference are then all $O(q^N)$. □

The order of these quantifiers is essential:

$$\forall(m,N)\ \exists J\ \exists\delta,C\ \forall q\in(0,\delta):\ |F(U)|_m\le Cq^N.$$

It does not assert one J that controls every derivative. It does not require summing the flat remainders $E_{J,m}$ over J. It also does not assert convergence of the uncut formal series. This distinction is the analytic counterpart of keeping a filtered algebra, its associated graded object, and a completed realization separate.

25.7 The finite correction cycle as a verifiable dependency graph

The source's construction provides the hypotheses needed above through separate analytic modules. The following locator table records where to check them; the table is a dependency guide, not a substitute proof [108, 109].

Interface	Mathematical obligation	Source location
Background	Smooth axis; exact exterior; supported compatible stress; positive cone margins	Theorem 4.6; Appendices A–C
Order recursion	A common radial interval, moment repair, derivative losses independent of formal order	Lemmas 5.1–5.4; Proposition 5.5
Phase geometry	Support separation and compatible finite covering coordinates	Section 6
Oscillatory solve	Transverse evolution, damping, real covariance, full curl correction	Propositions 7.2, 7.5–7.6; Lemma 7.7
Mean solve	Compact primitives, auxiliary inverse, five moment equations	Section 8
Iteration	Full residual recomputation, fixed-domain gains, all physical derivatives	Propositions 9.3–9.9
Completion	Smooth forcing, energy, comparison, all viscosities, periodization	Section 10

One can encode a correction state as a typed record rather than a loose list of estimated errors. Its fields should include a velocity, pressure, potential representation, support certificates, exact preserved moments, residual decomposition, and derivative bounds. A correction operation must produce a new such record. The pressure and residual are recomputed from the updated velocity, so one cannot reuse a previous covariance after altering a wave amplitude. This record-based reading is directly compatible with the typed graph approach developed earlier in the monograph.

The numerical exponents also permit an elementary exact check. Put $\kappa=10^{-5}$, $\sigma_j=1/5+j/10$, $B_j=1/2+\sigma_j$, and $C_j=1+\sigma_j$. All asserted gains of the source's cycle must exceed $1/10$ after derivative losses. For example,

$$\frac{18}{100}-2\kappa>\frac{17}{100}>\frac{1}{10},\qquad \sigma_j-3\kappa\ge\frac{1}{5}-3\kappa>\frac{17}{100},$$
$$\frac{1}{2}-2\kappa>\frac{1}{10},\qquad \frac{9}{10}-4\kappa>\frac{1}{10}.$$

These rational inequalities certify arithmetic slack once the functional estimates are supplied. They do not establish those functional estimates. This is a useful separation for learned proof search: a

model may propose an ordering or an exponent assignment, while exact arithmetic and analytic lemmas discharge different obligations.

25.8 Localization and smooth force extension

Proposition 25.11 (Localization without a divergence defect)**.** *Suppose A, Be_θ and p are smooth in Cartesian coordinates, put $u = \nabla \times A + Be_\theta$, and assume B is independent of θ. Let c be a smooth axisymmetric scalar. Set*

$$\widetilde{u} = \nabla \times (cA) + cBe_\theta, \qquad \widetilde{p} = cp, \qquad a = \nabla c \times A.$$

Then $\widetilde{u} = cu + a$ and $\nabla \cdot \widetilde{u} = 0$. Its residual is exactly

$$\begin{aligned}\mathcal{R}_\nu(\widetilde{u}, cp) = c\mathcal{R}_\nu(u,p) + (\partial_t c - \nu\Delta c)u - 2\nu\sum_j (\partial_j c)\partial_j u + p\nabla c \\ + (c^2 - c)(u\cdot\nabla)u + c(u\cdot\nabla c)u \\ + \partial_t a - \nu\Delta a + c(u\cdot\nabla)a + (a\cdot\nabla)(cu) + (a\cdot\nabla)a.\end{aligned}$$

Proof. The curl product rule gives $\widetilde{u} = cu + a$. Divergence of a curl is zero; the other term has divergence $r^{-1}\partial_\theta(cB) = 0$ away from the axis and hence everywhere by Cartesian smoothness. Expand $\partial_t(cu)$, $\Delta(cu)$, and $\nabla(cp)$ by the product rule. Finally, $(cu\cdot\nabla)(cu) = c^2(u\cdot\nabla)u + c(u\cdot\nabla c)u$. Adding the three advection terms containing a yields the formula. □

The displayed expression explains why a spatial cutoff requires endpoint regularity of the actual potentials as well as the velocity. On a region where $c = 1$ the extra terms vanish. On its transition region every term involving the uncut fields must remain smooth through the terminal time. A proof that estimates only $c\mathcal{R}_\nu(u,p)$ has omitted real terms.

Lemma 25.12 (Realization of terminal force jets)**.** *Let $f(x,t)$ be smooth for $t < T$, supported in a fixed compact set $K \subset \mathbb{R}^d$. Suppose every $\partial_x^\alpha \partial_t^j f$ has a uniform limit $\partial_x^\alpha F_j(x)$ as $t \uparrow T$, where $F_j \in C_c^\infty(\mathbb{R}^d)$ are supported in K. Then f has a smooth extension to $t \geq T$ supported in $K \times [T, T+1]$ on that side. If it already vanishes near an initial time, the combined extension retains that property.*

Proof. Choose a smooth cutoff χ, one near zero and zero for arguments at least one, and for $s = t - T \geq 0$ define

$$f_+(x, T+s) = \sum_{j=0}^{\infty} \chi(b_j s)\frac{s^j}{j!}F_j(x).$$

For fixed α, m, the product rule and $s \leq b_j^{-1}$ on the summand's support give a bound

$$C_{j,m} b_j^{m-j} \|F_j\|_{C^{|\alpha|}}.$$

Fix $b_0 = 1$. For $j \geq 1$, choose increasing $b_j \geq 1$ so that this is at most 2^{-j} whenever $|\alpha| + m \leq \lfloor j/2 \rfloor$. Each step imposes finitely many constraints with $m - j < 0$, so the choice exists. Every differentiated tail converges uniformly. At $s = 0$, the mth time derivative of a fixed summand is F_m when $j = m$ and zero otherwise. Uniform differentiated convergence gives the desired jets.

Uniform convergence of the original mixed derivatives makes their limits compatible: pass to the limit in the fundamental theorem of calculus along a coordinate segment, and then along a time segment. The two pieces therefore glue smoothly at T. All summands retain support in K, and vanish for $s \geq 1$; uniform convergence of every derivative also proves smoothness at that support boundary. □

Flatness at a single singular point is only part of this lemma's hypotheses. One must additionally control all terminal derivatives away from that point, including the cutoff transition region. Conversely, vanishing of every jet does not force a smooth function to vanish near the point: the function e^{-1/s^2} for $s > 0$, extended by zero for $s \le 0$, is the standard counterexample. Substituting "analytic" for "smooth" would change the problem and invalidate this extension method.

25.9 Comparison, quantitative residual checks, and proof strength

An explicit singular solution rules out a global competitor only after uniqueness has been established in the competitor's class. The whole-space class in the Clay problem permits no arbitrary decay assumption on the pressure. That is why a periodic energy calculation is not by itself a complete proof of alternative (C). The source's Lemma 10.5 handles this point using a pressure-gradient argument and expanding spatial cutoffs; its hypotheses must be preserved in any proof adapter [109].

The following separate estimate is useful for the project's computable fluid examples and for any learned approximate field.

Proposition 25.13 (Residual-to-solution estimate on the flat torus)**.** *Let u, v be smooth divergence-free periodic vector fields on $[0,T]$, and let periodic p, q satisfy*

$$\mathcal{R}_\nu(u,p) = f, \qquad \mathcal{R}_\nu(v,q) = f + r.$$

Write $w = v - u$ and $K(t) = \|\nabla u(t)\|_{L^\infty,\mathrm{op}}$. If K is integrable, then

$$\|w(t)\|_2 \le e^{\int_0^t K(a)\,da}\|w(0)\|_2 + \int_0^t e^{\int_s^t K(a)\,da}\|r(s)\|_2\,ds.$$

In particular, equal data and $r = 0$ imply $v = u$.

Proof. Subtract the equations in the form

$$\partial_t w + (v\cdot\nabla)w + (w\cdot\nabla)u - \nu\Delta w + \nabla(q-p) = r.$$

Periodic integration and zero divergence remove the transport and pressure pairings. Thus

$$\tfrac{1}{2}(\|w\|_2^2)' + \nu\|\nabla w\|_2^2 \le K\|w\|_2^2 + \|r\|_2\|w\|_2.$$

With $y_\delta = (\|w\|_2^2 + \delta^2)^{1/2}$ one has $y_\delta' \le K y_\delta + \|r\|_2$. Multiplication by $e^{-\int_0^t K}$ and integration prove the bound after $\delta \downarrow 0$. □

This inequality gives a precise role for a quantum or classical learner: it may produce v, but a certificate needs its divergence defect, momentum residual, initial error, and an appropriate stability coefficient. The estimate does not remain effective near blowup merely because the sampled residual is small. A pointwise grid evaluation does not bound the continuum L^2 residual without an interpolation or enclosure argument, and $\int K$ can become large. Numerical fitting and a posteriori proof are different stages of the same workflow.

The same issue explains why nearby classical results are not interchangeable. Nonuniqueness in a weak class does not give blowup of a smooth solution; a modified averaged nonlinearity is not the original Navier–Stokes nonlinearity; and rough forcing is not smooth forcing. These distinctions are visible respectively in the settings of Buckmaster–Vicol, Tao, and Albritton–Brué–Colombo [3, 28, 163].

25.10 Viscosity, periodization, and preservation of quantifiers

Proposition 25.14 (Viscosity conversion)**.** *If $\mathcal{R}_1(u,p) = f$ and $\nabla \cdot u = 0$, then for $\nu > 0$*

$$u_\nu(x,t) = \sqrt{\nu}\, u(x/\sqrt{\nu}, t), \quad p_\nu(x,t) = \nu p(x/\sqrt{\nu}, t), \quad f_\nu(x,t) = \sqrt{\nu}\, f(x/\sqrt{\nu}, t)$$

satisfy $\mathcal{R}_\nu(u_\nu, p_\nu) = f_\nu$ and zero divergence. The singular time is unchanged. In three dimensions,

$$\|u_\nu(t)\|_2^2 = \nu^{5/2}\|u(t)\|_2^2, \qquad \nu\|\nabla u_\nu(t)\|_2^2 = \nu^{5/2}\|\nabla u(t)\|_2^2.$$

Proof. Put $y = x/\sqrt{\nu}$. Each term of the momentum equation acquires a factor $\sqrt{\nu}$: two spatial derivatives contribute ν^{-1}, one contributes $\nu^{-1/2}$, and the amplitudes contribute the stated powers. The divergence equals $(\nabla_y \cdot u)(y,t)$. The Jacobian of $x = \sqrt{\nu}y$ is $\nu^{3/2}$, while the velocity square contributes ν and the gradient square contributes one. This proves both norm identities. The time argument was not changed. □

Proposition 25.15 (Disjoint-support periodization)**.** *Suppose smooth whole-space fields u, p, f solve Navier–Stokes with viscosity ν and have common spatial support in a compact set $K \Subset (-1/2, 1/2)^3$. Then their sums over integer translates define smooth periodic fields solving the same equation, with the same viscosity.*

Proof. The sums are locally finite. The distinct translated compact supports have positive separation, so at each point a neighborhood intersects at most one support, unless the fields vanish there. Differential operators commute with the sums. In the advection term every mixed product of different translates vanishes, leaving the sum of the individual advection terms. The same argument gives zero divergence and retains periodicity of both velocity and pressure. □

For a larger compact support one first uses the parabolic scaling

$$\widetilde{u}(x,t) = \lambda u(\lambda x, \lambda^2(t - t_0)), \quad \widetilde{p}(x,t) = \lambda^2 p(\lambda x, \lambda^2(t - t_0)), \quad \widetilde{f}(x,t) = \lambda^3 f(\lambda x, \lambda^2(t - t_0)).$$

Choosing λ large puts the support inside the cube. If the original fields vanish on an initial interval, they extend smoothly by zero before t_0. For original singular time one, $t_0 = 1 - \lambda^{-2}$ keeps the new singular time equal to one. These conditions explain why periodizing an arbitrary spatially decaying field would not suffice: overlapping copies would create additional nonlinear interactions.

25.11 Formal proof provenance and what this extension adds

The pinned upstream revision is

`f9e8bc5b38b6e212696e8a30e3e91517af887bbd`.

Its toolchain is Lean 4.34.0-rc2, with mathlib revision

`85e3a25e006c35636f0e53b0e9296caca2685bc0`.

The project's earlier selected Lean declarations use a different toolchain. They cannot be merged by silently substituting one compiler for the other. An isolated verification script preserves the upstream pins.

The final upstream declarations are

```
NavierStokes.Comparator.navier_stokes_breakdown_R3
NavierStokes.Comparator.navier_stokes_breakdown_periodic.
```

Their types quantify over every positive viscosity and retain the exact initial-data, forcing, regularity, energy, and periodicity predicates. The reference challenge modules intentionally contain unproved challenge statements. They are kept separate from the solution imports. A source scan that merely counts every occurrence of `sorry` in the whole repository would therefore misdescribe the solution's dependency boundary. Conversely, a source scan finding no such token cannot replace kernel checking or comparison of the definitions with the intended mathematics.

The inspection found 816 files under the Navier–Stokes source directory and 609 local modules in the final solution's import closure. Selected modules were read in detail; all 609 were not manually reviewed. The upstream metadata reports only the standard dependencies `propext`, `Classical.choice`, and `Quot.sound` for the final results. The accompanying verification ledger distinguishes that upstream report from any fresh build and independent Comparator execution performed in this extension [108].

In this review the exact compiler distribution was downloaded and its published SHA-256 checksum matched. Dependency cache retrieval then failed with HTTP 429 responses after partial decompression. A direct attempt to elaborate the final module stopped at missing local import artifacts. Consequently this pass did not complete a fresh build of the final upstream theorems, and did not run the independent Comparator certificate check. The source audit found no challenge import and no unproved-placeholder token in the 609-module local solution closure, after removing comments and string literals. That lexical result has the narrower meaning stated here; it is not a substitute for the uncompleted kernel verification.

The extension's new contribution is an explicit mathematical organization: quadratic residual calculus; exact information-loss identities; norm scaling that separates energy from regularity; a forbidden-word estimate that explains a common analytic domain; all-order realization with correct quantifiers; potential-based localization; and a quantitative residual comparison for learned fields. Chapter 26 then develops the arithmetic, exterior-power, and filtered-algebra connections, while Chapter 27 treats quantum representations and learning.

These are rigorous bridges because each has a specified map or estimate. They do not identify a finite surface-code complex with the continuum Navier–Stokes equation, or turn a learned function into an existence theorem without a stability certificate. They also preserve the released theorem's force. Removing that force, imposing a different physical constitutive law, or deriving the same singular flow from a microscopic quantum state would be a further theorem with additional analytic content.

Chapter 26

Orders, complexes and the algebra of a fluid singularity

Evidence status. This chapter gives original, explicit proofs of the finite algebraic and differential identities used to interpret the released construction. The construction of the singular Navier–Stokes solution remains the external theorem of OpenAI [109], reviewed in Chapter 25. The real auxiliary torus, the physical three-dimensional fluid, a Riemann surface and a quantum register are different objects. Every identification below specifies which structure it preserves.

26.1 What a useful reformulation must preserve

The central opportunity is to reorganize the proof into mathematical interfaces that the inherited project already studies: integral endomorphisms, exterior powers, complexes, quadratic tensor maps, moment determinants and filtered operator algebras. This is more informative than attaching algebraic terminology to a picture of a vortex. It also creates obligations. A change of notation must preserve the domain, positive viscosity, zero initial data, compact smooth forcing, bounded kinetic energy and the claimed divergence of velocity. An algebraic identity must not silently replace an analytic estimate that is uniform in the concentrating scale.

In the released proof the physical equation is on $\mathbb{R}^3$, with a subsequent periodic construction on $\mathbb{T}^3$. The auxiliary $\mathbb{T}^2$ carries oscillatory labels and is evaluated along a prescribed phase map. It is neither the physical domain nor a quantum phase space. A surface-algebra model can encode incidence, permitted paths and local tensor types. To obtain an actual fluid discretization one must additionally choose a metric, a volume form, boundary conditions, constitutive parameters and approximation operators. Those data will remain visible throughout this chapter.

26.2 The de Rham complex behind incompressible motion

Let M be an oriented flat three-dimensional manifold without boundary, or use fields on $\mathbb{R}^3$ for which the integrations below have no boundary terms. Write $\alpha = u^\flat$ for the velocity one-form, δ for the metric codifferential and $\Delta_H = d\delta + \delta d$ for the nonnegative Hodge Laplacian. With the usual Euclidean sign convention, Δ_H acts componentwise as minus the coordinate Laplacian. Thus incompressibility is $\delta\alpha = 0$.

Proposition 26.1 (The differential-form momentum equation)**.** *For a smooth divergence-free velocity, the usual forced Navier–Stokes equation is equivalent to*

$$\partial_t \alpha + \iota_u d\alpha + \nu \Delta_H \alpha + d\pi = f^\flat, \qquad \delta\alpha = 0, \qquad \pi = p + \tfrac{1}{2}|u|^2.$$

Its vorticity two-form $\omega = d\alpha$ satisfies

$$\partial_t \omega + \mathcal{L}_u \omega + \nu \Delta_H \omega = df^\flat.$$

Proof. The ith component of $\iota_u d\alpha$ is $\sum_j u_j(\partial_j u_i - \partial_i u_j) = (u \cdot \nabla)u_i - \partial_i(|u|^2/2)$. Adding $d\pi$ restores the advective and pressure terms. The sign of Δ_H gives $-\nu\Delta u$ on the left. Apply d to this equation. It commutes with ∂_t and with $d\delta + \delta d$ because $d^2 = 0$. Cartan's identity gives $d\iota_u\omega = \mathcal{L}_u\omega - \iota_u d\omega = \mathcal{L}_u\omega$. The pressure disappears by $d^2\pi = 0$. □

This places exterior powers in their proper role. The momentum flux $u \otimes u$ is symmetric; vorticity is alternating. The two occupy Sym^2 and $\bigwedge^2$ tensor types, respectively. They interact through differential operators and metric contractions. Exchanging their symmetry types would change the equation.

Proposition 26.2 (Energy and the work of a differential)**.** *Under the hypotheses just stated,*

$$\frac{d}{dt}\frac{1}{2}\|u\|_2^2 + \nu\|\nabla u\|_2^2 = \langle f, u\rangle_{L^2}.$$

In particular neither pressure nor transport contributes to the total kinetic-energy derivative.

Proof. The pressure contribution is $\langle dp, \alpha\rangle = \langle p, \delta\alpha\rangle = 0$. Pointwise, $\langle \iota_u d\alpha, \alpha\rangle = d\alpha(u, u) = 0$, and $\langle d(|u|^2/2), \alpha\rangle$ integrates to zero. Integration by parts gives $\langle -\Delta u, u\rangle = \|\nabla u\|_2^2$. Finally $\langle \partial_t u, u\rangle = \frac{1}{2}\partial_t\|u\|_2^2$. □

The proposition controls an integral. It does not bound $\|u\|_\infty$. This matters especially in the newly released construction: concentration makes the energetic cost of a large speed small. The scale computation in Chapter 25 is therefore compatible with this identity, rather than a violation of it. Nor does dissipation imply that every individual mode decreases: transport can redistribute energy between scales while its total work vanishes.

On a curved manifold one must specify which viscous operator is intended. The Hodge Laplacian, the rough Laplacian and the divergence of a symmetric strain tensor differ by curvature contributions under standard conventions. A topological surface model cannot choose between them. The flat torus examples below fix the Hodge convention explicitly. Their familiar three-dimensional torus drawing is a visualization of periodic coordinates, not an assertion that the induced metric of the embedded doughnut is flat.

26.3 A quadratic order occurs in the actual construction

Section 6 of OpenAI [109] chooses the integer matrix

$$J = \begin{pmatrix} 3 & 1 \\ 1 & 5 \end{pmatrix}, \qquad A = J - 4I = \begin{pmatrix} -1 & 1 \\ 1 & 1 \end{pmatrix}.$$

Its arithmetic controls two different derivative scales on an auxiliary torus. This is a direct meeting point with the orders developed earlier in the monograph.

Proposition 26.3 (The integral quadratic-order representation)**.** *One has* $A^2 = 2I$*,* $J^2 - 8J + 14I = 0$ *and* $\det J = 14$*. The map*

$$\rho : \mathbb{Z}[\sqrt{2}] \longrightarrow M_2(\mathbb{Z}), \qquad a + b\sqrt{2} \longmapsto aI + bA$$

is an injective unital ring homomorphism, and

$$\det \rho(a + b\sqrt{2}) = a^2 - 2b^2.$$

Consequently $J = \rho(4 + \sqrt{2})$ *and* $\det(J^k) = 14^k$ *for every integer* $k \geq 0$*.*

Proof. Matrix multiplication gives $A^2 = 2I$. Hence $(aI + bA)(cI + dA) = (ac + 2bd)I + (ad + bc)A$, the multiplication law of the quadratic order. Its matrix is $\left(\begin{smallmatrix} a-b & b \\ b & a+b \end{smallmatrix}\right)$, whose determinant is $a^2 - 2b^2$. If it is the zero matrix its off-diagonal entry gives $b = 0$, then $a = 0$, proving injectivity. Expanding $(J - 4I)^2 = 2I$ gives the polynomial identity. Determinants are multiplicative, including the case $k = 0$. □

The determinant is also the scalar by which the top exterior power acts: $\bigwedge^2 J = 14 \operatorname{id}_{\bigwedge^2 \mathbb{Z}^2}$. Thus the same alternating-tensor construction that gave the determinant and Cauchy–Binet earlier now measures a covering multiplicity. Since 14 is not a unit of $\mathbb{Z}$, this is a covering endomorphism, not an invertible integral change of basis.

Proposition 26.4 (Covering multiplicity and exact averaging)**.** *The map* $\pi_k : \mathbb{R}^2/\mathbb{Z}^2 \to \mathbb{R}^2/\mathbb{Z}^2$ *given by* $[y] \mapsto [J^k y]$ *is a surjective covering with* 14^k *points in each fiber. For normalized Haar measure,*

$$\int_{\mathbb{T}^2} F(J^k y)\, dy = \int_{\mathbb{T}^2} F(y)\, dy$$

for every continuous F*.*

Proof. Invertibility over $\mathbb{R}$ gives surjectivity. A fiber is a translate of $(J^{-k}\mathbb{Z}^2)/\mathbb{Z}^2$, isomorphic to $\mathbb{Z}^2/J^k\mathbb{Z}^2$. Integer row and column operations reduce the latter quotient to Smith normal form, whose order is $|\det J^k| = 14^k$. Locally the real linear map is a diffeomorphism, proving the covering assertion. For a character $e_n(y) = e^{2\pi i n \cdot y}$, pullback gives $e_{(J^k)^T n}$. Because J^k is invertible, this character has nonzero integral exactly when $n = 0$. Finite linear combinations of characters satisfy the integral identity. Uniform approximation by trigonometric polynomials extends it to continuous functions. □

This argument does not invoke a long-time ergodic limit. It is an exact identity for a finite covering. In the released proof one uses a common auxiliary torus for overlapping dyadic bands, keeping the difference of covering levels bounded. Its derivative factors are then uniformly bounded even though the absolute covering level grows. Omitting this distinction would lose the uniformity needed later in the estimates.

Proposition 26.5 (Why this is not a complex elliptic-curve endomorphism)**.** *There is no real linear complex structure* C *on* $\mathbb{R}^2$*, meaning* $C^2 = -I$*, that commutes with* J*.*

Proof. The real symmetric matrix J has distinct real eigenvalues $4 - \sqrt{2}$ and $4 + \sqrt{2}$. Every commuting map preserves each of its one-dimensional real eigenspaces. On each, C is multiplication by a real scalar c, so C^2 is multiplication by $c^2 \geq 0$ and cannot be -1. □

Therefore the auxiliary covering has an arithmetic description by a quadratic order, but cannot be advertised as a holomorphic endomorphism of one fixed complex elliptic curve. The fact that its underlying real manifold is a torus supplies too little structure for that conclusion. This distinction also prevents an unsupported identification with the elliptic-curve groups used by cryptography.

26.4 The number-field norm supplies an analytic estimate

Set $b = \sqrt{2} - 1$, $v_r = (1, -b)$ and $v_t = (b, 1)$. Direct calculation gives $Jv_r = (4 - \sqrt{2})v_r$ and $Jv_t = (4 + \sqrt{2})v_t$. Their irrational directions are useful because no nonzero integer frequency is perpendicular to either one, and the failure of perpendicularity has a quantitative lower bound.

Lemma 26.6 (A quadratic small-divisor bound)**.** *There is an explicit constant $c = (1 + (1 + \sqrt{2})^2)^{-1/2} > 0$ such that*

$$|v_r \cdot n| \geq c/|n|_2, \qquad |v_t \cdot n| \geq c/|n|_2 \quad (n \in \mathbb{Z}^2 \setminus \{0\}).$$

Proof. For the first expression put $a = n_1 + n_2$ and $d = -n_2$. Then $v_r \cdot n = a + d\sqrt{2}$. Its algebraic norm $a^2 - 2d^2$ is a nonzero integer: otherwise irrationality of $\sqrt{2}$ would force $a = d = 0$, and hence $n = 0$. Its absolute value is at least one. The conjugate is $n_1 + (1 + \sqrt{2})n_2$, of absolute value at most $c^{-1}|n|_2$ by Cauchy–Schwarz. Dividing the norm by this bound proves the first assertion. The other direction is $(n_2 - n_1) + n_1\sqrt{2}$, whose conjugate has the same coefficient norm. □

Proposition 26.7 (A tame inverse of a torus directional derivative)**.** *Let v be either direction above, and let F be a smooth zero-mean function on $\mathbb{T}^2$. There is a unique smooth zero-mean ϕ satisfying $v \cdot \nabla\phi = F$, namely*

$$\phi(y) = \sum_{n \neq 0} \frac{\widehat{F}(n)}{2\pi i\, v \cdot n} e^{2\pi i n \cdot y}.$$

For every real s there is C_s with $\|\phi\|_{H^s} \leq C_s \|F\|_{H^{s+1}}$. For every nonnegative integer m, one also has $\|\phi\|_{C^m} \leq C_m \|F\|_{C^{m+4}}$.

Proof. The lemma bounds the inverse multiplier by $C|n|$. Squaring and summing the weighted Fourier coefficients proves the Sobolev estimate. For the uniform derivative estimate, integration by parts bounds a Fourier coefficient of a C^{m+4} function by $C(1 + |n|)^{-m-4}\|F\|_{C^{m+4}}$. Multiplying by the m derivative factors and the inverse divisor leaves $C(1 + |n|)^{-3}$, summable in two dimensions. This proves uniform convergence of all required derivative series. Applying the same argument for arbitrarily large m proves smoothness. Uniqueness follows because every nonzero Fourier coefficient of a solution of the homogeneous equation vanishes, and the zero coefficient was prescribed. □

Here an integral norm identity yields a finite derivative loss, which is then paid for using smoothness at every fixed order. That is a genuine analytic contribution of the arithmetic structure. It is not an estimate uniform in derivative order, and it would not justify infinite-order flatness from a fixed finite regularity assumption. In Section 8 of the released proof, repeated integration by parts along such directions makes compact-support cutoff errors smaller than any specified power. The number of input derivatives increases with the requested power.

26.5 Symmetric tensors, covariance and the stress correction

For a real velocity perturbation w averaged by a positive probability measure, $\langle w \otimes w \rangle$ is positive semidefinite: for every a its quadratic form is $\langle (a \cdot w)^2 \rangle \geq 0$. It belongs to the symmetric tensor representation. A single $w \otimes w$ has rank at most one, so all its two-by-two minors vanish. The determinantal variety describes this algebraic rank constraint, but averaging rank-one matrices can produce a full-rank matrix. Rank should not be transported through an average without proof.

The released construction needs two specified radial cross-flux components rather than an arbitrary covariance matrix. If b_+ and b_- have disjoint auxiliary supports, their mixed products vanish pointwise. Writing

$$\mathcal{C}(w) = \langle w_r(w_\theta, w_z)\rangle\,, \qquad H = [\mathcal{C}(b_+) \mid \mathcal{C}(b_-)],$$

gives $\mathcal{C}(a_+b_+ + a_-b_-) = H(a_+^2, a_-^2)^T$. A cross-flux vector itself is not a positive semidefinite matrix; positivity concerns the coefficients of this two-ray representation.

Proposition 26.8 (Positive realization and a signed tangent inverse)**.** *Let $H \in \mathrm{GL}_2(\mathbb{R})$, let $s \in \mathbb{R}^2$ satisfy $y = H^{-1}s \in (0,\infty)^2$, and put $a_i = \sqrt{y_i}$. For $C(a) = H(a_1^2, a_2^2)^T$, define*

$$L_a\sigma = \mathrm{diag}(2a_1, 2a_2)^{-1}H^{-1}\sigma.$$

Then $C(a) = s$, $DC(a)L_a\sigma = \sigma$, and

$$C(a + L_a\sigma) = s + \sigma + H\big((L_a\sigma)_1^2, (L_a\sigma)_2^2\big)^T.$$

The last term is generally nonzero.

Proof. The first identity substitutes $a_i^2 = y_i$. Differentiating the coordinate squares gives $DC(a) = H\,\mathrm{diag}(2a_1, 2a_2)$, whose displayed right inverse exists since both amplitudes are positive. Expanding $(a_i + h_i)^2 = a_i^2 + 2a_ih_i + h_i^2$ proves the final identity. □

This is the same distinction between tangent information and an exact finite deformation that appeared in the varieties-of-complexes chapter. The determinant of H proves a local algebraic inverse exists. The PDE proof additionally needs a lower bound for that determinant, control of derivatives of H^{-1}, and flat edge factors that permit division by the leading amplitudes. It keeps those amplitudes fixed while applying signed increments and accounts for the quadratic remainder at each subsequent step. Replacing this procedure by the claim that every target stress is a positive covariance would be false.

The Schur-functor viewpoint is useful here without overreach. In characteristic zero, the flip on $V \otimes V$ yields the complementary projectors $(I+\mathrm{flip})/2$ and $(I-\mathrm{flip})/2$ onto symmetric and alternating tensors. The velocity covariance uses the first; the differential two-form uses the second. Higher tensor lifts use Young symmetries to remove redundant coordinates. Integral straightening theory controls the corresponding bases before denominators are inverted. None of these representation decompositions by itself supplies the scale estimates of the analytic construction.

26.6 Moment determinants and exact compatibility

Compactly supported primitives have obstructions. If $(\partial_r + e/r)\phi = g$ with ϕ compactly supported away from zero, multiplication by r^e and integration gives $\int r^e g\,dr = 0$. Conversely, for real e and $g \in C_c^\infty((0,\infty))$ satisfying that moment condition, $\phi(r) = r^{-e}\int_0^r s^e g(s)\,ds$ is compactly supported and solves the equation. This elementary exactness criterion is the prototype for the weighted radial inverses in the proof. Torus-dependent coefficients require the additional phase shifts and flat remainders described there.

Proposition 26.9 (Independence of ordered moment functionals)**.** *Let $\alpha_1, \dots, \alpha_m$ be distinct real numbers, and let β_j be nonzero, nonnegative smooth bumps supported inside disjoint ordered intervals $0 < I_1 < \cdots < I_m$. Then*

$$B_{ij} = \int_{I_j} x^{\alpha_i}\beta_j(x)\,dx$$

defines an invertible matrix. Smooth variation over a compact parameter set gives bounded inverse derivatives provided all the hypotheses, including distinctness of the exponents, persist there.

Proof. A nonzero linear combination of m distinct powers has at most $m-1$ distinct positive zeros. For induction divide by the smallest power. If the resulting expression is a nonzero constant, it has no zeros. Otherwise its nonzero derivative is a combination of at most $m-1$ distinct powers and has at most $m-2$ zeros by the induction hypothesis; Rolle's theorem then proves the assertion. Hence $\det[x_j^{\alpha_i}]$ never vanishes for $0 < x_1 < \cdots < x_m$. It has constant sign on that connected chamber. Multilinearity and Fubini give

$$\det B = \int_{I_1 \times \cdots \times I_m} \det[x_j^{\alpha_i}] \prod_j \beta_j(x_j)\, dx_1 \cdots dx_m.$$

The integrand has a fixed sign and is nonzero on a set of positive measure. Thus the determinant cannot vanish. Compactness bounds it away from zero for a compact family satisfying the hypotheses. The adjugate formula and repeated differentiation give the inverse-derivative bounds. □

For geometrically rescaled copies of one bump, the matrix becomes a Vandermonde matrix after row normalization. This connects the proof to the same alternating determinants and interpolation combinatorics used earlier. The quantitative limitation is crucial: if two exponents coalesce, the inverse can diverge. The parameter called λ in the source is fixed before the high radial frequency is chosen. Sending it to zero inside an argument that needs a bounded inverse would alter the proof.

Lemma 26.10 (A controlled quadratic moment correction)**.** *Suppose B is invertible, Q is bilinear, $\|B^{-1}\| \le \beta$, $\|Q\| \le \kappa$, and $\|d\| \le d_0$. If $8\beta^2\kappa d_0 \le 1$, the equation $Bc + Q(c,c) = d$ has a unique solution in the closed ball $\|c\| \le 2\beta d_0$.*

Proof. Set $r = 2\beta d_0$ and $T(c) = B^{-1}(d - Q(c,c))$. For $\|c\| \le r$, $\|T(c)\| \le \beta d_0 + \beta\kappa r^2 \le 3r/4 \le r$. For two points in the ball, bilinearity gives $\|T(c) - T(c')\| \le 2\beta\kappa r\|c - c'\| \le \frac{1}{2}\|c - c'\|$. The contraction theorem gives the unique fixed point in that ball. The fixed-point equation is exactly the required moment equation. □

The estimate is a local solvability statement. It does not establish that a whole family of discrepancies is small; that must be proved from the profile construction and its order of parameter choices. This is the appropriate division of labor between algebraic elimination and analysis.

26.7 From varieties of complexes to compatible fluid discretization

Let a finite oriented cell decomposition of a closed surface have real cochains C^0, C^1, C^2 and coboundary matrices $d_0 : C^0 \to C^1$, $d_1 : C^1 \to C^2$ with $d_1 d_0 = 0$. The matrix entries are signed incidence numbers. These matrices give a point of the affine variety of complexes defined by the entries of $d_1 d_0$. Additional rank conditions select the strata studied earlier. The underlying incidence relation is exact over the integers; a floating-point residual is only a computational check of its implementation.

Choose symmetric positive definite mass matrices M_0, M_1, M_2. Define

$$\delta_1 = M_0^{-1} d_0^T M_1, \qquad \delta_2 = M_1^{-1} d_1^T M_2, \qquad L_1 = d_0\delta_1 + \delta_2 d_1.$$

These are the chosen discrete adjoints and Hodge operator. Geometry enters through the M_j, not through the relation $d_1 d_0 = 0$ alone. Stable approximation requires additional commuting projection and norm estimates; finite element exterior calculus makes these requirements precise [7, 8].

Proposition 26.11 (Discrete Hodge energy and incompressible subspace)**.** *For every $x \in C^1$,*

$$x^T M_1 L_1 x = \|\delta_1 x\|^2_{M_0} + \|d_1 x\|^2_{M_2} \geq 0.$$

Moreover L_1 preserves $\ker \delta_1$*. A velocity of the form $\delta_2 \psi$ belongs to this kernel.*

Proof. The definitions identify d_0 and δ_1, and d_1 and δ_2, as adjoints in the indicated inner products. Move each adjoint across the pairing to obtain the sum of squares. Taking adjoints of $d_1 d_0 = 0$ gives $\delta_1 \delta_2 = 0$. Therefore $\delta_1 L_1 = (\delta_1 d_0)\delta_1$, which is zero on $\ker \delta_1$. The last assertion is the same zero-composition identity. □

An arbitrary algebraic change of basis preserves the complex equation. It preserves the numerical energy only if the metric is transported as well. If $x' = Tx$, then the corresponding matrix is $M_1' = T^{-T} M_1 T^{-1}$. Retaining the identity metric after a badly conditioned change of basis can change coercivity constants arbitrarily. This repeats, in the fluid setting, the distinction between homological equivalence and quantitative spectral control.

Theorem 26.12 (An energy-certified finite learned evolution)**.** *Work in orthonormal coordinates on a finite-dimensional divergence-free subspace. Let $K(a)$ be a locally Lipschitz skew-symmetric matrix, L a symmetric positive semidefinite matrix, and $B_\theta(a)$ a locally Lipschitz matrix. For continuous forcing $g(t)$, consider*

$$\dot{a} = -K(a)a - \nu La - B_\theta(a)^T B_\theta(a) a + g(t), \qquad \nu > 0.$$

It satisfies

$$\tfrac{1}{2}\frac{d}{dt}|a|^2 = -\nu a^T L a - |B_\theta(a)a|^2 + a \cdot g(t).$$

Every solution exists for all finite positive times, provided the displayed coefficients are defined and locally Lipschitz on the whole finite-dimensional space.

Proof. For a skew matrix $a^T K(a) a$ equals its own negative, hence is zero. The factorized term has work $|B_\theta(a)a|^2$. This proves the identity. Young's inequality yields $\frac{d}{dt}|a|^2 \leq |a|^2 + |g(t)|^2$. Gronwall gives a finite bound on every finite interval because g is continuous. A locally Lipschitz finite-dimensional ODE can be continued while its solution stays in a compact subset of its domain. The bound therefore rules out finite-time escape and proves global continuation. □

The theorem is a useful design rule for a learned closure, including one whose matrix entries are produced by quantum expectation values. It is also a useful warning: every fixed finite model can be globally well behaved while the continuum family develops unbounded pointwise velocity. The equivalence constant between a grid energy norm and the maximum grid value typically diverges under refinement. Global continuation of each finite ODE cannot be passed to the continuum without uniform regularity and consistency estimates.

The dissipative ansatz is deliberately restrictive. True interscale transfer can include backscatter, while $-B_\theta^T B_\theta$ never injects resolved energy. A trained model in this class must be assessed for that modeling bias. The algebraic certificate proves stability of its own evolution, not fidelity to the exact coarse dynamics.

26.8 Gentle algebras, path relations and quantitative interfaces

The project uses gentle surface algebras to organize permitted local paths and relations. A compatible use here is to let vertices label local field types or mesh patches, arrows label permitted transfers, and quadratic relations encode compositions known to vanish. The equation $d^2 = 0$ is one such zero composition. The derivative-block relation in the Volterra construction, discussed in Chapter 25, is another actual example. Counting admissible words then bounds how often a derivative-losing operator can appear.

One must verify the gentle conditions for the chosen quiver and relations. A free algebra with a single relation $b^2 = 0$ is not automatically gentle: allowed continuations at its other generator can violate the required local uniqueness. Nor does an arbitrary triangulated surface incidence quiver become gentle merely because it is drawn on a surface. The valid general statement is that a representation of a specified path algebra with relations encodes the corresponding exact operator identities. Calling it gentle requires the extra combinatorial checks already described in the earlier chapters.

An associated graded construction is likewise conditional on a specified filtration. If classes $\mathcal{M}_\alpha$ obey $\mathcal{M}_\alpha\mathcal{M}_\beta \subseteq \mathcal{M}_{\alpha+\beta}$, they track orders under multiplication. If a derivative maps $\mathcal{M}_\alpha$ into $\mathcal{M}_{\alpha-\kappa}$, its loss must be included in a correction cycle. Logarithmic factors and flat shell weights are part of the analytic definition. A finitely generated semigroup algebra, a flat toric degeneration, normality or a Schur-module decomposition does not follow solely from this notation. Such claims require the corresponding algebraic hypotheses and proofs.

Standard monomial theory is valuable when one introduces an actual finite coordinate ring of rank constraints, tensor symmetries or moment minors. It provides bases, relations and algebraic elimination. The Navier–Stokes construction further needs positivity margins, compact support and all-order differentiability. These conditions are not encoded by an affine scheme over $\mathbb{C}$. In particular positivity is real semialgebraic or analytic data, and flat nonzero functions cannot be recovered from a convergent analytic Taylor series at their flat endpoint.

26.9 Determinants, spectra and what the reformulation adds

For a finite positive operator L with eigenvalues λ_j, the identities

$$\det(sI + L) = \prod_j (s + \lambda_j), \qquad \frac{d}{ds} \log \det(sI + L) = \operatorname{tr}(sI + L)^{-1}$$

hold for $s > 0$ by diagonalization and ordinary differentiation. Exterior powers and Schur constructions organize coefficients of the determinant. A continuum determinant requires a specified regularization and domain; a spectral zeta function requires convergence or analytic continuation. No identification with an arithmetic L-function, its zeros, or the inherited GRH program follows from these finite identities.

The improvement supplied by this extension is therefore explicit: the integral torus covering is expressed as an order action; its norm becomes a small-divisor estimate; covariance realization becomes a quadratic map with a retained remainder; moment matching becomes a determinant and contraction problem; incompressibility becomes a complex relation; and learned finite models receive a verifiable energy interface. The released proof's analytic construction is preserved rather than replaced. A stronger theorem, such as a result for unforced Navier–Stokes or a quantitative algorithm producing its full singular field, would require additional work beyond these reformulations.

Chapter 27

Quantum information, exact reduction and learned fluid models

Evidence status. The reduction, channel and scalar Carleman statements below are proved explicitly. The computational companion evaluates finite classical simulations of the stated models. Quantum algorithms cited here have their own input, error and readout assumptions. No quantum speedup, trained fluid ToricGT model, or numerical reproduction of the released singular solution is claimed.

27.1 Information reduction is a map, not an autonomous dynamics

The microscopic-to-macroscopic hierarchy in Chapter 24 is naturally expressed by maps that retain observables. Examples include taking a partial trace, a one-particle marginal, a velocity moment, an angular average, or a projection onto a finite set of modes. These maps can be exact even when the subsequent closed evolution is approximate. The distinction is elementary: knowing a projected state need not determine the derivative of that state.

Proposition 27.1 (The obstruction to an exact instantaneous closure)**.** *Consider $\dot{x} = Ax + Bz$ on finite-dimensional vector spaces. If $By \neq 0$ for some hidden vector y, no function $F(x)$ can satisfy $F(x) = Ax + Bz$ for every pair (x, z).*

Proof. For fixed x, compare the hidden states $z = 0$ and $z = y$. The proposed equality gives both $F(x) = Ax$ and $F(x) = Ax + By$, contradicting $By \neq 0$. □

The statement does not rule out a closure on a specified invariant manifold, a conditional expectation under a prescribed ensemble, or an approximation with controlled error. It rules out omitting those hypotheses. It applies equally to a classical neural network and a quantum learned function: an expressive predictor cannot make an insufficient input state mathematically sufficient.

Theorem 27.2 (Exact memory from a linear hidden subsystem)**.** *Let A, B, C, D be constant compatible finite matrices and consider*

$$\dot{x} = Ax + Bz, \qquad \dot{z} = Cx + Dz.$$

Then

$$\dot{x}(t) = Ax(t) + Be^{tD}z(0) + \int_0^t Be^{(t-s)D}Cx(s)\,ds.$$

Conversely, any continuously differentiable x satisfying this equation, together with $z(t) = e^{tD}z(0) + \int_0^t e^{(t-s)D}Cx(s)\,ds$, satisfies the original system.

Proof. Multiply the hidden equation by e^{-tD} and differentiate the product. Integration gives the displayed expression for z. Substitution gives the equation for x. In the converse direction, differentiate the finite matrix-exponential integral using the fundamental theorem of calculus. This recovers $\dot z = Cx + Dz$, and substitution gives the other equation. □

This is a finite, fully explicit version of the memory mechanism associated with projection methods [99]. The memory kernel is $K(t) = Be^{tD}C$, and the hidden initial condition supplies a separate term. Dropping that term is not justified simply by calling it noise. Its distribution and correlations must be supplied if a stochastic description is intended.

For example, with scalar equations $\dot x = -2x + z$, $\dot z = x - 3z$, the exact kernel is e^{-3t}. Even when $z(0) = 0$, replacing the reduced equation by $\dot x = -2x$ discards a nonzero convolution. A fast-relaxation approximation can instead use $z \approx x/3$, but its error depends on the relative time scales and initial layer. Exact elimination and time-scale approximation are different steps.

Proposition 27.3 (Discrete memory and the resolvent determinant)**.** *For $x_{n+1} = Ax_n + Bz_n$, $z_{n+1} = Cx_n + Dz_n$,*

$$z_n = D^n z_0 + \sum_{j=0}^{n-1} D^{n-1-j}Cx_j,$$

and substitution gives the exact resolved recurrence. If $sI - D$ is invertible, the determinant of the full block resolvent satisfies

$$\det\begin{pmatrix} sI - A & -B \\ -C & sI - D \end{pmatrix} = \det(sI - D)\det\big(sI - A - B(sI - D)^{-1}C\big).$$

Proof. The hidden-state formula follows by induction, with the sum empty at $n = 0$. Multiplying by D and adding Cx_n advances the sum to $n + 1$. For the determinant, multiply the block matrix on the left by $\left(\begin{smallmatrix} I & B(sI-D)^{-1} \\ 0 & I \end{smallmatrix}\right)$, whose determinant is one. The upper-right block becomes zero and the upper-left block becomes the displayed Schur complement. The determinant of the resulting block triangular matrix is the product of its diagonal-block determinants. □

This links the project's determinant calculus to an actual information-reduction operation. The frequency dependence of the Schur complement is the transform of a memory kernel. Replacing it by a constant is a modeling approximation. The existence of its inverse does not provide a uniform condition-number bound, especially near a spectral pole.

27.2 Coarse-graining and the unresolved stress

For a normalized nonnegative averaging kernel, write $\bar u = \langle u\rangle$ and $R = \langle u\otimes u\rangle - \bar u\otimes\bar u$. The quadratic nonlinearity ensures that the equation for $\bar u$ contains a contribution involving R. A projection that retains only $\bar u$ has discarded information relevant to the momentum equation.

Proposition 27.4 (Covariance as the exact quadratic defect)**.** *For a probability average with finite second moments (at a fixed observation point, when the average is spatially local),*

$$R = \langle (u - \bar u)\otimes(u - \bar u)\rangle \succeq 0, \quad \operatorname{tr} R = \langle |u|^2\rangle - |\bar u|^2.$$

If a linear filter commutes with spatial derivatives, filtering $\nabla\cdot(u\otimes u)$ gives $\nabla\cdot(\bar u\otimes\bar u) + \nabla\cdot R$.

Proof. Expand the centered product, use $\langle u \rangle = \bar{u}$ and $\langle 1 \rangle = 1$, and take traces. Testing the matrix against any vector gives an average of a square. The final identity follows by linearity, the definition of R, and the assumed derivative commutation. □

For a convolution evaluated at x, the centered integrand in this proof is $u(y) - \bar{u}(x)$, averaged over y. It is not $u(y) - \bar{u}(y)$ for a general non-idempotent filter.

A sharp Fourier truncation need not be a positive averaging kernel. Its quadratic defect is still defined, but the positive-semidefinite conclusion requires the probability-average hypothesis. Similarly, a filter near a wall need not commute with differentiation. Those are concrete places where otherwise plausible information-theoretic arguments fail if their maps are left unspecified.

In the released blowup proof, oscillatory perturbations generate selected averaged quadratic fluxes. The authors construct those fluxes and bound their errors; they do not infer them from an assumed generic turbulence closure. The covariance inverse in Proposition 26.8 captures one finite algebraic component of that construction. Its analytic realization retains the carrier phase, exact curl correction, support separation and derivative estimates.

27.3 A quantum observable can have the same closure obstruction

For a finite quantum system with density matrix ρ and Hermitian observable O, the retained datum $\operatorname{tr}(\rho O)$ is only one linear functional of the state. Even noiseless knowledge of it need not determine its future derivative.

Example 27.5 (Identical observed value, opposite future derivative)**.** Take $\hbar = 1$, Hamiltonian $H = X$ and observable $O = Z$, with Pauli matrices in the conventions of the earlier quantum chapters. The states $\rho_\pm = (I \pm Y)/2$ both have $\operatorname{tr}(\rho_\pm Z) = 0$. Unitary evolution gives

$$\frac{d}{dt}\langle Z \rangle = i\langle [X, Z] \rangle = 2\langle Y \rangle,$$

so their derivatives are $+2$ and -2.

The commutator is obtained by direct multiplication, $[X, Z] = -2iY$. Thus the example requires no interpretive claim about measurement disturbance. The retained statistic simply omits a degree of freedom that matters dynamically. A reduced density matrix can likewise acquire memory when a subsystem is coupled to an environment. A time-homogeneous Markov semigroup requires additional assumptions; it does not follow from partial trace alone.

The same point controls proposed quantum information explanations of viscosity. Closed unitary evolution preserves von Neumann entropy. Effective irreversible evolution arises after a specified reduction and limiting or modeling assumptions. A completely positive trace-preserving map provides a consistent finite open-system operation, but an arbitrary such map is not a derivation of the Navier–Stokes constitutive law.

27.4 An explicit heat channel on a periodic surface grid

Let the position register have orthonormal basis $\{|j, k\rangle : j, k \in \mathbb{Z}/L\mathbb{Z}\}$. Let T_x, T_y be the unitary permutations that translate by one grid point. For $0 \le r \le 1/4$ define

$$\Phi_r(\rho) = (1 - 4r)\rho + r(T_x \rho T_x^\dagger + T_x^\dagger \rho T_x + T_y \rho T_y^\dagger + T_y^\dagger \rho T_y).$$

Proposition 27.6 (A trace-preserving quantum diffusion step)**.** *The map Φ_r is completely positive, trace preserving and unital. Its diagonal probabilities obey*

$$p'_{jk} = (1-4r)p_{jk} + r(p_{j-1,k} + p_{j+1,k} + p_{j,k-1} + p_{j,k+1}).$$

They remain nonnegative, sum to one, and their Shannon entropy cannot decrease.

Proof. Use Kraus operators $\sqrt{1-4r}I$, $\sqrt{r}T_x$, $\sqrt{r}T_x^\dagger$, $\sqrt{r}T_y$, $\sqrt{r}T_y^\dagger$. Their adjoint products sum to I, as do the reversed products. This proves complete positivity, trace preservation and unitality. Taking diagonal entries gives the stencil. It is a convex combination of translations of the probability vector. The function $-s\log s$ is concave on $[0,1]$, with its continuous value zero at zero. Jensen's inequality at each site, followed by summation and translation invariance, proves the entropy assertion. □

Proposition 27.7 (Discrete heat energy and eigenvalues)**.** *Let L be a positive integer, $h>0$ and $\nu \geq 0$. On the periodic $L \times L$ grid, define $\Delta_h p = h^{-2}(T_x + T_x^{-1} + T_y + T_y^{-1} - 4I)p$. Then*

$$-\langle p, \Delta_h p\rangle = h^{-2}\sum_{j,k}((p_{j+1,k}-p_{jk})^2 + (p_{j,k+1}-p_{jk})^2).$$

The Fourier mode $e^{2\pi i(aj+bk)/L}$ has eigenvalue $-4h^{-2}[\sin^2(\pi a/L) + \sin^2(\pi b/L)]$. The heat equation $\dot p = \nu\Delta_h p$ therefore contracts squared Euclidean norm and preserves the constant mode.

Proof. Expand each squared difference and reindex periodic sums to obtain the quadratic identity. Translation acts on the Fourier mode by $e^{\pm 2\pi i a/L}$ or $e^{\pm 2\pi i b/L}$. The formula follows from $2\cos\theta - 2 = -4\sin^2(\theta/2)$. Differentiating $\frac{1}{2}\|p\|^2$ along the heat equation and using the first identity proves contraction. All shifts fix the constant vector. □

With $r = \nu\Delta t/h^2$, the channel gives an explicit heat step. Its positivity requires the displayed time-step bound. The continuous semigroup is the limit of these steps, equivalently the exponential of the random-translation generator. This is an exact finite quantum-channel interpretation of scalar diffusion. The browser computes it classically on a 16×16 grid. The position space has dimension 256, representable by eight qubits, but the computation is not an eight-qubit hardware experiment or a solution of nonlinear Navier–Stokes. A signed velocity field is also not a probability distribution; any encoding of it must preserve signs and normalization by an additional construction.

27.5 What the Madelung transformation does and does not derive

On a simply connected open region where a smooth scalar wave function does not vanish, write $\psi = \sqrt{\rho}\, e^{iS/\hbar}$. For a particle of mass $m>0$ in a real potential V, substitute this into $i\hbar\partial_t\psi = -(\hbar^2/2m)\Delta\psi + V\psi$.

Proposition 27.8 (The local quantum hydrodynamic equations)**.** *With $v = \nabla S/m$ and $Q = -(\hbar^2/2m)\Delta\sqrt{\rho}/\sqrt{\rho}$, one obtains*

$$\partial_t\rho + \nabla\cdot(\rho v) = 0, \qquad \partial_t S + \frac{|\nabla S|^2}{2m} + V + Q = 0.$$

Consequently $m(\partial_t v + (v\cdot\nabla)v) = -\nabla(V+Q)$ and $\nabla\times v = 0$ in that region.

Proof. Differentiate the amplitude-phase product twice. Its Laplacian is

$$e^{iS/\hbar}\left[\Delta\sqrt{\rho}+\frac{2i}{\hbar}\nabla\sqrt{\rho}\cdot\nabla S+\frac{i}{\hbar}\sqrt{\rho}\,\Delta S-\frac{1}{\hbar^2}\sqrt{\rho}\,|\nabla S|^2\right].$$

Equate imaginary parts and multiply by $2\sqrt{\rho}$ to obtain continuity. Divide the real part by $\sqrt{\rho}$ to obtain the second equation. Its gradient gives the velocity equation because the Hessian of S is symmetric, so $\nabla(|v|^2/2)=(v\cdot\nabla)v$. The curl of a gradient is zero. □

There is no viscous term in this calculation. At zeros of ψ the amplitude-phase chart fails; around a closed loop with a single-valued nonzero wave function the phase winding can give circulation $2\pi\hbar n/m$. One scalar phase on a nonvanishing simply connected patch therefore cannot represent an arbitrary smooth vortical incompressible velocity. More general spinor, mixed-state or multi-component encodings may represent richer currents, but their dynamical correspondence must be separately derived. The Born–Oppenheimer and kinetic limits reviewed earlier are distinct routes, not consequences of the scalar Madelung identity alone.

27.6 Carleman lifting and the symmetric algebra

Consider a finite quadratic system $\dot{x}=Ax+B(x\otimes x)+c$, with $x\in\mathbb{R}^n$. Differentiating a tensor power gives

$$\frac{d}{dt}x^{\otimes m}=\sum_{j=1}^{m}x^{\otimes(j-1)}\otimes\dot{x}\otimes x^{\otimes(m-j)}.$$

Substitution couples tensor degrees $m-1,m,m+1$. Thus the infinite collection of powers evolves linearly in a larger space. This is an identity along an existing solution, not a conclusion that the infinite system has a bounded solution for all time.

All tensors $x^{\otimes m}$ are symmetric. In commuting coordinates the degree-m monomials form a basis of dimension $\binom{n+m-1}{m}$ rather than n^m. The symmetric algebra and its standard monomials therefore provide an exact reduction of redundant tensor coordinates. Differentiation acts as a derivation on this algebra. Exterior powers have a different role: they encode determinants and rank constraints, not powers of one commuting vector, whose repeated wedge is zero.

Proposition 27.9 (An exact scalar truncation calculation)**.** *Let $N\geq 1$ be an integer, $\alpha>0$, $\beta\geq 0$, $x_0>0$, and consider $\dot{x}=-\alpha x+\beta x^2$. Put $r=\beta x_0/\alpha$ and $z(t)=r(1-e^{-\alpha t})$. Before the denominator vanishes,*

$$x(t)=\frac{x_0e^{-\alpha t}}{1-z(t)}.$$

Truncate the lifted equations $\dot{y}_m=-m\alpha y_m+m\beta y_{m+1}$ at $m=N$ by setting $y_{N+1}=0$ and $y_m(0)=x_0^m$. Their first component is

$$y_1^{(N)}(t)=x_0e^{-\alpha t}\sum_{j=0}^{N-1}z(t)^j,\qquad x(t)-y_1^{(N)}(t)=x_0e^{-\alpha t}\frac{z(t)^N}{1-z(t)}.$$

If $r<1$, the error is bounded by $x_0r^N/(1-r)$ for all $t\geq 0$. If $r>1$, the nonlinear solution blows up at $t_*=-\alpha^{-1}\log(1-1/r)$; the finite lifted system does not remove that singularity.

Proof. For the exact equation, $w = 1/x$ satisfies $\dot{w} = \alpha w - \beta$. Solving this linear equation yields the displayed formula. The truncated triangular system is solved by downward variation of constants. More explicitly its mth component is

$$y_m^{(N)}(t) = x_0^m e^{-m\alpha t} \sum_{j=0}^{N-m} \binom{m+j-1}{j} z(t)^j.$$

Differentiate this expression, use $z' = \beta x_0 e^{-\alpha t}$ and $j\binom{m+j-1}{j} = m\binom{m+j-1}{j-1}$, and check the initial values. Uniqueness for the finite linear system proves the formula. At $m = 1$ it is the geometric partial sum. Subtracting from the exact rational expression gives the remainder. The bound follows from $0 \le z(t) \le r < 1$. When $r > 1$, the denominator first vanishes at the stated positive time, while the numerator remains positive. □

This calculation is a small, rigorous model of the convergence issue faced by quantum linearization. The finite lift is globally solvable even when it no longer approximates the nonlinear solution. It follows that solving a truncated lifted system, classically or quantum mechanically, cannot establish continuum regularity without a truncation estimate valid in the relevant regime.

27.7 What the quantum-algorithm literature actually supplies

Liu et al. [87] analyze a dissipative quadratic ODE algorithm with a parameter measuring nonlinear and forcing strength relative to linear dissipation. Their stated favorable regime requires $R < 1$, together with sparse access, state preparation, error and solution-norm assumptions. The final output is a normalized quantum state encoding a vector; obtaining a full classical velocity field has an additional readout cost. Their numerical Burgers examples outside the sufficient regime are evidence about those examples, not a theorem covering every Reynolds number. The retrieved arXiv source is version 4, dated 7 July 2026. That correction adds $\|F_0\| \le \|F_2\|$ to the relevant convergence statements. It does not repair a separate eigenvalue-to-norm inference in the printed Lemma 1, which must not be imported into this project.

Proposition 27.10 (A counterexample to eigenvalue-only norm decay)**.** *For the real quadratic system $\dot{u} = F_1 u + F_2(u \otimes u) + F_0$, choose*

$$F_1 = \begin{pmatrix} -1 & 4 \\ 0 & -2 \end{pmatrix}, \quad F_2 = \begin{pmatrix} 1/100 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad F_0 = \begin{pmatrix} 1/1000 \\ 0 \end{pmatrix}, \quad u(0) = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix}.$$

The matrix F_1 is diagonalizable with strictly negative eigenvalues. In the Euclidean and induced spectral norms, the source's parameter is $R = 3\sqrt{2}/500 < 1$, and $\|F_0\| \le \|F_2\|$. Nevertheless

$$\frac{d}{dt}\|u(t)\|^2\Big|_{t=0} = \frac{1007}{2000} > 0.$$

In particular these hypotheses do not imply $\|u(t)\| < \|u(0)\|$ for every positive sufficiently small t.

Proof. The distinct eigenvalues are $-1, -2$, the three relevant norms are $\|u(0)\| = 1/\sqrt{2}$, $\|F_2\| = 1/100$, $\|F_0\| = 1/1000$, and substitution gives the stated R. Direct evaluation gives $\dot{u}(0) = (3007/2000, -1)^T$. Therefore $2u(0)^T\dot{u}(0) = 1007/2000$. The bound asserted by the source's equation (4.2), with spectral real part -1, would instead be $-1 + 3\sqrt{2}/500 < 0$. The exact positive derivative contradicts that bound and yields an interval of initial norm increase by differentiability. □

The problem is that $u^*(F_1 + F_1^*)u$ is controlled by the largest eigenvalue of the Hermitian part, not by twice the largest real part of the eigenvalues of an arbitrary nonnormal matrix. The source specifies Euclidean and spectral norms and does not first transform the system into an eigenbasis. Its scalar rescaling leaves F_1 unchanged. The counterexample concerns that particular printed inference; it is not a claim that every application or a suitably repaired quantum algorithm fails.

Proposition 27.11 (A sufficient dissipative norm hypothesis)**.** *Suppose* $(F_1 + F_1^*)/2 \preceq -\alpha I$ *with* $\alpha > 0$, $\|F_2\| \le \beta$, *and* $\|F_0(t)\| \le \gamma$, *with continuous forcing. If a radius* $a > 0$ *satisfies* $\beta a^2 - \alpha a + \gamma < 0$, *every solution starting in* $\|u(0)\| \le a$ *remains in that ball and continues for all positive times.*

Proof. At a nonzero state of radius r, differentiation and Cauchy–Schwarz give $\dot r \le -\alpha r + \beta r^2 + \gamma$. At the boundary $r = a$ the right side is strictly negative. Continuity makes it negative on a sufficiently small annulus about that radius, so a trajectory cannot first cross that annulus outwards. A boundary initial state immediately moves inward by the same inequality. Thus the ball is invariant. The polynomial vector field is locally Lipschitz in the state, and continuity of the forcing on each finite time interval and boundedness of the trajectory give global finite-dimensional continuation. □

This repaired condition is stronger than spectral stability and is exactly the kind of additional metric information demanded throughout the monograph. A weighted inner product can sometimes establish such a condition, but its conditioning and the transformed nonlinear norm must then enter the quantitative and algorithmic bounds. We use only the independently proved estimates here, not the unrepaired general norm-decay lemma.

Non-dissipative extensions also need hypotheses. Theorem 1.1 of the reviewed version of Wu et al. [179] assumes a diagonalizable F_1, eigenvalues with nonpositive real parts, a uniform solution bound $\|x(t)\|_1 \le \mu$, and the strictly positive resonance gap

$$\Delta = \inf_{k,\, m \in \mathbb{N}^n,\, |m| \ge 2} \left| \lambda_k - \sum_j m_j \lambda_j \right| > 0.$$

Here $\mathbb{N}$ includes zero. Its stated estimate uses $R_r = 4e\mu\kappa_1(W)\|F_2\|_1/\Delta$; geometric convergence requires $R_r < 1$. Merely asserting that each individual resonance expression is nonzero does not establish a positive infimum over infinitely many expressions. The condition also has a particularly restrictive consequence for real systems.

Proposition 27.12 (Real nonresonance excludes neutral oscillatory modes)**.** *Let* F_1 *be a real matrix, with all eigenvalues in the closed left half-plane. If no eigenvalue equals a sum of at least two eigenvalues, allowing repetition, then every eigenvalue has strictly negative real part.*

Proof. A zero eigenvalue violates the hypothesis because $0 = 0 + 0$. A nonzero purely imaginary eigenvalue $i\omega$ of a real matrix has the conjugate eigenvalue $-i\omega$. Then $i\omega = 2i\omega + (-i\omega)$ is a sum with three terms, again violating the hypothesis. These are all possibilities on the imaginary axis, so the nonpositive real parts are strictly negative. □

Consequently, the title of that paper cannot be used to assert a theorem for arbitrary real conservative or neutral fluid systems. Its numerical examples and its precise theorem assumptions must be assessed separately. This observation does not refute its nonresonant estimate; it identifies the real-system domain in which that estimate could apply. We do not use the source's stronger estimate as a dependency of the released singularity proof. A direct estimate adequate for the present finite comparison is as follows.

Proposition 27.13 (A norm-controlled short-time Carleman estimate)**.** *Let $N \geq 1$ be an integer. Suppose $\dot{x} = F_1 x + F_2 x^{\otimes 2}$ on $[0, T]$, with $\|x(t)\| \leq M$. Use compatible Euclidean tensor norms or compatible ℓ^1 tensor norms, and assume $\|e^{tF_1}\| \leq 1$ for $t \geq 0$. For the degree-N Carleman truncation initialized by $y_j(0) = x(0)^{\otimes j}$,*

$$\|x(t) - y_1(t)\| \leq M\big(M\|F_2\|t\big)^N, \qquad 0 \leq t \leq T.$$

In particular the right side converges geometrically in N when $M\|F_2\|T < 1$.

Proof. The diagonal degree-j generator is $A_{j,j} = \sum_{i=0}^{j-1} I^{\otimes i} \otimes F_1 \otimes I^{\otimes(j-1-i)}$. Its semigroup is $(e^{tF_1})^{\otimes j}$ and is contractive under the stated norms. The adjacent block is the analogous sum with F_2 replacing one input slot; its norm is at most $j\|F_2\|$. Write $\eta_j = x^{\otimes j} - y_j$. All initial errors vanish, and the sole terminal forcing is $A_{N,N+1}x^{\otimes(N+1)}$. Repeated variation of constants expresses η_1 as an N-fold ordered integral. The integrand norm is at most $N!\|F_2\|^N M^{N+1}$. The ordered simplex has volume $t^N/N!$, giving the bound. □

Diagonalizability with nonpositive spectral real parts alone does not imply contractivity in the original norm. If $F_1 = W\Lambda W^{-1}$, one can apply the proposition to $z = W^{-1}x$, with the transformed nonlinearity $\widetilde{F}_2 = W^{-1}F_2(W \otimes W)$ and a bound on z. Returning to x costs $\|W\|$. This makes the conditioning visible instead of silently discarding it. Semisimplicity of neutral eigenspaces also differs from claiming that every eigenvalue has geometric multiplicity one. These distinctions correct misleading shortcuts in the preliminary estimates of the reviewed preprint without importing them into our argument.

For our project the required interface is a bound for the particular lifted operator and initial state, on the requested time interval, with its dependence on discretization made explicit.

The September 2026 end-to-end study by Jennings et al. [69] is especially relevant. It accounts for Carleman truncation, time stepping, conditioning, unitary block construction and observable extraction in a lattice-Boltzmann approach. Its shifted formulation and discrete lifting address specific earlier bottlenecks, but its conclusions concern selected observables and parameter regimes. The paper reports numerical convergence and nonconvergence behavior under changing truncation order, and distinguishes proven bounds from fitted scaling. We retain that distinction. Neither its quantum encoding nor its finite Reynolds-number experiments are a proof of the released three-dimensional singularity theorem.

Two further approaches deserve separate names. Variational quantum methods can represent nonlinear functionals using several copies of a state and a chosen ansatz [90]. Their practical result depends on optimization, ansatz approximation, measurement variance and implementation costs. Tensor-network methods such as Gourianov et al. [56] can exploit interscale correlations on classical computers; their compression depends on the ranks needed by the particular flow. A tensor-network simulation is not a quantum-device execution. Both can be useful research tools without claiming a universal exponential speedup.

For every proposed fluid algorithm, the complete cost should specify at least the spatial approximation, time discretization, nonlinear lift or ansatz error, state preparation, success probability, conditioning and requested observable. Normalizing a large vector does not make its small physically relevant component free to read out. A statement polynomial in the number of qubits can still have a large or exponential dependence on physical accuracy, evolution time or a hidden condition number.

27.8 Exact finite repairs for collision block encodings

The detailed algorithm of Jennings et al. [69] makes it possible to inspect the finite maps underlying its collision encoding. In the reviewed arXiv version 2, Appendix D, equation (D56), gives a purported unitary for an equality map, while Appendix I, equation (I65), gives a matrix dilation without its necessary contraction normalization. We supply explicit finite replacements. These proofs concern the two interfaces; they do not establish Carleman convergence, a total resource bound, or quantum advantage for the complete fluid algorithm.

Let $\mathcal{H}$ have orthonormal basis $\{|a\rangle : a \in \{0,1\}^n\}$, with $n \geq 1$. Padding a position register to a power of two permits this notation. Define

$$S = \sum_a |a,0\rangle\langle a,a|, \qquad \Pi_{\mathrm{eq}} = \sum_a |a,a\rangle\langle a,a|, \qquad \Pi_0 = \sum_a |a,0\rangle\langle a,0|.$$

The register 0 is the all-zero bit string. The map S retains equal input positions and resets the second register on that subspace.

Proposition 27.14 (Equality map as a partial isometry)**.** *One has $S^\dagger S = \Pi_{\mathrm{eq}}$, $SS^\dagger = \Pi_0$, and $\|S\| = 1$. The expression*

$$B = S \otimes I + (S - I) \otimes X$$

on two position registers and a flag qubit is not unitary.

Proof. Multiply the two finite sums for S and $S^\dagger$. Orthonormality eliminates every cross term, giving the two projectors. Hence S is isometric on the equality subspace and zero on its orthogonal complement, which proves the norm statement. For any nonzero a,

$$B|a,a,0\rangle = |a,0,0\rangle + |a,0,1\rangle - |a,a,1\rangle.$$

The three output vectors are pairwise orthogonal. Thus a unit input has output squared norm three, contradicting unitarity. □

This counterexample applies to the expression printed in (D56), with its stated square embedding $S = \sum_a |a,0\rangle\langle a,a|$. It is not an obstruction to constructing an equality block encoding.

Proposition 27.15 (Reversible equality encoding and its success probability)**.** *Define $U|a,b\rangle = |a, b \operatorname{xor} a\rangle$. Then U is unitary and $\Pi_0 U = S$. A flag measurement can realize this block with success probability $\|Sx\|^2$ on any normalized input $x \in \mathcal{H} \otimes \mathcal{H}$.*

Proof. The basis permutation U is its own inverse. Applying Π_0 to its output retains exactly the inputs with $b = a$, proving $\Pi_0 U = S$. For an explicit flag, define the reversible basis permutation

$$V|a,c,z\rangle = |a,c,z \operatorname{xor} \mathbf{1}_{c\neq 0}\rangle, \qquad W = V(U \otimes I).$$

Both factors are unitary. Starting the flag in $|0\rangle$ and retaining outcome zero gives

$$(I \otimes \langle 0|)W(x \otimes |0\rangle) = Sx.$$

The Born rule yields success probability $\|Sx\|^2$. Writing $x = \sum_{a,b} x_{ab}|a,b\rangle$ makes this probability $\sum_a |x_{aa}|^2$. When it is positive, the conditional output is $Sx/\|Sx\|$; when it is zero, that outcome never occurs. □

No copying of an arbitrary unknown quantum state occurs here. The XOR operation copies a basis label into a zero register, and on superpositions acts coherently, generally producing correlations. Applied to $x \otimes x$, the equality branch contains squared amplitudes, but preparing two copies already presupposes repeated access to the state preparation. Its success probability remains part of the operation.

Proposition 27.16 (Contraction-normalized rectangular dilation)**.** *Let $A : \mathcal{H}_{\mathrm{in}} \to \mathcal{H}_{\mathrm{out}}$ be a finite complex matrix and let $\alpha > 0$ satisfy $\alpha \geq \|A\|$. Put $C = A/\alpha$ and*

$$D_{\mathrm{out}} = (I - CC^{\dagger})^{1/2}, \qquad D_{\mathrm{in}} = (I - C^{\dagger}C)^{1/2}.$$

Then

$$\mathcal{U}_C = \begin{pmatrix} C & D_{\mathrm{out}} \\ D_{\mathrm{in}} & -C^{\dagger} \end{pmatrix} : \mathcal{H}_{\mathrm{in}} \oplus \mathcal{H}_{\mathrm{out}} \longrightarrow \mathcal{H}_{\mathrm{out}} \oplus \mathcal{H}_{\mathrm{in}}$$

is unitary. On normalized input $(x, 0)$, projection onto the first output summand succeeds with probability $\|Ax\|^2/\alpha^2$ and, when successful, produces $Ax/\|Ax\|$. Conversely, a block of a unitary representing A/α requires $\alpha \geq \|A\|$.

Proof. Take a singular-value decomposition $C = L\Sigma R^{\dagger}$. Every singular value s_j lies in $[0, 1]$. The positive square roots therefore exist, and the SVD gives $CD_{\mathrm{in}} = D_{\mathrm{out}}C$: on each paired singular direction both sides multiply by $s_j\sqrt{1 - s_j^2}$, and both vanish on the unpaired directions. Consequently

$$\mathcal{U}_C\mathcal{U}_C^{\dagger} = \begin{pmatrix} CC^{\dagger} + D_{\mathrm{out}}^2 & CD_{\mathrm{in}} - D_{\mathrm{out}}C \\ D_{\mathrm{in}}C^{\dagger} - C^{\dagger}D_{\mathrm{out}} & D_{\mathrm{in}}^2 + C^{\dagger}C \end{pmatrix} = I.$$

The same calculation gives $\mathcal{U}_C^{\dagger}\mathcal{U}_C = I$. Its action on $(x, 0)$ is $(Cx, D_{\mathrm{in}}x)$, proving the measurement formulas. Finally, the norm of a compression of a unitary is at most one, so any such block must satisfy $\|A/\alpha\| \leq 1$. □

Thus (I65) must use the contraction Σ_2/α, including the same scaling inside both square roots. The reviewed paper's equation (D58a) gives $\alpha = 6/\bar{\tau}^{\star}$ for its $D = 2$ collision matrix. A normalization cannot be removed merely because the requested output is eventually renormalized: it determines the success amplitude and enters the resources of the larger block-encoding algorithm.

Proposition 27.17 (Linear instrument, conditional normalization)**.** *The successful and unsuccessful branches of the preceding dilation define linear completely positive maps*

$$\rho \longmapsto C\rho C^{\dagger}, \qquad \rho \longmapsto D_{\mathrm{in}}\rho D_{\mathrm{in}}.$$

Their total trace equals $\operatorname{tr}\rho$*. Normalizing the successful branch does not turn it into deterministic nonlinear quantum evolution.*

Proof. Each branch has a single Kraus operator and is completely positive: after any ancillary extension it sends a positive operator to a conjugate of that operator. Since $C^{\dagger}C + D_{\mathrm{in}}^2 = I$, cyclicity of trace proves the trace identity. The conditional state is divided by $\operatorname{tr}(C\rho C^{\dagger})$ only after the corresponding random outcome occurs. Discarding that probability would change the physical operation. □

27.9 A global compatibility test for a spin representation

Meng and Yang [97] give a two-component spin representation of viscous flow. Their evolution is nonlinear and non-Hermitian. In the constant-density formulation, their Eq. (15) requires an auxiliary scalar ϕ satisfying

$$\omega \cdot \nabla\phi = -\frac{2\nu}{\rho_0}|\nabla\psi|^2(u\cdot\omega), \qquad \omega = \nabla \times u. \tag{27.1}$$

The source explicitly leaves global solvability open. The following calculation makes a necessary global condition concrete. It concerns a particular global lift, not the existence of the velocity field.

Proposition 27.18 (Periodic compatibility for a smooth global lift)**.** *Let $\mathbb{T}^3 = (\mathbb{R}/2\pi\mathbb{Z})^3$, let $\nu, \rho_0, m, \hbar > 0$, and fix a time. Suppose $\psi : \mathbb{T}^3 \to \mathbb{C}^2$ is a globally defined C^3 periodic spinor with $\psi^*\psi = \rho_0$, and $\phi : \mathbb{T}^3 \to \mathbb{R}$ is a globally defined C^2 periodic scalar. Define*

$$u_j = \frac{\hbar}{m\rho_0}\operatorname{Im}(\psi^*\partial_j\psi), \qquad |\nabla\psi|^2 = \sum_{j=1}^{3}\|\partial_j\psi\|^2, \qquad \omega = \nabla\times u.$$

If (27.1) holds, then

$$\int_{\mathbb{T}^3}|\nabla\psi|^2(u\cdot\omega)\,dx = 0, \qquad |\nabla\psi|^2 \geq \frac{m^2\rho_0}{\hbar^2}|u|^2.$$

Proof. Periodicity and $\nabla\cdot\omega = 0$ give $\int\omega\cdot\nabla\phi = -\int\phi\nabla\cdot\omega = 0$. Integrating (27.1) and using $\nu/\rho_0 > 0$ proves the first assertion. For each coordinate, the complex Cauchy–Schwarz inequality gives

$$|\operatorname{Im}(\psi^*\partial_j\psi)|^2 \leq |\psi^*\partial_j\psi|^2 \leq \rho_0\|\partial_j\psi\|^2.$$

Summing and substituting the definition of u gives the second assertion. □

Proposition 27.19 (A smooth fluid solution excluded by these global-lift assumptions)**.** *For every $\nu > 0$, the velocity and pressure*

$$u(t,x,y,z) = e^{-\nu t}(\sin z, \cos z, 0), \qquad p = 0$$

are a smooth periodic solution of $\partial_t u + (u\cdot\nabla)u = -\nabla p + \nu\Delta u$, $\nabla\cdot u = 0$, for all $t \geq 0$. At no finite time can this velocity satisfy all the global-lift assumptions of Proposition 27.18.

Proof. The components depend only on z, and the third component vanishes. Hence the divergence and advective derivative are zero. Direct differentiation gives $\Delta u = -u$, $\partial_t u = -\nu u$, and $\nabla\times u = u$, establishing the equations. Set $a = e^{-\nu t} > 0$. Then $|u|^2 = u\cdot\omega = a^2$ everywhere. The preceding pointwise bound would imply

$$\int_{\mathbb{T}^3}|\nabla\psi|^2(u\cdot\omega)\,dx \geq \frac{m^2\rho_0}{\hbar^2}a^4(2\pi)^3 > 0,$$

contradicting its necessary zero value. □

The hypotheses include a single-valued periodic auxiliary scalar and a single globally defined smooth periodic spinor. Patchwise lifts, transition functions, singularities, other boundary conditions, or a different auxiliary system require a separate analysis and are not ruled out by this proposition. The example has an explicitly smooth solution: it provides no evidence of fluid blowup. Nor does it contradict the source's acknowledged global qualification. It identifies a concrete obligation before treating a local or conditional representation as a global equivalence for all periodic flows. These two propositions have complete ordinary proofs here; they are not included among the companion's Lean-checked declarations.

A separate algebraic obligation. For a finite real matrix A, let A^+ be its Moore–Penrose pseudoinverse. Substitution shows that $f = A^+S$ solves $Af = S$ exactly if and only if $AA^+S = S$. A singular-value decomposition makes AA^+ the orthogonal projector onto range A; otherwise the residual is $(I - AA^+)S$. At positive rank,

$$\|A^+S\| \le \frac{\|S\|}{\sigma^+_{\min}(A)},$$

where $\sigma^+_{\min}$ is the smallest positive singular value. The compatible example $A_\varepsilon = \operatorname{diag}(\varepsilon, 0)$, $S = (1,0)^T$, $\varepsilon > 0$, has $A^+_\varepsilon S = (\varepsilon^{-1}, 0)^T$. Thus the notation of a pseudoinverse does not remove conditioning problems near rank changes. When applying such an auxiliary solve to a spin-field model, range compatibility and uniform estimates must be established for its actual structured right-hand side. The generic example alone does not decide that structured question.

27.10 A mathematically constrained quantum ToricGT extension

The previously implemented quantum ToricGT was trained and checked on tokenization-path data. Its weights cannot be presented as a validated fluid model. The reusable component is its architecture: incidence-aware vertex and edge tokens, local quantum features, classical aggregation and an explicit certificate after prediction. This section specifies a fluid adaptation and proves the structural certificate independently of training.

For a surface mesh, tokens should include oriented incidence, the chosen mass matrices, local coordinates or invariant metric features, boundary tags, resolved velocity and any retained memory state. A surface's embedding alone is insufficient. Gauge-like changes of cochain basis should transform the fields and metric together. Under mesh relabeling, the local circuit input and classical aggregation must obey a stated equivariance rule; otherwise predictions can depend on arbitrary vertex numbering.

Let $|\psi_\theta(x)\rangle = U_\theta(x)|0\rangle$ be a normalized finite circuit and let $q_j(x) = \langle\psi_\theta(x)|O_j|\psi_\theta(x)\rangle$ for Hermitian O_j with $\|O_j\| \le 1$. Then $q_j \in [-1, 1]$. A classical readout can construct a matrix $B_\theta(x)$ and a skew matrix $K_\theta(x) = \frac{1}{2}(A_\theta(x) - A_\theta(x)^T)$. In divergence-free coordinates, use the vector field of Theorem 26.12. The resulting dissipation identity is true for every choice of circuit parameters; it does not rely on the optimizer finding a good model.

Proposition 27.20 (A parameterized positive constitutive coefficient)**.** *Let $0 < \nu_{\min} < \nu_{\max}$ and let $\ell_\theta(x)$ be any finite real readout of circuit features. Then*

$$\nu_\theta(x) = \nu_{\min} + (\nu_{\max} - \nu_{\min})\frac{1}{1 + e^{-\ell_\theta(x)}}$$

lies strictly between the two bounds. If the readout is locally Lipschitz, using this coefficient in a factorized finite viscosity preserves the energy sign in Theorem 26.12.

Proof. The exponential is positive, so the logistic factor is strictly between zero and one. Its derivative is bounded. Composition with a locally Lipschitz readout is locally Lipschitz. Multiplication of a squared norm by a positive scalar preserves nonnegativity, giving the asserted sign. □

Proposition 27.21 (The circuit derivative and a nonlinear loss)**.** *Suppose one parameter occurs in one gate $e^{-i\theta P/2}$ with $P = P^*$ and $P^2 = I$. For any expectation $q(\theta)$ measured after the full circuit,*

$$q'(\theta) = \tfrac{1}{2}[q(\theta + \pi/2) - q(\theta - \pi/2)].$$

For a differentiable loss $L(q_1, \dots, q_m)$, its derivative is the chain-rule sum $\sum_j (\partial L/\partial q_j) q'_j$. Shifting the entire nonlinear loss is not generally the same operation.

Proof. Expand the gate as $\cos(\theta/2)I - i\sin(\theta/2)P$. The expectation is of the form $a + b\cos\theta + c\sin\theta$ after the fixed preceding and following operations are absorbed into the state and observable. Differentiation and the two shifted trigonometric identities prove the first assertion. The second is ordinary multivariable differentiation. For example, if $q = \cos\theta$ and $L = q^2$, shifting L by $\pm\pi/2$ gives equal values and an erroneous zero, whereas $L' = -\sin 2\theta$ in general. □

Shared parameters in multiple gates require a sum of gate-occurrence derivatives. The previous ToricGT experiments already distinguish this chain rule from an incorrect whole-loss shift. For a fluid extension it is equally essential: a conservation or energy penalty is a nonlinear loss and must be differentiated correctly.

27.10.1 A concrete training and validation specification

A future trained fluid ToricGT should be assessed against the same dynamical problem solved by its classical baselines. Data must identify viscosity, forcing, domain, mesh, time step, reference solver and truncation error. Split by trajectories and physical parameter families, not by randomly mixing nearby snapshots of the same trajectory. Include independently held-out meshes and resolutions. Compare with a classical model of similar trainable size, a frozen-circuit model, the physical baseline without learned closure and an exact manufactured solution.

Report resolved trajectory error, derivative or stress error where available, divergence, energy balance, pressure residual and total wall time including data transfer and readout. A projection or repair can make a constraint pass even when a predictor is inaccurate, so pre-repair and post-repair results must be distinguished. A small training residual is not a certified PDE residual unless its sampling, quadrature and interpolation errors have also been bounded.

Memory is a plausible place to use learning: approximate a stable auxiliary realization of a kernel, with enough state to distinguish histories that share the same instantaneous resolved velocity. This is more defensible than asking a memoryless network to invert an information-losing projection. For a linear stable hidden subsystem, Theorem 27.2 supplies a reference target and a complete exact comparison. For a nonlinear flow, the corresponding error and stability hypotheses require new analysis.

27.11 The executable examples and their proof boundaries

The new laboratory separates four tasks. The Taylor–Green field is an exact two-dimensional periodic Navier–Stokes solution with known pressure and decay, used to test derivatives and normalization. The integer covering panel evaluates the actual order matrix from the released proof. The circuit panel computes a two-qubit state and uses its bounded features to select a positive heat-channel parameter. The concentrating-core panel displays the released scaling exponents, without claiming to compute the source's fully corrected singular field.

The heat circuit is an untrained parameterized example. It is not a fluid-trained ToricGT, and the finite heat channel does not include the nonlinear momentum term. The code records actual finite invariants and numerical discrepancies. The companion's selected Lean declarations formalize energy, projection, finite memory and order identities, with precise component mappings; they do not certify every browser operation or every theorem in these chapters.

Taken together, these examples make information reduction operational. One can see which quantities are preserved by a complex, which are lost under projection, which reappear as a memory kernel or a quadratic stress, and which remain controlled by a positive quantum-channel construction. The released Navier–Stokes proof then occupies its proper place: a continuum construction whose

all-order analytic estimates establish a forced singularity, and whose algebraic components can be understood and checked within this wider framework.

Chapter 28

Exact reduction contracts and the information lost by tropicalization

Evidence status. This chapter adds exact factorization, discrete trajectory, error, and signed polynomial results. The companion module `ReductionContracts.lean` contains sixteen kernel-checked declarations with a per-declaration axiom manifest. The cited analytic-forcing obstruction is an external theorem, not a result formalized by this module.

Information reduction becomes a mathematical claim when the retained data determine the quantity being predicted. This chapter supplies a necessary and sufficient condition, then a quantitative obstruction when the condition fails. It also shows precisely why the orders of two signed terms do not determine the order of their sum. These statements make the project's graph, order, and tropical interfaces testable without identifying them with a continuum proof.

28.1 The exact closure criterion

Let X be a state space, Y the space of retained observations, and Z an output space. A map $P : X \to Y$ describes what the representation retains; $Q : X \to Z$ describes what a model must predict. No topology or probability distribution is needed for the following statement.

Theorem 28.1 (Factorization through retained information)**.** *Suppose P is surjective. There is a map $G : Y \to Z$ satisfying $G(Px) = Qx$ for every x if and only if*

$$Px = Py \quad \Longrightarrow \quad Qx = Qy \qquad (x, y \in X).$$

When it exists, G is unique.

Proof. If G exists, equal observations have equal images under G, proving the condition. Conversely, for each $b \in Y$ choose x_b with $P(x_b) = b$ and define $G(b) = Q(x_b)$. The fiber condition makes this independent of the choice and gives $G(Px) = Qx$. If G, H both factor Q, surjectivity gives, for every $b = Px$, the equality $G(b) = Qx = H(b)$. □

Surjectivity is a modeling convention, not a demand that a compressed representation cover an arbitrary ambient tensor space: one can take Y to be the image of P. The theorem proves existence of a set-theoretic map. It does not prove that G is continuous, measurable, Lipschitz, efficiently computable, or representable by a particular neural architecture. Those are separate requirements on a usable model.

For a discrete update $F : X \to X$, take $Q = P \circ F$. Then the criterion is exactly preservation of the equivalence relation $x \sim y$ defined by $Px = Py$. This gives a precise meaning to an exact reduced dynamics.

Proposition 28.2 (Propagation of an exact discrete reduction)**.** *Suppose $P \circ F = G \circ P$, $u_{n+1} = F(u_n)$, $v_{n+1} = G(v_n)$, and $Pu_0 = v_0$. Then $Pu_n = v_n$ for every natural number n.*

Proof. The initial equality is given. If it holds at n, then $Pu_{n+1} = PF(u_n) = G(Pu_n) = G(v_n) = v_{n+1}$. Induction proves the claim. □

For a differentiable observation of a continuous-time trajectory $\dot{x} = F(x)$, the chain rule instead gives

$$\frac{d}{dt}P(x(t)) = DP(x(t))F(x(t)).$$

The quantity to factor is therefore $Q(x) = DP(x)F(x)$. Only for a linear observation does this become $Q(x) = PF(x)$. A nonlinear tokenizer or learned coordinate map cannot silently use the linear rule. The factorization theorem applies to this Q as a map, while differentiability and existence of trajectories remain analytic premises. The Lean trajectory declaration in this chapter concerns discrete time only.

For example, $P(x, y) = x$ discards the variable needed to predict $Q(x, y) = y$. States $(0, 0)$ and $(0, 1)$ share their observation and have distinct outputs. No deterministic function of the observation alone can give both outputs. Adding a memory state or retaining y can change the representation and remove this obstruction; additional network capacity on the same input cannot.

28.2 A lower bound on every approximate closure

The preceding obstruction has a quantitative version, independent of how G is trained or parameterized.

Proposition 28.3 (Unresolved diameter bounds approximation error)**.** *Let Z be a seminormed additive commutative group. For any $P : X \to Y$, $Q : X \to Z$, and $G : Y \to Z$, if $Px = Py$, then*

$$\|Qx - Qy\| \le \|Qx - G(Px)\| + \|Qy - G(Py)\|.$$

Consequently, if each error on the right is at most ε, then $\|Qx - Qy\| \le 2\varepsilon$.

Proof. Put $g = G(Px) = G(Py)$ and write $Qx - Qy = (Qx - g) + (g - Qy)$. The triangle inequality and $\|g - Qy\| = \|Qy - g\|$ give the first assertion. Bounding the two terms separately gives the second. □

This is a worst-case bound on a specified pair. A probabilistic average-loss claim additionally needs a distribution and a loss; a sampled dataset does not establish a diameter bound on every unseen fiber. The theorem provides a diagnostic for proposed reductions: find equal retained states with separated target stresses or derivatives. A quantum expectation readout, graph embedding, or toric feature map is subject to the same criterion once its accessible classical observation has been specified.

28.3 Signed cancellation is invisible to orders alone

For $p \in \mathbb{Q}[X]$, let $\operatorname{ord}_0(p)$ be its smallest exponent with nonzero coefficient, with $\operatorname{ord}_0(0) = \infty$. The finite-order statements below use nonzero inputs and nonzero residuals, so they do not exploit a convention at zero. Tropical order calculations are useful bookkeeping, but their additive rule must retain the possibility of cancellation; compare the positive-semiring setting of Litvinov [86].

Proposition 28.4 (Arbitrarily large cancellation gain)**.** *For every integer $k \geq 0$, put*

$$p = X, \qquad q_k = -X + X^{k+2}.$$

Then $p, q_k, p + q_k$ are nonzero and

$$\operatorname{ord}_0(p) = \operatorname{ord}_0(q_k) = 1, \qquad \operatorname{ord}_0(p + q_k) = k + 2 > 1.$$

In particular, the identity $\operatorname{ord}_0(p + q) = \min\{\operatorname{ord}_0 p, \operatorname{ord}_0 q\}$ is false in general even for nonzero summands and a nonzero sum.

Proof. The constant coefficients vanish. The coefficient of X in p is one and that in q_k is minus one, since $k + 2 \geq 2$. Each summand has order one and is nonzero. Their sum is exactly X^{k+2}, which is nonzero and has order $k + 2$. This proves all assertions. □

The integer k is arbitrary. Thus no fixed correction to the minimum of the input orders can recover all cancellation gains. Signed leading coefficients must first be compared; when they cancel, higher coefficients must be examined. For complex amplitudes, phases carry analogous information. This issue is especially consequential for a momentum residual: its transport, pressure, temporal, and viscous terms carry signs, and smallness of their exact sum is a different property from the size of each term. An order graph can organize candidate balances. A proof must still establish the coefficient identities and bounds on the remaining terms.

28.4 A current analytic obstruction to removing the force

Theorem 1.1 of Constantin et al. [38], version 2, addresses suitable weak solutions smooth before the candidate singular time. It assumes a uniform spatial C^2 bound on the force, spatial analyticity locally uniform on interior time intervals, the anisotropic angular-mean bounds in its equations (1.7a)–(1.7b) through derivative order two, and exact axisymmetry on a positive-radius core at every earlier time. Under these hypotheses the candidate point is regular. Zero forcing meets the force hypotheses. Consequently, preserving this structural class while deleting the force cannot retain the released singularity. This is an obstruction to that specified class, not a global unforced regularity theorem and not a disproof of the smooth forced construction. The cited analytic theorem is used as an external result; this extension has not formalized its proof.

For the project's future developments, this gives three concrete obligations. A proposed information reduction must state its fibers and the outputs they determine. A tropical residual certificate must retain signed coefficients where leading orders tie. A transformed fluid theorem must preserve its forcing and regularity hypotheses, or separately prove the replacements. These obligations are compatible with graph-based proof search and quantum learning, but they cannot be discharged by an analogy between representations.

28.5 Formal coverage and reproducibility

The companion module proves the two directions and uniqueness in Theorem 28.1, its discrete-closure specialization and a mismatched-fiber obstruction, Proposition 28.2, both inequalities in Proposition 28.3, and every polynomial assertion in Proposition 28.4. The declarations use ordinary set-theoretic functions, actual seminorms, and actual rational polynomials; no placeholder predicate stands in for the missing analytic theorem. The verification manifest records the source hash, exact compiler and mathlib pins, process exit code, and each declaration's foundational dependencies. Its sixteen checked declarations are additional to the earlier finite fluid and quantum modules.

Chapter 29

Corrected transport, information and quantum contracts for tropical quivers

Evidence status. This chapter adds explicit corrections to selected claims in the pinned Tropical Quivers and ToricGT research manuscripts. Seventeen finite declarations in TransportContracts.lean and six in WeightedGainContracts.lean are kernel checked; the latter include a signed matrix-action estimate and a time-varying residual budget. The information-theoretic and stochastic-limit discussions are mathematical arguments with stated hypotheses, not a formalization of all source papers. None of the statements below extends the external forced Navier–Stokes theorem to the unforced problem.

29.1 Source scope and a quantitative contract

The source revisions are:

- Tropical_Quivers_of_Archs [142], commit `2b6ab79a685e2c6233fb8892e59589951feeddad`;
- ToricGT-Program, commit `f91ba1811b5a4f70d53c7120acf8ca2a0049f85b`.

The pinned primary sources are available online.[1] The companion review gives the exact files, inspected sections, and unresolved coverage. The first repository contains a manuscript and its PDF. The second contains five research-program papers, earlier versions, reference PDFs, and a ToricGT gitlink. The gitlink points at a separate code revision. A proposed architecture, a numerical experiment and a proved theorem remain different kinds of evidence.

The useful common interface is a finite typed representation with a specified metric, update map, certificate and error budget. A surface cellulation provides incidence maps; a gentle algebra additionally imposes path relations; a physical fluid requires constitutive laws and analytic estimates; a quantum channel requires complete positivity and trace preservation. Each additional structure has its own contract. The following corrections make those contracts compositional without identifying them.

[1] Primary sources: https://github.com/amelie-iska/Tropical_Quivers_of_Archs/tree/2b6ab79a685e2c6233fb8892e59589951feeddad and https://github.com/amelie-iska/ToricGT-Program/tree/f91ba1811b5a4f70d53c7120acf8ca2a0049f85b.

29.2 Suffix transport is the correct cyclic coboundary

Let $V_0, \dots, V_k$ be vector spaces, let $A_j : V_{j-1} \to V_j$ be linear maps, and let $x_j \in V_j$. Define the edge residual $r_j = x_j - A_j x_{j-1}$ and the suffix transport

$$T_j = A_k A_{k-1} \cdots A_{j+1} : V_j \longrightarrow V_k, \qquad T_k = I_{V_k}.$$

Composition acts right to left. In particular T_j transports a residual from the target of edge j to the endpoint of the path. No commutativity, inverse, equal dimension or flatness assumption is required.

Theorem 29.1 (Typed transport telescope)**.** *With the preceding notation,*

$$\sum_{j=1}^{k} T_j r_j = x_k - A_k \cdots A_1 x_0.$$

For a closed path, $V_k = V_0$ and $x_k = x_0$, this equals $(I - H)x_0$, where $H = A_k \cdots A_1$ is the based holonomy. It vanishes for every initial vector if and only if $H = I$.

Proof. Expand the left side as

$$\sum_{j=1}^{k} A_k \cdots A_{j+1} x_j - \sum_{j=1}^{k} A_k \cdots A_j x_{j-1}.$$

The term in the first sum with index $j < k$ equals the term in the second sum with index $j + 1$. These cancel in their common target space. The two remaining terms are x_k and $-A_k \cdots A_1 x_0$. On a closed path, the residual is zero on every vector precisely when the linear map $I - H$ is zero. For a particular vector it is enough that $Hx_0 = x_0$. □

The equal-space specialization is formalized by induction in `TransportContracts.suffix_transport_telescope`. Its recursive form is

$$s_0 = 0, \qquad s_{j+1} = A_{j+1} s_j + r_{j+1}.$$

This is also an efficient implementation: the accumulated discrepancy is transported once per edge, rather than storing all suffix matrices.

Proposition 29.2 (A flat-cycle counterexample to prefix transport)**.** *Take scalar arrows $A_1 = 2$, $A_2 = 1/2$ and states $x_0 = x_2 = 0$, $x_1 = 1$. Then $A_2 A_1 = 1$. Nevertheless the alternative prefix-weighted expression*

$$A_1 r_1 + (A_1 A_2) r_2 = \frac{3}{2}$$

is nonzero. The suffix expression $A_2 r_1 + r_2$ is zero.

Proof. Here $r_1 = 1$ and $r_2 = -1/2$. Substitution gives $2 - 1/2 = 3/2$ for prefixes and $1/2 - 1/2 = 0$ for suffixes. □

The pinned Tropical Quivers section on elliptic complexes used forward prefix products as transport back to a cycle basepoint. The counterexample shows that its asserted flatness implication does not follow from that formula. Replacing prefixes by the suffixes above repairs the direction and multiplication order simultaneously. The source is retained unchanged; this chapter is an explicit amendment.

29.3 A positive Laplacian does not itself supply a complex

For finite inner-product spaces, let $D_0 : C^0 \to C^1$ and $D_1 : C^1 \to C^2$ be arbitrary linear maps. Define $L_1 = D_0 D_0^* + D_1^* D_1$. Positivity holds with no complex assumption:

$$\langle w, L_1 w \rangle = \|D_0^* w\|^2 + \|D_1 w\|^2 \geq 0, \qquad \ker L_1 = \ker D_0^* \cap \ker D_1.$$

However, an orthogonal exact/coexact/harmonic decomposition requires $D_1 D_0 = 0$.

Proposition 29.3 (Exact hypothesis for orthogonal Hodge components)**.** *The subspaces* $\operatorname{im} D_0$ *and* $\operatorname{im} D_1^*$ *are orthogonal if and only if* $D_1 D_0 = 0$*. Under that hypothesis,*

$$C^1 = \operatorname{im} D_0 \ \oplus^{\perp} \operatorname{im} D_1^* \ \oplus^{\perp} \ \ker L_1.$$

Proof. For every x, y, $\langle D_0 x, D_1^* y \rangle = \langle D_1 D_0 x, y \rangle$. Vanishing of all these pairings is equivalent to $D_1 D_0 = 0$ by nondegeneracy. The orthogonal complement of the sum of the two images is $\ker D_0^* \cap \ker D_1 = \ker L_1$. Finite dimensionality gives the orthogonal direct sum. For arbitrary maps the two images can overlap; positivity of L_1 does not remove that overlap. □

For covariant quiver differences, Theorem 29.1 identifies the curvature defect of each based face as $I - H$. An approximately flat numerical connection is therefore not an exact chain complex. Its residual norm can be reported and penalized, but its "cohomology" must not be treated as the cohomology of an exact differential unless the relation is established. For gentle algebras the corresponding path relations must likewise be checked in the actual representation. Continuum ellipticity, mesh-uniform coercivity and convergence remain additional analytic requirements.

29.4 What a prefix-derived certificate can add

The mutual-information identities in ToricGT II–IV need an explicit distinction between the entire available prefix and a lossy model state. The source definition makes its certificate measurable from the full prefix, so it cannot add Bayes information to that same full prefix.

Proposition 29.4 (Deterministic certificates have zero extra Bayes information)**.** *Let* B, X, Z *be finite random variables and suppose* $Z = f(X)$*. Then*

$$H(B \mid X, Z) = H(B \mid X), \qquad I(B; Z \mid X) = 0.$$

The same conditional-expectation conclusion holds when Z *is measurable with respect to a sigma-field* $\mathcal{F}$ *and* B *has a finite alphabet.*

Proof. For every x of positive probability, the event $X = x$ already fixes $Z = f(x)$. Thus the conditional distribution of B given (X, Z) is the conditional distribution given X. Their conditional entropies coincide, and their difference is the displayed mutual information. In the sigma-field form, $\sigma(\mathcal{F}, Z) = \mathcal{F}$, so both conditional expectations condition on the same sigma-field. □

This does not imply that computable certificates are useless. Let $S = g(X)$ be the representation actually available to a bounded predictor. In general $I(B; Z \mid S)$ can be positive, even though $I(B; Z \mid X) = 0$. For example, take X uniform on $\{0, 1, 2, 3\}$, $B = Z = X \bmod 2$, and $S = \lfloor X/2 \rfloor$. Given S, B is a fair bit; given (S, Z), it is known. The corresponding improvement is one bit, while the improvement beyond complete X is zero.

Proposition 29.5 (Actual predictor improvement includes approximation error)**.** *Let $q_0(\cdot \mid S)$ and $q_1(\cdot \mid S, Z)$ be normalized predictors positive on the support of the target, with finite expected log losses L_0, L_1. In base-two units,*

$$\begin{aligned} L_0 - L_1 = I(B; Z \mid S) \\ + \mathbb{E}\, D_{\mathrm{KL}}\big(p(\cdot \mid S) \| q_0(\cdot \mid S)\big) \\ - \mathbb{E}\, D_{\mathrm{KL}}\big(p(\cdot \mid S, Z) \| q_1(\cdot \mid S, Z)\big). \end{aligned}$$

Proof. For each conditioning value, insert and subtract $-\log_2 p(B \mid \cdot)$ in the expected log loss. This gives conditional entropy plus conditional KL divergence for each predictor. Subtract the two decompositions and use the defining entropy difference for conditional mutual information. □

A deterministic topology, toric or proof certificate can therefore reduce the approximation or optimization error of a restricted neural predictor. That is a measurable training benefit, not new information beyond its complete source data. A deployable certificate should be evaluated by held-out predictive loss and its separately accounted artifact and runtime cost. Leakage control still requires every certificate to be available before the predicted observation; it does not require its usefulness to come from a positive conditional mutual information given all that data. This distinction also prevents a finite topological invariant from being mistaken for extra evidence about a physical system.

29.5 Damping and stochastic limits require their own hypotheses

For a map F with Lipschitz bound L and $0 < \eta \le 1$, the relaxed map $G_\eta = (1-\eta)I + \eta F$ has the upper bound $\mathrm{Lip}(G_\eta) \le (1-\eta) + \eta L$. When $L > 1$, that upper bound is greater than one for every positive η. It cannot certify contraction by decreasing η.

Proposition 29.6 (Damping cannot repair every expanding cycle)**.** *For $F(x) = 2x$ on $\mathbb{R}$, $G_\eta(x) = (1+\eta)x$, so its exact Lipschitz constant is greater than one for every $\eta > 0$.*

Proof. For distinct x, y, the ratio $|G_\eta(x) - G_\eta(y)|/|x - y|$ equals $1 + \eta > 1$. □

The pinned manuscript's assertion that the displayed bound restores contraction for $L > 1$ is therefore incorrect. Damping can still stabilize other maps: for $F(x) = -2x$, the relaxed multiplier is $1 - 3\eta$, whose absolute value is less than one for $0 < \eta < 2/3$. The correct test concerns the actual map, a valid norm bound, or a justified local Jacobian criterion; an unnecessarily large upper bound is not an exact Lipschitz constant.

Proposition 29.7 (A sufficient dissipativity condition for relaxation)**.** *Let b be a map on a real Hilbert space and suppose that, for all x, y,*

$$\langle b(x) - b(y), x - y\rangle \le -\mu\|x - y\|^2, \qquad \|b(x) - b(y)\| \le L_b\|x - y\|,$$

where $0 < \mu \le L_b$. If $0 < \eta < 2\mu/L_b^2$, then $x \mapsto x + \eta b(x)$ is a contraction with bound $\sqrt{1 - 2\mu\eta + L_b^2\eta^2} < 1$.

Proof. For $d = x - y$ and $e = b(x) - b(y)$, expand $\|d + \eta e\|^2 = \|d\|^2 + 2\eta\langle d, e\rangle + \eta^2\|e\|^2$. The assumptions bound this by $(1 - 2\mu\eta + L_b^2\eta^2)\|d\|^2$. The coefficient is nonnegative since its minimum over η is $1 - \mu^2/L_b^2 \ge 0$, and it is less than one on the stated interval. Taking square roots proves the assertion. □

For $G_\eta = (1-\eta)I + \eta F$, apply the proposition to $b = F - I$. This supplies a usable sufficient condition rather than an inference from $\mathrm{Lip}(F) > 1$. A learned skew-plus-dissipative fluid closure may preserve the energy of one trajectory while failing this two-trajectory condition when its coefficients depend on the state; that additional derivative control must be established separately.

The stochastic "embedding fluid" passage needs a second qualification. The continuous/hybrid GFlowNet theory of Lahlou et al. [79] does not, by its balance identities alone, establish a small-step limit. Let a small update have conditional mean $\mathbb{E}[\Delta Z \mid Z = z] = b(z)\Delta t + o(\Delta t)$ and conditional second moment $\mathbb{E}[\Delta Z \Delta Z^T \mid Z = z] = a(z)\Delta t + o(\Delta t)$. For a smooth compactly supported test function φ, Taylor expansion, with a controlled remainder, gives

$$\frac{\mathbb{E}[\varphi(Z+\Delta Z) - \varphi(Z) \mid Z = z]}{\Delta t} \longrightarrow b \cdot \nabla\varphi + \frac{1}{2}\sum_{i,j} a_{ij}\partial_i\partial_j\varphi.$$

Under the additional tightness, moment and convergence assumptions needed for a diffusion limit, the density consequently satisfies

$$\partial_t \rho = -\nabla \cdot (b\rho) + \frac{1}{2}\sum_{i,j} \partial_i\partial_j(a_{ij}\rho).$$

This conservative Markov limit has no independently generated source or sink. Terms $\sigma_{\mathrm{src}} - \sigma_{\mathrm{sink}}$ require separately specified branching, killing or injection dynamics, with their own scaling and convergence hypotheses; they do not follow from the displayed conditional first and second moments alone. The pure continuity equation follows only when the second-order term vanishes in the chosen scaling, or when the current has been defined to include that diffusive contribution. A GFlowNet balance objective by itself does not supply this limit. Likewise an antisymmetric-Jacobian penalty is a useful proposed regularizer, not a theorem that semantic circular reasoning has been eliminated or that physical Navier–Stokes dynamics is reproduced.

29.6 Quantum and gentle-algebra versions of the transport contract

The transport identity extends directly to linear maps between matrix spaces. Let Φ_j be quantum channels and let ρ_j be proposed density operators. Set $r_j = \rho_j - \Phi_j(\rho_{j-1})$. The r_j are Hermitian trace-zero discrepancies, not density operators. Applying Theorem 29.1 to the underlying real vector spaces of Hermitian matrices gives

$$\sum_j (\Phi_k \circ \cdots \circ \Phi_{j+1})(r_j) = \rho_k - (\Phi_k \circ \cdots \circ \Phi_1)(\rho_0).$$

Trace-norm contractivity of each channel on Hermitian discrepancies then bounds the final discrepancy by $\sum_j \|r_j\|_1$. This uses actual channel assumptions; an arbitrary learned real connector does not acquire them by being called quantum. A noisy cycle need not have identity holonomy, and no inverse quantum operation is assumed.

For completeness, this contractivity has a short finite-dimensional proof. For Hermitian Y, spectral decomposition gives $\|Y\|_1 = \max_{-I \le K \le I} \mathrm{Tr}(KY)$. The adjoint of a positive trace-preserving map is positive and unital, so $-I \le K \le I$ implies $-I \le \Phi^*(K) \le I$. Therefore $\mathrm{Tr}(K\Phi(Y)) = \mathrm{Tr}(\Phi^*(K)Y) \le \|Y\|_1$. Taking the maximum proves $\|\Phi(Y)\|_1 \le \|Y\|_1$; the triangle inequality applied to the telescope gives the stated error bound. Complete positivity supplies the stronger physical requirement of positivity after adjoining an arbitrary ancillary system.

Proposition 29.8 (Trace preservation obstructs a zero path relation)**.** *Suppose* Φ *and* Ψ *are composable trace-preserving linear maps on matrix spaces containing a normalized state* ρ*. Then* $\Psi \circ \Phi$ *is not the zero map. Hence a nontrivial relation* $ba = 0$ *of a gentle path algebra cannot be represented by making both composable arrows* a, b *into trace-preserving channels on those full spaces.*

Proof. Trace preservation gives $\operatorname{Tr}(\Psi(\Phi(\rho))) = \operatorname{Tr}\rho = 1$. The zero map would instead produce trace zero. This contradiction uses only linearity, normalization and trace preservation, so it applies in particular to completely positive trace-preserving channels. □

The proposition is formalized, with an abstract linear trace functional, as

`TransportContracts.trace_preserving_composite_nonzero`.

A legitimate quantum interpretation of forbidden paths can use completely positive trace-non-increasing instrument branches, explicit rejection flags, or linear amplitude maps whose ranges and kernels enforce the relation. For instrument branches, a forbidden composition has zero probability; the complete instrument, including all outcomes, is trace preserving. Normalizing a postselected branch is a separate nonlinear state update and its success probability must be retained. This distinction aligns the gentle algebra with operational quantum information without assigning impossible channel semantics to its arrows.

29.7 Tropical approximation is not an automatic universality transfer

Finite ReLU networks give continuous piecewise-affine functions. With arbitrary real slopes they can be expressed using differences of maxima of real affine functions. Classical lattice-based toric geometry imposes integrality or rationality data beyond this statement. Nonzero biases also produce polyhedral complexes rather than necessarily conical fans. These are compatible descriptions once the category and homogenization convention are specified.

Softmax is smooth and is not a finite piecewise-affine map on an open nontrivial logit region. A local piecewise-affine surrogate therefore needs an approximation error bound. Moreover, dot-product logits with both query and key variable are bilinear. A maximum of such scores need not be piecewise affine in the joint inputs. The finite maximum-reduction identity used by ring evaluation remains exact, independently of whether the score functions are tropical polynomials.

The graph-universality argument in the Tropical Quivers manuscript also requires a measure-specific step. The cited primary transformer theorem is Yun et al. [189]; the separate TokenGT construction is Kim et al. [76]. Yun et al.'s Theorems 2–3 use L^p, $1 \le p < \infty$, rather than the uniform norm asserted by the fixed-cell theorem in the program's Paper I. Approximation in ambient Lebesgue L^p does not in general imply approximation on a lower-dimensional tokenized graph family. For example, $g_\epsilon(t) = \max(1 - |t|/\epsilon, 0)$ has $\|g_\epsilon\|^p_{L^p(\mathbb{R})} = 2\epsilon/(p+1)$, tending to zero, while $g_\epsilon(0) = 1$ for every $\epsilon > 0$. A tokenized domain supported at that point sees no convergence. A valid transfer must specify the graph-domain measure or prove uniform approximation there, and must match the graph-relabeling action to the token symmetry group. Continuous detokenization, finite padded sizes and admissible output types must also be retained. Expressive approximation under those assumptions does not establish that gradient descent finds the approximant or that unseen physical distributions generalize.

The bundled transformer-geometry reference is version 1 of `2604.14727`. Its current version 2 [156] explicitly distinguishes fixed-key top-one routing, a log-lifted value parameterization, and

finite-temperature approximations. That distinction should be retained when importing its geometry into a trainable quiver with jointly varying keys and queries. The separate tropical-attention paper [64] proposes a specified algebraic kernel and reports algorithmic benchmarks. Its Appendix C also contains the localized errors repaired next, so this extension does not import that proof package wholesale or transfer its claims to a different dot-product-plus-maximum operation.

29.8 Independent repairs to the tropical-attention reference

Version 2 of Hashemi et al. [64], Appendix C's proof of Lemma 3.1, uses $v_\lambda(a) = \log a - \lambda$ for $a > 0$ and claims a non-Archimedean additive inequality. This logarithm is an Archimedean coordinate change. Neither a constant offset nor projective normalization turns ordinary addition into exact maximum.

Proposition 29.9 (Exact logarithmic identities and dequantization error)**.** *For $a, b > 0$ and real λ,*

$$v_\lambda(ab) = v_\lambda(a) + v_\lambda(b) + \lambda.$$

For any $m \geq 1$, $z_1, \ldots, z_m \in \mathbb{R}$ and $\tau > 0$,

$$M \leq \tau \log \left(\sum_{j=1}^{m} e^{z_j/\tau} \right) \leq M + \tau \log m, \qquad M = \max_j z_j.$$

In particular, for $a = b = 1$, $v_\lambda(a + b) = \log 2 - \lambda > \max\{v_\lambda(a), v_\lambda(b)\}$.

Proof. The product identity follows by expanding $\log(ab) = \log a + \log b$. Factoring $e^{M/\tau}$ out of the sum leaves a number between one and m. Apply the increasing logarithm and multiply by $\tau > 0$ to obtain the two bounds. The displayed counterexample follows from $\log 2 > 0$. □

Thus the source's product offset has the wrong sign, and its additive inequality omits a strictly positive term even on positive inputs. The valid limiting statement is convergence to maximum as $\tau \downarrow 0$ for a fixed finite number of terms, with the explicit error above. If the number of terms varies, $\tau \log m$ must also be controlled. For signed or complex terms, cancellation requires additional information, as proved in Proposition 28.4. These corrections do not invalidate ordinary maximum-plus dynamic programming.

The same Appendix C, Lemma C.1, uses a large finite penalty as if it were the exact value $-\infty$. A finite mask can be exact on a bounded input domain, but it is not an input-independent annihilator on unbounded values.

Proposition 29.10 (When a finite tropical mask is exact)**.** *Suppose a desired maximum has at least one admissible candidate of score at least L. Suppose every excluded unpenalized candidate has score at most U. Giving every excluded candidate an additive penalty $-M$ makes it strictly unable to attain the maximum whenever $M > U - L$. Without an upper bound on excluded values, no fixed finite M guarantees this conclusion.*

Proof. Every penalized excluded score is at most $U - M < L$, below an admissible score. Conversely, with an admissible score zero and any finite penalty M, an excluded raw score $M + 1$ becomes one and wins. □

An explicit $-\infty$ structural mask, or a proved range and sufficient finite margin, repairs the construction. An all-masked query needs its own defined output. Bounds on learned input features and their perturbations must be included in a claimed finite-margin certificate. The source also assigns $-\infty$ to excluded value entries in that proof. If such entries are genuinely enforced, they remove those candidates independently of the finite score penalty. The defect identified here is the inference from a large finite penalty alone; it is not a counterexample to a separately and exactly masked maximum.

Proposition 29.11 (Exact-step and bounded-length tropical paths)**.** *Write $\oplus = \max$ and $\otimes = +$ for scalar tropical operations, with their induced matrix operations. Let D be a finite tropical matrix and let I have zero diagonal and $-\infty$ elsewhere. Then*

$$v_{t+1} = D \otimes v_t \quad \Longrightarrow \quad v_t = D^{\otimes t} \otimes v_0,$$

whereas

$$w_{t+1} = w_t \oplus (D \otimes w_t), \quad w_0 = v_0 \quad \Longrightarrow \quad w_t = \left(\bigoplus_{j=0}^{t} D^{\otimes j} \right) \otimes v_0.$$

Thus a pure t-step update and a finite Kleene closure are different operations in general.

Proof. The first identity is induction and associativity. For the second, distribute D over the inductive sum; adding the previous sum retains precisely powers zero through $t+1$. Idempotence removes duplicate terms. For a direct counterexample take the one-by-one matrix $D = (-1)$ and $v_0 = 0$. At $t = 2$ the pure update is -2, while the maximum of the zeroth, first and second powers applied to v_0 is zero. Even the source's convention summing powers zero through $t-1$ gives zero, not -2. □

This repairs the mismatch in Theorem C.5 of the reference. For maximum-plus longest-path closure, deleting a repeated cycle can only improve a path when every cycle has nonpositive total weight. Under that hypothesis a best finite path can be chosen with at most $n-1$ edges, with the length-zero path allowed. For minimum-plus shortest paths the corresponding hypothesis is absence of negative-weight cycles. The sign and semiring cannot be interchanged. On graphs with profitable cycles, a bounded-horizon computation is still defined, but it is not the unbounded closure.

These are complete finite arguments for the operations actually used. They do not assert that a trained attention stack has learned the exact parameters, nor that any method solves arbitrary NP-hard instances in polynomial resources. The end-to-end scientific contract remains a measured prediction together with independently checked algebraic structure.

29.9 A nonnormal obstruction to a Jacobian power certificate

The Tropical Quivers manuscript's section on Jacobian spectral-radius estimation iterates $v \mapsto Jv/\|Jv\|$ but labels the resulting estimate as the largest singular value. For nonnormal matrices these are different quantities. The following exact witness isolates the error without any floating-point approximation.

Proposition 29.12 (A nilpotent Jacobian with arbitrarily large gain)**.** *For any real $M > 1$, put*

$$J_M = \begin{pmatrix} 0 & M \\ 0 & 0 \end{pmatrix}, \qquad e_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Then $J_M^2 = 0$, $\rho(J_M) = 0$, and $\|J_M\|_2 = M$. Nevertheless $\|J_M e_2\|_2^2 = M^2 > 1$, while the next application of J_M is zero. In particular, a vanishing or small eigenvalue power probe does not certify one-step Euclidean contraction.

Proof. Direct multiplication gives $J_M^2 = 0$ and $J_M^T J_M = \operatorname{diag}(0, M^2)$, whereas $J_M J_M^T = \operatorname{diag}(M^2, 0)$. Thus J_M is nonnormal and both eigenvalues are zero. For $x = (x_1, x_2)$, $\|J_M x\|_2^2 = M^2 x_2^2 \le M^2 \|x\|_2^2$, with equality at $x = e_2$. This proves the norm formula and all remaining assertions. □

For a general real Jacobian, singular-norm power iteration uses $J^T J$, requiring both a Jacobian-vector product and its adjoint product. Even that correction does not turn a finite numerical estimate into a certified upper bound: a Rayleigh quotient samples a direction and can miss the maximizing singular direction. A contraction certificate needs a proved upper bound, a validated enclosure, or an explicitly controlled operator parametrization. The nilpotent example can have powers that eventually vanish while still amplifying perturbations at the first step.

The companion formalizes its explicit matrix products, nonnormality and squared coordinate gain. The spectral-radius and operator-norm conclusions are the written finite-dimensional argument above; the corresponding Lean components are not represented as a complete analytic norm formalization.

29.10 Path multiplicity separates tropical growth from operator stability

The max-plus cocycle in Schreiber [142] records the largest logarithmic path weight. An ordinary matrix product adds contributions from many paths. Its operator growth therefore needs a multiplicity term. This is a qualification of the source's contraction interpretation, not a rejection of its max-plus subadditivity calculation.

Let $A_t : \mathbb{C}^{d_{t-1}} \to \mathbb{C}^{d_t}$, $1 \le t \le T$ and $T \ge 1$, with each $d_t \ge 1$, and set $P_T = A_T \cdots A_1$. An admissible path from j to i is a sequence $(i_0, \dots, i_T)$ with $i_0 = j$, $i_T = i$ and all $A_t(i_t, i_{t-1}) \ne 0$. Let $C_T(i,j)$ count these paths and let $W_T(i,j)$ be their maximum absolute product weight, taking both values to be zero when no admissible path exists. Thus W_T is a maximum-product matrix product; its logarithm, with $\log 0 = -\infty$, is the corresponding max-plus product. Let S_t be the zero–one support matrix of A_t.

Theorem 29.13 (Finite path-multiplicity bound)**.** *With the preceding notation,*

$$|P_T(i,j)| \le (|A_T| \cdots |A_1|)(i,j) \le C_T(i,j) W_T(i,j), \qquad C_T = S_T \cdots S_1.$$

Writing $N_T = \max_i \sum_j C_T(i,j)$ and $m_T = \max_{i,j} W_T(i,j)$ gives the induced row-sum norm bound

$$\|P_T\|_\infty \le \max_i \sum_j C_T(i,j) W_T(i,j) \le N_T m_T.$$

For square n-dimensional arrows, $N_T \le n^T$. Consequently, whenever $m_T > 0$ eventually and $\limsup_{T\to\infty} T^{-1} \log m_T + \log n < 0$, the ordinary matrix products decay exponentially in this norm.

Proof. Expand each entry of P_T as the sum over the intermediate indices. The triangle inequality bounds its absolute value by the sum of the absolute path weights. Each nonzero summand is at most $W_T(i,j)$, and exactly $C_T(i,j)$ such paths occur. Multiplying the support matrices counts those paths, proving the entrywise assertions. Sum over j and maximize over i to obtain the norm inequalities. In dimension n, a fixed final index leaves T freely chosen earlier indices, so there are at most n^T paths ending at that index. The strict limsup inequality then gives constants $c > 0$ and T_0 with $n^T m_T \le e^{-cT}$ for $T \ge T_0$. □

If $m_T = 0$ at any time, every path of that length vanishes, so all later products are zero and the conclusion holds directly. More precise bounds can retain N_T or the entire entrywise table instead of n^T. Using $\log(\epsilon + |A_t(i,j)|)$ with $\epsilon > 0$, as in the source, increases path weights and produces a conservative maximum-product upper bound; it does not remove the number of summed paths. Conversely, signed or complex cancellation can make P_T smaller than the positive bound. These inequalities do not recover such cancellation from magnitudes.

Proposition 29.14 (Decaying maximal paths with an expanding linear map)**.** *For $n \geq 2$ and $1/n < a < 1$, let $A = a\mathbf{1}\mathbf{1}^T$ on $\mathbb{R}^n$. Every length-k path has weight a^k, so its tropical growth exponent is $\log a < 0$. Nevertheless*

$$A\mathbf{1} = na\mathbf{1}, \qquad A^k = n^{k-1}a^k\mathbf{1}\mathbf{1}^T, \qquad \|A^k\|_\infty = (na)^k > 1 \quad (k \geq 1).$$

The path-multiplicity norm bound is attained.

Proof. Each entry is a, and each row has n entries. The eigenvector identity follows by row summation. Since $(\mathbf{1}\mathbf{1}^T)^2 = n\mathbf{1}\mathbf{1}^T$, induction gives the power formula. Its row sum is $(na)^k$. There are n^{k-1} paths between any prescribed pair of endpoints and n^k paths ending at a prescribed final vertex, so both multiplicity inequalities are equalities. □

For example, $n = 2$, $a = 3/4$ gives individual path decay $(3/4)^T$ and ordinary growth $(3/2)^T$. Thus a negative maximal-path exponent alone cannot certify contraction of a learned fluid closure, a quiver update, or a quantum-amplitude surrogate. A separately proved norm inequality or a preserved stochastic/channel structure is required. The exact regression computes support counts and maximum weights by dynamic programming and compares them against exhaustive path enumeration on small, possibly rectangular, signed rational examples. It is a validation of finite products, not an estimate of an infinite-time physical exponent.

29.11 Weighted small-gain certification at a fixed rate

The weighted small-gain criterion in Schreiber [142] is useful, but its instruction to minimize the weighted row bound by a linear program needs a qualification. Suppose $G \geq 0$ is a finite matrix of valid block Lipschitz gains. For positive weights w_i define $\|x\|_{\infty,w} = \max_i \|x_i\|/w_i$. Its row bound is

$$q(w) = \max_i \frac{(Gw)_i}{w_i}.$$

Proposition 29.15 (Fixed-rate feasibility is linear)**.** *For fixed $q \geq 0$, $q(w) \leq q$ is equivalent to the componentwise linear inequalities $Gw \leq qw$, $w_i > 0$. If fixed bounds $0 < \ell_i \leq w_i \leq u_i$ are imposed, feasibility is a linear-program problem. Feasibility is monotone in q. Jointly minimizing q over w in this representation is not a linear program because qw_i contains products of variables.*

Proof. Multiply each row-ratio inequality by its positive denominator w_i. At fixed q, every resulting coefficient is fixed, so the inequalities and box bounds are linear. If a weight vector is feasible at q, it remains feasible at every larger rate because its coordinates are positive. If q is also variable, qw_i is bilinear, which establishes the final qualification. □

One may bisect a certified bracket of rates and solve the fixed-rate feasibility problems. A returned numerical witness should still be checked with exact arithmetic or validated error bounds

before it is used as a proof of contraction. The checker supplied here verifies a given rational witness and does not claim to solve the optimization problem. For

$$G = \begin{pmatrix} 0 & 2 \\ 1/8 & 0 \end{pmatrix}, \qquad w = \begin{pmatrix} 4 \\ 1 \end{pmatrix}, \qquad Gw = \tfrac{1}{2}w,$$

the unweighted row bound is two, while the weighted bound is $1/2$. For any positive feasible weights, the two inequalities multiply to $1/4 \le q^2$, so a proposed rate $q = 2/5$ is impossible. These exact equalities and the rejection are included in the regression. Bounds on G must hold over the claimed domain; a few sampled Jacobians do not establish them.

29.11.1 Propagating a certified gain through time

The feasibility certificate becomes useful for dynamics only after a comparison inequality has been established for the actual update. The following result states that missing link explicitly and retains every per-step perturbation.

Theorem 29.16 (Common weighted gain and accumulated residuals)**.** *Let I be a nonempty finite index set, $w_i > 0$, and $q \ge 0$. For each $n \ge 0$ let G_n be an entrywise nonnegative real matrix with $G_n w \le qw$. If real matrices A_n satisfy $|(A_n)_{ij}| \le (G_n)_{ij}$, then*

$$|x_j| \le r w_j \ (j \in I), \quad r \ge 0 \quad \Longrightarrow \quad |(A_n x)_i| \le qr w_i \ (i \in I).$$

More generally, suppose $r_0 \ge 0$, $\delta_n \ge 0$ and nonnegative block errors $d_{n,i}$ obey

$$d_{0,i} \le r_0 w_i, \qquad d_{n+1,i} \le \sum_j (G_n)_{ij} d_{n,j} + \delta_n w_i.$$

Then

$$d_{n,i} \le R_n w_i, \qquad R_n = q^n r_0 + \sum_{k=0}^{n-1} q^{n-1-k} \delta_k. \tag{29.1}$$

In particular, if $u_{n+1} = A_n u_n + e_n$, $|u_{0,i}| \le r_0 w_i$ and $|e_{n,i}| \le \delta_n w_i$, the same conclusion holds with $d_{n,i} = |u_{n,i}|$. For $0 \le q < 1$ and $\delta_n \le \bar{\delta}$, this gives

$$\|u_n\|_{\infty,w} \le q^n r_0 + \bar{\delta} \frac{1 - q^n}{1 - q}.$$

The finite-sum formula also applies at $q = 0$ and $q = 1$; no division by $1 - q$ is required there.

Proof. For a single signed matrix action, the triangle inequality gives

$$|(A_n x)_i| \le \sum_j |(A_n)_{ij}|\, |x_j| \le r \sum_j (G_n)_{ij} w_j \le rq w_i.$$

For the error comparison, $R_0 = r_0$ and the finite sum in (29.1) satisfies $R_{n+1} = qR_n + \delta_n$. All R_n are nonnegative. If $d_{n,j} \le R_n w_j$, the nonnegative entries of G_n imply

$$d_{n+1,i} \le R_n (G_n w)_i + \delta_n w_i \le (qR_n + \delta_n) w_i = R_{n+1} w_i.$$

Induction proves the comparison. In the signed recurrence, $|u_{n+1,i}| \le \sum_j (G_n)_{ij} |u_{n,j}| + \delta_n w_i$ follows from the triangle inequality, so the comparison applies. The uniform-perturbation assertion follows by bounding every δ_k by $\bar{\delta}$ and summing a finite geometric series. □

For a nonlinear map between normed blocks, one must prove

$$\|F_{n,i}(x) - F_{n,i}(y)\| \le \sum_j (G_n)_{ij} \|x_j - y_j\|$$

on a domain containing both trajectories. A uniform derivative bound on a convex domain is one sufficient route; a switching or nonsmooth map needs its own comparison argument. Invariance of that domain is a separate obligation. Time-varying matrices need not commute, and individual spectral-radius bounds do not replace a common positive weight vector. For $q < 1$, the unforced comparison contracts initial discrepancies; uniformly bounded persistent errors give the displayed asymptotic error upper bound. For $q \ge 1$, a valid finite-time bound is still useful but does not establish contraction.

The six exact declarations in `extension/lean/WeightedGainContracts.lean` prove the row-ratio equivalence, rate monotonicity, gain comparison, signed matrix estimate, and the two finite residual-convolution statements over an ordered field. The rational and real instances therefore share the same algebraic proof. The nondividing comparison declarations use weaker hypotheses where possible. The interpretation as a weighted norm, the nonlinear domain obligation, and the displayed closed geometric-series corollary are written arguments; the whole composite theorem is not registered as fully formalized.

29.11.2 An exact certificate has an exact input domain

The boolean API in `extension/code/verify_toric_transport.py` now accepts only literal Python integers and `Fraction` values, with Boolean values excluded. It checks the input types before any arithmetic, as well as square shape, nonnegative gains and rate, and strictly positive weights. An earlier implementation documented rational inputs but did not enforce that domain: a subnormal floating-point weight could round both sides of a false inequality to the same value, and an infinite weight could make two infinite quantities compare equal. The continuation record preserves both rejected historical cases, including the exact rational reconstruction of the floating-point values. This corrects an executable acceptance condition; the rational theorem is unchanged.

The companion `weighted_gain_witness` serializes the matrix, weights, proposed rate, every row product, right-hand side, signed slack, row ratio, and maximal ratio. Each rational is a reduced pair of decimal integer strings with positive denominator. Thus an integer larger than the exact range of a JavaScript number remains exact in JSON. A separate `check_weighted_gain_witness.py` recomputes the entire certificate using integer addition, multiplication, greatest common divisors and cross multiplication; it imports neither the producer nor `Fraction`. It rejects duplicate keys, noncanonical rational representations, missing or additional fields, and altered derived data. The checker uses explicit conditions that remain active under Python optimization.

Its exit code is zero for an accurately recorded holding bound, two for an accurately recorded failed bound, and one for malformed or inconsistent evidence. A holding rate at least one retains the explicit flag `certifies_contraction=false`. Parser resource limits are separate from the mathematical statement. This replay checks the supplied finite inequalities; it does not authenticate a file, prove a global model-gain bound, verify the Python interpreter, or certify an optimizer. The source hashes, positive and negative replay examples, and adversarial regression results are retained with the release evidence.

29.12 Proper scoring must include its probability boundary

The cross-entropy theorem in Schreiber [141] optimizes over $q(Z) \in (0,1)$ but does not exclude deterministic conditional labels. When $\eta(Z) = \mathbb{P}(Y = 1 \mid Z)$ equals zero or one, its stated minimizer lies outside that open interval.

Proposition 29.17 (Binary log loss with correct endpoint conventions)**.** *For $\eta \in [0,1]$, extend*

$$R_\eta(q) = -\eta \log_2 q - (1-\eta)\log_2(1-q)$$

to $q \in [0,1]$ by assigning zero to a zero-probability term and $+\infty$ to a positive-probability event assigned probability zero. Then R_η is uniquely minimized at $q = \eta$. Over the restricted domain $q \in (0,1)$ the minimizer exists precisely when $0 < \eta < 1$; for $\eta = 0$ or 1, its infimum is zero and is not attained.

Proof. If $0 < \eta < 1$, both endpoints have infinite risk. In the interior, $R'_\eta(q) = (q-\eta)/(q(1-q)\ln 2)$, which is negative below η and positive above it. Thus the interior minimum is unique. If $\eta = 0$, the risk is $-\log_2(1-q)$, uniquely minimized at zero. Every interior q has positive risk and is improved by replacing it with $q/2$; the risk tends to zero as $q \downarrow 0$. If $\eta = 1$, the risk is $-\log_2 q$, uniquely minimized at one. Every interior q is improved by $(1+q)/2$, and the risk tends to zero as $q \uparrow 1$. These cases prove all assertions. □

Applying this argument conditionally gives the Bayes statement for measurable forecasts with the corresponding admissible range. A sigmoid with finite logits produces only interior probabilities, so deterministic targets need not admit a finite-logit optimum. The proper-scoring theorem does not by itself prove empirical calibration of a restricted, finitely trained oracle. Calibration and useful held-out control remain measured properties, with the prefix and artifact accounting established earlier.

29.13 Predictive synergy is not determined by geometric overlaps

The last context-utility theorem of Schreiber [140] proposes an additive information decomposition over independent sheaf components, with overlap corrections determined by Cech data. The usual chain rule for mutual information contains conditional statistical terms; independence of source components alone does not remove predictive synergy.

Proposition 29.18 (Independent disjoint components can carry one bit jointly)**.** *Let X, Y be independent fair bits, let $B = X \operatorname{xor} Y$, and let the additive statistic be $T = X + Y$. Then B is determined by T, but*

$$I(B;X) = I(B;Y) = 0, \qquad I(B;T) = I(B;X,Y) = 1.$$

The two source components may be stored on disjoint pieces with no overlap restriction data.

Proof. The four pairs (X,Y) have probability $1/4$. Conditional on either value of either source bit, B remains a fair bit, so both individual mutual informations vanish. The target is a fair bit overall and equals $T \bmod 2$, so its conditional entropy given T is zero. It is also determined by the pair (X,Y). Its unconditional entropy is one bit, which proves the other equalities. The storage locations of the two source bits impose no restriction on this joint probability law. □

For arbitrary finite variables and conditioning context C, the exact decomposition is

$$\begin{aligned} I(B; X, Y \mid C) &= I(B; X \mid C) + I(B; Y \mid X, C) \\ &= I(B; X \mid C) + I(B; Y \mid C) + I(X; Y \mid B, C) - I(X; Y \mid C). \end{aligned}$$

Both identities follow by expanding conditional mutual information as differences of conditional entropies and applying the entropy chain rule. Even if $I(X; Y \mid C) = 0$, conditioning on the target can make the sources dependent. In the example, $I(X; Y \mid B) = 1$, and there is no source overlap whose absence would force this correction to vanish. Additional probabilistic factorization assumptions are needed to replace these terms by a particular geometric-overlap model. Compatibility of local sections alone is insufficient. If X, Y are already determined by the complete conditioning prefix, all extra conditional information is zero, as proved earlier; the example applies when the actual predictor has retained only a smaller context.

The regression checks the complete rational probability table, all pairwise-independence equalities, deterministic recovery from T, and the conditional dependence given B. Its reported one-bit values follow from these fair and deterministic distributions. The general entropy and log-loss arguments in this section are written proofs; they are not added to the count of Lean-checked declarations.

29.14 Implications for the research program

The repaired transport identity provides a shared finite contract for surface cochains, quiver connectors, quantum instruments and learned trajectory residuals. It can support a structure-preserving fluid model, but a Navier–Stokes proof still needs the external construction's analytic estimates. The corresponding GRH/Hilbert–Pólya program still requires a precise arithmetic spectral identification and operator-domain control; finite self-adjoint matrices do not identify their spectra with zeta zeros. A positive finite Hodge spectrum likewise does not establish a uniform continuum Yang–Mills mass gap. Finally, a path-algebra certificate can authenticate a structural computation only when an actual cryptographic protocol binds it to a message and a threat model; difficult-looking tropical equations are not, by themselves, hardness assumptions.

The exact regression script tests 288 generic rational paths, additional rectangular paths, 128 flat cycles with noncommuting arrows, the prefix counterexample, positive damping examples and the finite information example. It also compares tropical dynamic programming against exhaustive bounded-path enumeration and records the finite-mask and horizon counterexamples. The new path-multiplicity checks also compare exact products with an independent enumeration, including changing dimensions, and retain a dense example attaining the bound. Fixed-rate small-gain witnesses, the proper-scoring endpoint inequalities and the full XOR probability table are checked separately. The logarithmic statements are ordinary mathematical proofs above; the exact regression encodes their positive-input counterexamples using the monotonicity of the logarithm, without treating floating-point logarithms as formal certificates. The Lean log records the exact theorem types and axiom dependencies for its seventeen selected declarations. Those records state what has been checked and leave the larger scientific and analytic obligations visible.

Chapter 30

Quantum–tropical fluid realizations: operators, control and substrates

Evidence status. This chapter develops an explicitly proposed realization of ToricGT-style graph models, together with finite mathematical results and an exact two-qubit feedback example. Its operator, tropical and control layers have separate contracts. The physical literature supplies demonstrated building blocks, not an experimental realization of the combined architecture. Generative substrate design and the mappings to natural systems below are research proposals. The new finite arguments are written proofs; only the weighted-gain components identified in the formal type map have Lean coverage. No new fluid-trained model, fabricated substrate or quantum speedup is reported.

30.1 What a realization must specify

Here a *quantum–tropical fluid realization* means a model with three explicit ingredients: quantum states and admissible operations on them; graph-based routing or a logarithmic asymptotic description; and dynamics whose relation to a fluid or other physical substrate is quantitatively specified. It is a proposed interface, rather than a claim that these three descriptions are interchangeable. A useful realization must identify an encoding, an evolution, a readout and the approximation error at each interface.

The reviewed ToricGT implementation contributes graph tokens, endpoint incidence, finite Weyl operators, a mixture of ordinary and max-plus attention, and expert routing. The program and Tropical Quivers manuscripts add proposed algebraic and dynamical constraints. Their exact revisions and the corrections to selected statements are recorded in Chapter 29. We retain that pinned meaning of "ToricGT-style"; the name alone imposes neither a hydrodynamic limit nor a physical quantum representation.[142, 145, 146]

A concrete specification can be written

$$(Q, \mathcal{R}, \mathcal{A}, \mathcal{E}, \Phi, \mathcal{M}, \pi, \mathcal{P}).$$

Here Q is a finite quiver, $\mathcal{R}$ its intended relations, $\mathcal{A}$ the operator algebra, $\mathcal{E}$ an encoding, Φ an admissible evolution, $\mathcal{M}$ a measurement instrument, π a classical or coherent control law, and $\mathcal{P}$ a physical model with units, boundary conditions and parameter uncertainty. A compiler must say which of these structures it preserves. The fluid discretization and its continuum consistency belong to $\mathcal{P}$; the predictor's training objective belongs to a further statistical specification.

For example, a gentle-algebra relation declaring a composition of two arrows zero is meaningful for a representation by linear maps on vector spaces. A composition of trace-preserving channels on a nonzero state space cannot be the zero map: it still preserves trace. Thus a forbidden path may be enforced in a classical routing graph, or represented by trace-decreasing branches of an instrument, but cannot simply be imposed as a zero relation between unconditional trace-preserving channels. Relations must be checked in the category in which the proposed hardware actually operates.

30.2 A finite operator algebra with conservative quantum transport

30.2.1 Classical position and quantum internal state

Assign a finite-dimensional Hilbert space $\mathcal{H}_i$ to each vertex and set

$$\mathcal{A}_Q = \bigoplus_{i \in V} \mathcal{B}(\mathcal{H}_i).$$

A state is a tuple $\rho = (\rho_i)$ with $\rho_i \succeq 0$ and $\sum_i \operatorname{Tr} \rho_i = 1$. The mass at vertex i is $m_i = \operatorname{Tr} \rho_i$; when $m_i > 0$, ρ_i / m_i is its conditional internal state. The direct sum describes a classical location register with quantum internal degrees of freedom. Coherences between locations instead require the full algebra on $\bigoplus_i \mathcal{H}_i$, including off-diagonal blocks. This choice changes the state space and therefore changes the physics. Neither finite description is automatically a many-body quantum fluid: a tensor-product or Fock-space model has additional particle and correlation structure. An encoding into finitely many blocks must state which observables it retains and bound truncation or leakage effects.

Let $P_{ij} \geq 0$, $\sum_i P_{ij} = 1$, and let $\Phi_{ij} : \mathcal{B}(\mathcal{H}_j) \to \mathcal{B}(\mathcal{H}_i)$ be completely positive trace-preserving maps for each allowed transition. Define

$$(\mathcal{T}\rho)_i = \sum_j P_{ij} \Phi_{ij}(\rho_j). \tag{30.1}$$

This finite construction is related to open quantum walks, whose operator-valued transition rules combine quantum internal states with classical position records.[17]

Proposition 30.1 (Conservative transport on a direct-sum algebra)**.** *Under the preceding assumptions, $\mathcal{T}$ is completely positive and trace preserving. Its masses satisfy $m' = Pm$. For Hermitian inputs $X = (X_j)$,*

$$\sum_i \|(\mathcal{T}X)_i\|_1 \leq \sum_j \|X_j\|_1.$$

In particular the total trace distance between two states cannot increase.

Proof. Write $\Phi_{ij}(X) = \sum_\alpha K_{ij\alpha} X K_{ij\alpha}^*$ with $\sum_\alpha K_{ij\alpha}^* K_{ij\alpha} = I_{\mathcal{H}_j}$. Regard each $\sqrt{P_{ij}} K_{ij\alpha}$ as a block operator supported from j to i. Their Kraus sum defines a completely positive map, including after tensoring with any finite ancillary matrix algebra. On input block j, the sum of their adjoint products is $\sum_i P_{ij} I_{\mathcal{H}_j} = I_{\mathcal{H}_j}$, which proves trace preservation. Taking the trace of each output block gives $m' = Pm$.

For a Hermitian X_j, write its positive and negative parts as $X_j = X_j^+ - X_j^-$. Positivity and trace preservation imply

$$\|\Phi_{ij}(X_j)\|_1 \leq \operatorname{Tr} \Phi_{ij}(X_j^+) + \operatorname{Tr} \Phi_{ij}(X_j^-) = \|X_j\|_1.$$

The triangle inequality and column sums of P then give

$$\sum_i \Big\| \sum_j P_{ij} \Phi_{ij}(X_j) \Big\|_1 \le \sum_{i,j} P_{ij} \|X_j\|_1 = \sum_j \|X_j\|_1 .$$

□

Complete positivity is an algebraic constraint on a map and its extensions to entangled ancillary systems. A numerical test on a few positive input matrices does not establish it. A Kraus representation with a checked normalization, or an exactly positive Choi matrix with the correct partial trace, gives a finite certificate. The Heisenberg dual $\mathcal{T}^*$ acts linearly on observables and is unital. These statements concern the same channel in two dual descriptions.

30.2.2 The extra requirements imposed by a fluid interpretation

A local transition matrix depending on a prescribed velocity field can transport the masses m_i. Adding internal channels can represent collisions, relaxation or a controlled dissipative feature map. This is a finite transport model. To obtain a specified fluid equation one must also identify physical density, momentum, energy, length and time scales; prove the required conservation laws; specify constitutive assumptions; and control the discretization and kinetic limits. A positive mass-conserving map alone has no specified pressure law, viscosity or incompressibility. The reduction contracts in the preceding chapters state the corresponding closure and convergence obligations.

For a continuous-time finite quantum model, a standard admissible generator has the form

$$\dot\rho = -i[H,\rho] + \sum_\alpha \Big(L_\alpha \rho L_\alpha^* - \tfrac{1}{2}\{L_\alpha^* L_\alpha, \rho\} \Big), \qquad H = H^* .$$

With finite matrices and time-independent coefficients, this generates a completely positive trace-preserving semigroup.[85] Local choices of the L_α encode dissipative graph transitions; off-diagonal Hamiltonian couplings can support coherent transport. A forward Euler step need not itself be completely positive, even when the continuous generator is admissible. A numerical realization must retain that distinction or use a structure-preserving step.

Quantum fluids of light provide a particularly relevant physical family: interactions, coherent driving and loss lead to rich collective dynamics, with different descriptions appropriate to mean-field and correlated regimes. Their equations supply a possible physical model for coupled tokens; a generic optical neural-network nonlinearity does not automatically reproduce the correlated quantum regime.[5, 33]

30.3 Toric features, tropical routing and the information they discard

On $\mathbb{C}^d$, define $X|k\rangle = |k+1 \bmod d\rangle$ and $Z|k\rangle = \omega^k |k\rangle$, where $\omega = e^{2\pi i/d}$. Then $ZX = \omega XZ$. Products $Z^a X^b$ give finite Weyl features, while expectations $\mathrm{Tr}(O\rho)$ give real features for Hermitian O. A unitary change of basis transports both O and ρ and leaves their pairing unchanged. A learned interface may therefore use these finite operator features while retaining a classical graph controller. Additional claims about a toric variety, a torus action or a BGG complex require their own explicitly preserved structure.

The tropical part can enter as a routing optimizer or as a logarithmic envelope. For $\tau > 0$ and $m \geq 1$ real scores s_j, define

$$L_\tau(s) = \tau \log \sum_{j=1}^{m} e^{s_j/\tau}, \qquad p_j^\tau = \frac{e^{s_j/\tau}}{\sum_k e^{s_k/\tau}}.$$

Proposition 30.2 (Controlled tropical limit of a positive sum)**.** *For $M = \max_j s_j$,*

$$M \leq L_\tau(s) \leq M + \tau \log m.$$

If the maximizer j_* is unique and $s_{j_*} - s_j \geq \gamma > 0$ for all $j \neq j_*$, then

$$p_{j_*}^\tau \geq \frac{1}{1 + (m-1)e^{-\gamma/\tau}}.$$

Proof. Factor $e^{M/\tau}$ from the sum. The remaining sum is between one and m, giving the first bounds after taking its logarithm. Divide the normalizing sum for $p_{j_*}^\tau$ by $e^{s_{j_*}/\tau}$. Each other term is at most $e^{-\gamma/\tau}$, which gives the second bound. □

This statement concerns a positive exponential sum. With several maximizers the limiting probability is supported on that set; there is no uniquely selected route without an extra rule. The reviewed tropical-attention code uses scores that depend on dot products and applies a max-plus aggregation. Even when the aggregation is piecewise affine in its supplied scores, the complete map need not be piecewise affine in all original inputs.

For a product of nonnegative gain matrices, logarithms turn multiplication along one path into addition. Taking the largest path gives a max-plus envelope, with $\log 0 = -\infty$ and zero-weight paths absent. Ordinary matrix multiplication also sums over paths; its norm requires the path-multiplicity factor proved in Chapter 29. Signed or complex amplitudes add phase cancellation as another issue. A tropical magnitude envelope can upper-bound selected effects once these factors are supplied, but does not retain the interference pattern or determine entanglement.

Min-plus arithmetic also arises directly in control. A deterministic finite-horizon route problem has the Bellman recursion

$$V_n(i) = \min_{a \in \mathcal{U}(i)} \{c_n(i,a) + V_{n+1}(f(i,a))\}.$$

This is a precise optimization layer that can select graph operations. A stochastic, risk-sensitive or entropy-regularized controller has a different recursion, whose transition probabilities and normalization must be retained. A zero-temperature approximation can reduce exploration but also creates switching boundaries. The common gain certificate below is useful only if it covers every allowed branch and the effects of switching.

30.4 Nonlinear dynamics on an operator state space

30.4.1 Where the nonlinearity lives

An operator algebra is a linear space with a product and an adjoint. There are several mathematically distinct ways to obtain nonlinear dynamics associated with it. First, a classical controller z can select a channel or a generator:

$$\rho_{n+1} = \Phi_{\theta(z_n)}(\rho_n), \qquad z_{n+1} = f(z_n, y_n), \qquad y_n = \mathcal{M}_n(\rho_n).$$

For a fixed classical setting, Φ_θ is linear on density operators. The coupled controlled system can be nonlinear. If y_n denotes an ideal expectation rather than an actual measurement record, the displayed recursion is an ensemble-level model and needs a readout approximation before it describes hardware. For actual measurement records, the conditional measurement backaction must be included in the state update, or readout must use separately prepared copies. It cannot be omitted while retaining a claim about the same undisturbed quantum system.

Second, a measurement instrument has unnormalized branches $\rho \mapsto K_y \rho K_y^*$. Each is linear and completely positive. Conditioning on a nonzero-probability outcome divides by $\mathrm{Tr}(K_y \rho K_y^*)$ and produces a nonlinear normalized update. A physically complete account retains the probabilities and the classical outcome register. Replacing that instrument with a deterministic normalized map on an unknown state changes the operation.

Third, effective mean-field equations may contain a state-dependent Hamiltonian, for example $\dot\rho = -i[H(\rho), \rho]$, or a cubic order-parameter equation. Such equations approximate selected observables or many-body regimes under stated assumptions. Their nonlinearity does not by itself supply a fundamental nonlinear quantum channel. Driven polariton and atomic condensate models make this distinction experimentally relevant. [22, 33]

These three cases give a practical interpretation of a nonlinear operator perspective: learn and control dynamics *on* a physically admissible state space while retaining the linear completely positive operations appropriate to each fixed experimental setting. Positivity, trace, symmetry, causality and accessible controls become architectural constraints.

30.4.2 Koopman lifting and its closure error

For a deterministic state-space map $F : \mathcal{X} \to \mathcal{X}$, the Koopman operator is $(Uf)(x) = f(F(x))$. It is linear in the observable f even when F is nonlinear. On an algebra of ordinary scalar observables it also preserves products. For a stochastic transition, conditional expectation instead gives a positive unital operator, generally without product preservation. Neither construction is automatically the Heisenberg dual of a channel on the noncommutative algebra $\mathcal{A}_Q$: their domains and their notions of positivity must first be matched.

A finite dictionary $\psi(x)$ may satisfy

$$\psi(F(x)) = K\psi(x) + r(x).$$

An exact finite representation requires $r = 0$ on the claimed domain. Learning K from trajectories ordinarily leaves a residual, and a small training residual is not a uniform domain bound. Finite invariant subspaces that also reconstruct the state impose substantial restrictions, as the primary Koopman-control analysis explains.[26] The residual belongs in an error recurrence such as Theorem 29.16. This connects operator learning to a checked control budget without declaring a finite closure exact.

30.5 Control, observability and robustness

30.5.1 A common gain for coupled classical and quantum blocks

Let d_ρ measure quantum-state error, for example trace distance, and let d_z measure controller error. Suppose the actual closed-loop map satisfies on an invariant domain

$$\begin{pmatrix} d'_\rho \\ d'_z \end{pmatrix} \le G \begin{pmatrix} d_\rho \\ d_z \end{pmatrix} + \begin{pmatrix} \eta_\rho \\ \eta_z \end{pmatrix}, \qquad G = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \ge 0.$$

For $q > \max(a, d)$ and $b, c > 0$, the certificate $Gw \le qw$ with $w_\rho, w_z > 0$ is equivalent to

$$\frac{b}{q-a} \le \frac{w_\rho}{w_z} \le \frac{q-d}{c}.$$

Indeed the first row rearranges to the left inequality and the second to the right. Feasibility is therefore equivalent to $bc \le (q-a)(q-d)$ at that rate. If $a, d < 1$ and $bc < (1-a)(1-d)$, continuity supplies a feasible $q < 1$ sufficiently close to one. Cases with zero coupling can be checked directly using $Gw \le qw$. These formulas expose the cost of feedback: individually stable components can still have an excessive loop gain.

Once the gains are justified, a common w and q give the full finite-time budget of Theorem 29.16. An optimizer may propose w; the independent rational checker verifies the proposed inequalities. Separate certificates can cover model discrepancy, channel approximation, integration error and readout error as explicit residuals. For a uniform weighted perturbation δ and $q < 1$, the eventual bound is $\delta/(1-q)$. A smaller q improves this robustness bound, while rapid contraction also suppresses memory of past inputs. Reservoir design therefore needs a task-dependent tradeoff between memory, separation of inputs and noise sensitivity.[52]

30.5.2 Accessible controls and finite readout

For a closed finite quantum system, accessible Hamiltonians generate a Lie algebra that constrains the reachable unitaries. Full unitary control still preserves the eigenvalues of a density matrix. Changing its spectrum requires measurement, dissipation or an enlarged system and subsequent discarding. Thus the ability to fit a target matrix numerically does not prove that a specific device can implement it.[2] For an open device, the allowed jump operators, rates, locality, noise and bandwidth similarly constrain attainable channels.

Observability has a corresponding cost. One cannot obtain every expectation of an unknown state from one undisturbed specimen. A hardware proposal must specify repeated preparations, ensemble access, weak measurement or an instrument with a state estimator. A statistical error guarantee over a finite number of measurements can be converted into a deterministic residual bound on the event that those estimates are accurate. The event's confidence, the number of shots and the finite time horizon remain part of the result; an unlimited-horizon guarantee does not follow automatically.

For example, suppose M measurement settings are fixed in advance and each uses N independent shots with outcomes in $[-1, 1]$. Hoeffding's inequality and a union bound give simultaneous error at most ε with probability at least $1 - \delta$ when

$$N \ge \frac{2}{\varepsilon^2} \log \frac{2M}{\delta}.$$

For one setting, the two-sided failure probability is at most $2e^{-N\varepsilon^2/2}$; adding these M bounds proves the assertion. The confidence concerns sampling error for those fixed settings, not calibration bias or arbitrary temporally correlated reservoir outputs. Adaptive feedback requires a corresponding conditional sampling argument and a specified finite horizon.[65]

Model predictive control can optimize a finite horizon subject to channel admissibility, actuator limits and a certified invariant set. For a classical energy model

$$\dot{x} = (J(x) - R(x))\nabla E(x) + B(x)u, \qquad J^T = -J, \quad R \succeq 0,$$

the chain rule gives $\dot{E} = -\nabla E^T R \nabla E + y^T u$ with $y = B^T \nabla E$. This is a useful energy-balance constraint for a fluid, elastic or electrical surrogate. A quantum dissipative model needs its own

thermodynamic assumptions: an arbitrary driven channel need not decrease a chosen energy or increase entropy. The proposed controller must preserve the balance appropriate to its physical model.

30.6 An exact nonlinear quantum-state feedback example

The following small construction instantiates the control contract without assuming a gain bound from sampled derivatives. It uses a tuple of two qubit density matrices; it does not encode entanglement between the two blocks. Let $Z = \operatorname{diag}(1, -1)$,

$$z_j = \operatorname{Tr}(Z\rho_j), \quad B = \begin{pmatrix} 0 & 2 \\ 1/8 & 0 \end{pmatrix}, \quad b = \begin{pmatrix} 1/4 \\ 0 \end{pmatrix}, \quad u_i = \operatorname{clip}((Bz + b)_i, -1, 1).$$

Set $a = 1/4$, $\tau(u) = (I + uZ)/2$, and define

$$F_i(\rho) = a\rho_i + (1 - a)\tau(u_i). \tag{30.2}$$

The clipped controller is nonlinear but globally Lipschitz. Its expectations are assumed known exactly for the ideal model. A device implementation would estimate them through the readout process just described.

Proposition 30.3 (A physical channel family and a contracting feedback law)**.** *For every fixed* $u \in [-1, 1]$,

$$\Phi_u(X) = aX + (1 - a)\tau(u) \operatorname{Tr} X$$

is completely positive and trace preserving. The ideal feedback map F *preserves pairs of density matrices and satisfies, with* $D_i = \frac{1}{2}\|\rho_i - \sigma_i\|_1$,

$$D'_i \le \sum_j G_{ij} D_j, \qquad G = \begin{pmatrix} 1/4 & 3/2 \\ 3/32 & 1/4 \end{pmatrix}, \qquad G \begin{pmatrix} 4 \\ 1 \end{pmatrix} = \tfrac{5}{8} \begin{pmatrix} 4 \\ 1 \end{pmatrix}.$$

It therefore contracts the weighted maximum trace distance at rate $5/8$*. Its fixed point has Bloch vectors* $(0, 0, 1/3)$ *and* $(0, 0, 1/24)$*. The feedback map itself is not affine in the pair of input states.*

Proof. Since $|u| \le 1$, $\tau(u)$ is a density matrix. In input–output order, the unnormalized Choi matrix is

$$J(\Phi_u) = a|\Omega\rangle\langle\Omega| + (1 - a)I \otimes \tau(u), \qquad |\Omega\rangle = |00\rangle + |11\rangle.$$

Both terms are positive, and tracing over the output gives $aI + (1 - a)I = I$. Equivalently the replacement map has Kraus operators $\sqrt{t_k}|k\rangle\langle l|$, where t_k are the diagonal entries of $\tau(u)$ and $k, l \in \{0, 1\}$. Scaling these replacement operators by $\sqrt{1 - a}$ and adjoining $\sqrt{a}I$ gives the same channel. Convexity then proves state-space invariance for every chosen control.

Measurement duality gives $|\Delta z_j| \le 2D_j$. Clipping is one-Lipschitz and $\frac{1}{2}\|\tau(u) - \tau(v)\|_1 = |u - v|/2$. Thus

$$D'_i \le aD_i + \frac{1 - a}{2}|\Delta u_i| \le aD_i + (1 - a)\sum_j |B_{ij}| D_j.$$

Substitution gives G, and direct multiplication gives its certificate. The state space is complete in this finite-dimensional metric, so the strict contraction has at most one fixed point and converges to it. At a fixed point the transverse Bloch coordinates vanish; the displayed z values solve $z = Bz + b$ and lie inside the clipping interval. They therefore give existence as well as uniqueness.

For nonaffinity, hold the first input at $z_1 = 0$ and choose second inputs with $z_2 = 0$ and $z_2 = 1$. The first output z values are $3/16$ and $3/4$. Their average is $15/32$. At the midpoint input $z_2 = 1/2$, the clipped control is one and the first output is $3/4$. The difference is $9/32$, so a single affine channel cannot equal this ideal feedback map on all inputs. □

The fixed-control channels are physically admissible. The ideal feedback recursion describes what happens when a controller has the required expectations; it is not an instruction for obtaining them without readout cost. The associated exact implementation uses rational Bloch coordinates and rational Choi matrices. It retains a coherent initial state, tests an actually saturated step, and checks trace-distance bounds by squaring nonnegative quantities, avoiding floating-point square-root tolerances. These replacement channels are generally nonunital: $\Phi_u(I) = I + (1 - a)uZ$. The entropy-monotonicity statement for the earlier unital heat channel consequently does not transfer to this example.

There is also an exact tropical comparison. Weighted normalization gives

$$K = \operatorname{diag}(w)^{-1} G \operatorname{diag}(w) = \begin{pmatrix} 1/4 & 3/8 \\ 3/8 & 1/4 \end{pmatrix}.$$

Every row of K^n sums to $(5/8)^n$, whereas the largest weight of one length-n path is $(3/8)^n$. The latter is attained by alternating the vertices. The difference comes from adding the other paths, exactly as in the path-multiplicity analysis. The proposed rate $3/5$ fails the supplied weighted certificate and is retained as negative replay evidence.

Finally, if each measured expectation has deterministic error at most ε, the extra weighted one-step discrepancy is at most $3\varepsilon/16$: the two component bounds after division by w_i are $3\varepsilon/16$ and $3\varepsilon/64$. Therefore the same residual theorem gives the asymptotic envelope

$$\frac{3\varepsilon/16}{1 - 5/8} = \frac{\varepsilon}{2}.$$

For $\varepsilon = 1/100$, this is $1/200$. This is conditional error control under the stated readout bound. It is not an experimentally measured noise floor or a substitute for a finite-shot confidence analysis.

30.7 Candidate engineered and natural substrates

A substrate is relevant when its controlled degrees of freedom implement the specified encoding, operation and readout with bounded error. Its microscopic quantum origin alone does not establish useful coherent quantum processing. The following examples identify both the usable mechanism and the experiment needed to connect it to the proposed architecture.

Substrate	Candidate role	Evidence and realization test
Programmable photonic networks	Interferometers implement coherent mode mixing; phase shifters and detectors can supply a graph operator and readout. A separate controller can update routing.	The cited programmable-transport experiment used continuous-wave laser illumination and intensity readout. A few-photon channel would need its own nonclassical-state specification and validation, alongside loss and control-drift measurements. [63]

Substrate	Candidate role	Evidence and realization test
Exciton–polariton fluids	Coupled cavities or spatial modes supply driven, lossy, interacting fields. They are candidates for nonlinear reservoirs and graph-coupled fluid features.	Quantum-fluid and superfluid phenomena motivate the mechanism. A realization must identify its mean-field or correlated regime, its input/output map, and the controller's error bounds. [5, 33]
Ultracold atomic gases	Trapped modes or lattice sites provide quantum matter with controllable couplings and, in selected settings, engineered local dissipation.	Dissipative many-body dynamics have been studied experimentally. A ToricGT mapping would additionally need accessible graph couplings, repeated state preparation and an observable task with a calibrated readout budget.[22]
Superconducting cavity circuits	Microwave modes, nonlinear elements and engineered loss can implement bosonic channels, memory and controlled dissipation.	Engineered dissipation has confined dynamics to a quantum-state manifold. The proposed graph realization must quantify truncation, leakage, decoherence and the attainable channel family.[83]
Magnonic and magnetic networks	Spin-wave propagation and nonlinear responses can provide a compact dynamical reservoir with graph-like couplings.	Magnetic reservoir experiments support this computing mechanism. The usual wave-regime demonstration establishes a physical reservoir; quantum-state processing would need its own preparation and coherence evidence.[102]
Soft bodies and fluid–structure systems	Elastic motion, damping and hydrodynamic coupling can distribute nonlinear memory across a body, with sensors and actuation supplying the interface.	Soft-body reservoir experiments demonstrate task-relevant dynamics. Natural morphologies motivate graph constraints, but every generated body needs a measured transfer map and robustness test. [101]
Photosynthetic molecular networks	Excitonic couplings, vibrations and environmental relaxation offer a natural example of open-system energy transport.	The functional role and character of observed coherence require care. Electronic and vibrational contributions, disorder and open-system conditions must be distinguished before proposing an information processor.[47, 191]

Natural vascular, porous and transport networks provide another source of geometries and conservation constraints for a classical graph-fluid realization. Their branching and feedback suggest a design grammar; they do not supply a trained ToricGT model by identification alone. Likewise, a crystal or a biological assembly discovered in nature can inspire a coupling pattern without providing arbitrary programmable control. The useful question is which functions the accessible dynamics can implement and which invariants survive the readout.

There is no universal ordering of these substrates. Interferometric systems favor controlled linear mixing; strongly interacting quantum platforms favor selected many-body phenomena; dissipative mechanical and magnetic reservoirs may provide economical nonlinear memory. A comparison needs a common task, input bandwidth, output accuracy, reset protocol, energy accounting and training budget. Device operation, control and measurement must all enter the cost.

30.8 De novo substrate design with generative models

30.8.1 Generate a physical candidate for a specified operator

A generative model can propose a material or geometry conditioned on a target response. Inorganic crystal generation and generative photonic inverse design provide concrete examples of this approach. Their outputs are candidates for a physical validation pipeline; the generative likelihood is not an operator-realization certificate.[70, 190]

Let s describe composition, geometry, connectivity and fabrication parameters, and let u describe accessible controls. A forward model produces an effective evolution $\mathcal{F}_{s,u}$ and readout $\mathcal{M}_s$. One can seek

$$\min_{s,u,w,q} \quad L_{\text{task}}(s,u) + \lambda_E E(s,u) + \lambda_C C(s) + \lambda_R R_{\text{uncertainty}}(s,u)$$

subject to physical admissibility, an observable target interface, and $G(s,u)w \le qw$ with $w > 0$, $q < 1$ when contraction is appropriate. The objective terms need measured definitions: prediction error, energy per usable output, fabrication or control cost, and a robustness penalty. A generative prior can constrain plausible s without replacing the physical constraints.

The specification should begin with a target channel, scattering map, response kernel or controlled fluid update, rather than with a visual shape. For an optical device, relevant constraints include a realizable material response, loss, bandwidth and fabrication tolerances. For a magnetic, elastic or fluid device, stability, conservation laws, boundary conditions and sensor coupling matter. For a proposed quantum device, state preparation, available interactions and admissible measurement are part of the target. Complete positivity constrains a reduced channel; additional locality and controllability constraints determine whether the chosen substrate can realize it.

30.8.2 A falsifiable co-design loop

The following sequence separates design exploration from validation.

1. Specify the graph task, allowed algebraic relations, encoding, observables, error tolerance and resource budget. Decide whether the target requires coherent quantum states or only a wave/dynamical response.

2. Generate candidate structures conditioned on those requirements. Apply exact geometry, connectivity and composition constraints where available, and retain the generator version and random seed.

3. Evaluate candidates with the appropriate physical solver and uncertainty model. Reject violations of conservation, positivity, stability or attainable controls; a good surrogate score cannot override these checks.

4. Identify a reduced input–state–output model on a stated domain. Measure its residual on held-out conditions and obtain a justified gain envelope. If only sampled derivatives are available, record that empirical scope rather than promoting it to a uniform bound.

5. Co-design the controller and positive weights. Replay finite rational certificates independently, and propagate solver, readout and fabrication errors through the common-gain budget.

6. Evaluate prospective fabrication variation, then fabricate or assemble a selected candidate only under an actual experimental program. Calibrate it, test closed-loop behavior and compare with the predicted operator. A simulation-only result remains labeled as such.

7. Compare against conventional numerical solvers and simpler physical or classical learning baselines on the same task. Include state preparation, measurement, reset, calibration and controller costs.

MatterGen establishes that conditional generative modeling can propose inorganic crystals with target properties, including an experimentally tested candidate. It does not supply the mapping from a desired quantum channel to a fabricable quantum device.[190] Generative photonic inverse design directly targets an optical response, but its particular demonstrated response still needs an additional mapping to this graph-fluid architecture.[70] Those missing mappings are useful research questions: they specify what a new project must actually demonstrate.

Natural systems can enter this loop as a prior on topology, material motifs or dynamics. A learned grammar might favor sparse transport networks, hierarchical elastic couplings or molecular geometries. The design must then be evaluated with the same operator, control and readout contracts as a wholly synthetic candidate. This makes biomimetic inspiration testable without treating an analogy as a realization.

30.9 Connections to the wider research program

The construction offers a practical role for the project's order and surface-algebra language: encode legal compositions and symmetries, select observable features, and demand that a realization preserve specified relations. Quantum information supplies admissible state transformations and operational distances. Tropical mathematics supplies routing and logarithmic bounds. Control theory connects these finite operations to robust dynamical behavior. Each connection has an explicit map and an error obligation.

For Hilbert–Pólya research, a controllable finite operator family can test spectral statistics or evaluate a proposed discretization. A finite fit to zeros does not establish the required self-adjoint operator, spectral identification or analytic limiting theorem. For Yang–Mills, an engineered finite dissipative or Hamiltonian gap is a property of that system and has no automatic implication for the continuum mass-gap problem. For cryptography, a generated physical response may be a useful device primitive, but a security claim additionally requires a threat model, entropy or hardness assumptions, extraction and implementation analysis. The existing cryptographic implementations and their security scope remain separate evidence.

For Navier–Stokes, the useful next experiments are bounded: preservation of discrete invariants, closure-error prediction, controllable manufactured solutions and comparisons with conventional solvers at stated resolution. The externally released forced singularity theorem retains its own continuum hypotheses and independent-build status. The present finite control construction neither supplies a stronger continuum theorem nor removes the unresolved independent upstream build.

The immediate implementation milestone is the exact feedback example and its independent gain replay. The next research milestone is a trained graph-fluid surrogate with a documented invariant domain and held-out error budget. A physical milestone requires one specified substrate, calibrated encoding and readout, measured operator agreement and a fair resource comparison.

Chapter 31

Coverings, orders and learned quantum renormalization

31.1 Voltage covers preserve the specified local presentation

Let Q be a finite quiver and assign each arrow $a : u \to v$ a voltage $\nu(a)$ in the cyclic group $\mathbb{Z}/d\mathbb{Z}$. The lift has vertices (u, s) and arrows

$$(a, s) : (u, s) \longrightarrow (v, s + \nu(a)).$$

Translation of every sheet index by the same element is a deck action. A path lifts uniquely after its initial sheet is supplied. A monomial zero relation lifts to the corresponding zero relation on each lifted path.

Proposition 31.1 (Finite gentle cover)**.** *If a finite-dimensional quadratic monomial bound quiver is gentle, the voltage lift with the lifted monomial relations is gentle and finite-dimensional.*

Proof. At each lifted vertex, incoming and outgoing arrows correspond bijectively to those at the base vertex. Thus the degree-at-most-two conditions are unchanged. For each lifted arrow, allowed and forbidden successor and predecessor choices correspond bijectively to the base choices, so their uniqueness conditions persist. Any allowed lifted path projects to an allowed base path of the same length. Since the base's allowed-path lengths are bounded, the lifted allowed-path lengths are bounded as well. For a finite quadratic monomial presentation this gives finite dimensionality. □

The implementation checks every voltage assignment in selected small covers, not only one favorable drawing. It also constructs explicit unramified rectangular-torus cell covers, checks both cellular coboundary pullback identities and verifies $d_1 d_2 = 0$. The degree-d torus cover still has genus one. This is consistent with multiplicativity of Euler characteristic for an unramified cover, not evidence that an arbitrary quiver cover automatically extends to a marked, branched or holomorphic surface cover.

For a surface with boundary, marked points or ramification, the covering data must include those structures. The gentle algebras associated with unpunctured marked surfaces in [G1] have specific hypotheses. A general categorical covering functor [G2], a voltage cover of a presentation, and a topological covering of a dessin surface are related constructions but not interchangeable definitions.

31.2 Integral relations can obstruct a special-fiber cover

The rank-four order of the preceding volume has generators satisfying

$$uv = \pi e_1, \qquad vu = \pi e_2.$$

A lift that assigns independent sheet shifts to u and v makes their composite act by the total sheet shift. Over a torsion-free module with nonzero π, preserving the displayed relation requires that total shift to be the identity. In the explicit matrix realization the residual is

$$\pi(S_{\nu(u)+\nu(v)} - I).$$

Reduction modulo π turns the relation into a zero relation. The special fiber therefore admits covers that fail the original integral relation. The code checks this obstruction at degree two: voltages $(1, 0)$ fail at $\pi = 3$ but pass after specialization to zero; voltages $(1, 1)$ preserve the integral relation. In this two-arrow example, zero cycle voltage also makes the resulting cover split into components. A connected nontrivial special-fiber cover is not thereby a connected cover of the same integral order.

More general order coverings require a declared grading, central-parameter action, local lattice data and compatibility with all nonmonomial relations. The present obstruction gives an exact test to apply before extending a neural or categorical construction over the generic fiber. Forgetting the deformation parameter can genuinely remove information needed to validate the lift.

31.3 Neural pullback, descent and the cost of arbitrary identifiers

For a d-sheet finite cover let U pull a base feature to a constant feature on each fiber. Over $\mathbb{R}$, put $P = d^{-1}U^{\mathsf{T}}$. Then

$$PU = I, \qquad (UP)^2 = UP.$$

The second map projects onto fiber-constant signals. In characteristic dividing d, this average is not defined by division; replacing it by an unnormalized sum can even send the constant-one feature to zero. The implementation retains that counterexample.

Proposition 31.2 (Local layer transport)**.** *Suppose a base message-passing update uses a shared message function, a shared node update, and a sum or normalized local mean over incoming arrows. On pullback features, the same update on the cover commutes with U.*

Proof. Every lifted vertex has one incoming lift of each incoming base arrow. The message arguments are equal because the source and target features are pulled back. Therefore the local sum is the base sum; the local degree is also the same, so normalized means agree. Applying the common node function preserves equality. Induction gives the result for every finite composition of such layers. □

There is a related global-attention identity. If each key/value is replicated exactly d times, then each softmax numerator and its denominator receive the same factor d, leaving the weighted value unchanged. Repetition also leaves a coordinatewise maximum unchanged. Consequently normalized softmax and max-plus operations commute with the declared uniform replication. This statement requires transported or tied features and parameters. Arbitrary absolute sheet identifiers break it, as the negative test demonstrates. Global attention on arbitrary non-fiber-constant states is not covered by the identity.

The normalized linear embedding $J = U/\sqrt{d}$ is an isometry. It is a direct-sum encoding of amplitudes, not a map that clones an unknown quantum state into d independent tensor copies. That

distinction preserves the no-cloning constraints already proved in the CCC chapter. A quantum circuit descends through a cover only after its state embedding and operator intertwining are specified.

Covering alone also does not preserve a positive lower spectral gap. A connected d-fold cover of a finite cycle is a longer cycle; its smallest positive graph-Laplacian eigenvalue is $4\sin^2(\pi/(dN))$, which tends to zero. Base eigenvectors lift, but new non-fiber-constant modes appear. Thus cover degree is a geometric limit requiring its own estimates, not a substitute for the infinite-width limit of a neural parameterization.

31.4 A learned hierarchy of physical quantum channels

The implemented network coarse grains four qubits to two and then one. Each stage applies explicitly specified two-qubit unitary blocks and traces out the declared environment qubits. The first stage shares three rotation parameters across its two pairs, and the second stage uses another three. A scalar readout adds slope and bias; it is a classical postprocessor, not part of the quantum channel.

For a stage unitary W and an environment basis $\{|e\rangle\}$, the Kraus operators are

$$K_e = (I \otimes \langle e|)W.$$

Their completeness follows directly:

$$\sum_e K_e^* K_e = W^* \left(I \otimes \sum_e |e\rangle\langle e| \right) W = I.$$

If L_f describes the next complete stage, the combined operators are $L_f K_e$ and

$$\sum_{e,f} (L_f K_e)^* (L_f K_e) = \sum_e K_e^* \left(\sum_f L_f^* L_f \right) K_e = I.$$

The general rectangular-matrix identity is checked in Lean. An independent numerical implementation constructs the Kraus operators, while the training path uses a Heisenberg observable and another reference performs explicit partial traces. Their agreement checks both wire order and discarded-subsystem semantics.

A discarded qubit need not be physically erased during a simulation. If no later gate acts on it, postponing its trace to the end gives the same retained density matrix. This allows the complete four-qubit unitary circuit to be replayed independently on CPU and cuStateVec. It does not postselect favorable outcomes, and it does not replace a channel by a nonlinear normalization of a selected branch.

31.5 What the learned renormalization map actually learns

Training uses exact Gibbs states of a four-site periodic transverse-field Ising Hamiltonian at inverse temperature one. The 33 coupling values are split by a declared index rule into train, validation and test. The objective retains the mean nearest-neighbor ZZ observable at the final readout. The six gate angles and two readout parameters are actually optimized; validation selects the checkpoint and test couplings do not select it.

Seed	Selected step	Test MSE	Constant baseline
17	120	0.0000615	0.0260293
29	120	0.0009693	0.0260293

This is a learned coarse-graining channel with a measurable target, not an assertion that every relevant operator is preserved. The complete first-stage two-qubit density matrix is retained, and its effective Hamiltonian is reconstructed as

$$H_{\text{eff}} = -\log \rho_{\text{coarse}} + cI.$$

Every finite-temperature input is positive definite. Unitary conjugation preserves that property, and a partial trace of a positive-definite operator is positive definite: for nonzero v, sum the strictly positive quadratic forms on $v \otimes e$. Thus the matrix logarithm is defined for these exact states.

The implementation expands H_{eff} in all sixteen two-qubit Pauli operators. Projecting it back to only IX, XI and ZZ is a separate model reduction. The trace distance between the actual coarse state and the Gibbs state of that truncated Hamiltonian is reported at every coupling. Nonzero generated interactions are not discarded while claiming exact closure within the original Ising family.

A useful stability statement follows from the integral formula for the matrix logarithm. If positive definite A, B have minimum eigenvalue at least $\mu > 0$, the resolvent identity gives

$$\|\log A - \log B\| \le \|A - B\| \int_0^\infty (\mu + t)^{-2}\, dt = \frac{\|A - B\|}{\mu}.$$

A learned coarse state close to a singular state can therefore have an unstable effective Hamiltonian even when its density-matrix error is small. The recorded minimum eigenvalues expose this condition instead of hiding it in a renormalization plot. An infinite-system fixed point, scaling exponent or universal phase transition needs a system-size analysis beyond these four-site experiments.

31.6 An explicit infinite-width theorem

The width study uses independent one-qubit feature circuits, not the number of sheets in a cover. Let $w_r \sim N(0, I_d)$ and $b_r \sim \text{Unif}[0, 2\pi]$ independently, and define

$$\varphi_r(x) = \sqrt{2}\, \langle Z \rangle_{R_Y(w_r \cdot x + b_r)|0\rangle} = \sqrt{2} \cos(w_r \cdot x + b_r).$$

Set $K_M(x, y) = M^{-1} \sum_{r=1}^{M} \varphi_r(x)\varphi_r(y)$.

Theorem 31.3 (Independent quantum features and their limiting kernel)**.** *For fixed x, y,*

$$K_M(x, y) \longrightarrow K(x, y) = \exp(-\|x - y\|^2/2)$$

almost surely. On a fixed set of n inputs,

$$\Pr\left(\max_{i,j} |K_M(x_i, x_j) - K(x_i, x_j)| > \epsilon \right) \le 2n^2 e^{-M\epsilon^2/8}.$$

Every empirical kernel is positive semidefinite.

Proof. Averaging over b gives $\mathbb{E}_b[2\cos(a+b)\cos(c+b)] = \cos(a-c)$. The characteristic function of a standard Gaussian then gives the displayed K. The strong law applies to the independent bounded products. Each product lies in $[-2, 2]$, so Hoeffding's inequality gives $2e^{-M\epsilon^2/8}$ for a fixed pair. A union bound over the n^2 pairs proves the finite-set estimate. Positive semidefiniteness follows because K_M is a feature Gram matrix. □

The empirical study uses widths 16, 64, 256, 1,024 and 4,096 and four fixed seeds, preserving all errors rather than requiring every realization to improve monotonically. The random-feature circuit is classically tractable; the example establishes a controlled quantum-feature limit, not a computational speedup.

Readout-only kernel ridge regression has a corresponding stability estimate. If $\|K_M - K\|_2 \le \delta$ and regularization is λI with $\lambda > 0$, the resolvent identity gives

$$\|(K_M + \lambda I)^{-1} y - (K + \lambda I)^{-1} y\|_2 \le \delta \|y\|_2 / \lambda^2.$$

This is a specific trained readout regime. It does not prove that training every gate of an unrestricted deep quantum network leaves the kernel fixed.

The trained-network Gaussian-process results in [Q2] are relevant to broader architectures, but their initialization, light-cone dependence, normalization, limiting kernel and learning-rate assumptions must travel with the theorem. A deep hierarchy with expanding causal cones, shared parameters or max-plus branches does not inherit the result merely by increasing a register count. Cover limits, feature-width limits, physical thermodynamic limits and optimizer-time limits remain four distinct operations.

31.7 Cover symmetry and the dependence structure of wide quantum networks

Exact symmetry and independent random features need not be compatible under an arbitrary parameter-sharing scheme. If every sheet uses the same circuit parameters, repeating a feature does not create independent samples. In particular, for one random feature φ, the empirical expression $M^{-1} \sum_{r=1}^{M} \varphi(x)\varphi(y)$ is just $\varphi(x)\varphi(y)$ for every M. It does not converge to the averaged RBF kernel by increasing the number of identical copies. The additional executable control retains this rank-one counterexample alongside a genuinely nested sequence of independent features.

A useful general criterion can nevertheless be stated in terms of light-cone dependence. Consider centered random circuit readouts $X_{M,1}, \dots, X_{M,M}$, uniformly bounded by a constant, with a dependency graph of maximum degree D_M. This means that two collections with no edge between them are independent. Suppose $M^{-1} \operatorname{Var}(\sum_r X_{M,r}) \to \sigma^2$ and $D_M = o(M^{1/4})$.

Proposition 31.4 (A sufficient Gaussian finite-dimensional limit)**.** *Under these assumptions, $M^{-1/2} \sum_r X_{M,r}$ converges in distribution to $N(0, \sigma^2)$. The same assertion holds jointly on any fixed finite input set when the bound and dependency graph apply to the corresponding readout vectors and their normalized covariance matrices converge.*

Proof. For each fixed order k, the cumulant of the sum is a sum of joint cumulants. A joint cumulant vanishes when its variables split into independent groups. Thus only tuples connected in the dependency graph, allowing repeated indices, contribute. There are at most $C_k M (D_M + 1)^{k-1}$ such tuples. Boundedness gives a uniform fixed-order cumulant bound by the moment–cumulant formula. After normalization, the order-k bound is therefore $C'_k M^{1-k/2} (D_M + 1)^{k-1}$.

For $k = 3$ this tends to zero by $D_M = o(M^{1/4})$; for every fixed $k > 3$ it also tends to zero. The first cumulant is zero and the second converges by assumption. The moment–cumulant relations therefore give convergence of all moments to the Gaussian moments. The Gaussian distribution is moment-determinate, so the method of moments gives convergence in distribution. For a finite vector of inputs, apply the same argument to each fixed linear combination and use the Cramér–Wold criterion. □

For randomly initialized circuits, disjoint parameter light cones can supply the required dependency graph. A global shared parameter or a growing cone can violate the bound. Training all gates with a common data-dependent loss can introduce new dependencies, so the initialization theorem is not automatically a theorem about the trained network. The trained-QNN results in [Q2] address this transition under their additional kernel, light-cone and scaling assumptions.

One symmetry-compatible construction uses independent feature channels, each of which is averaged over the same finite deck group. Every channel is then exactly equivariant or invariant as specified, while different channels remain independent. The width limit is taken over those independently initialized channels on a fixed covered graph. Increasing the covering degree is a different limit and can change both the dependency bound and the spectrum. This provides an explicit design rule rather than conflating replication, geometric lifting and statistical width.

31.8 Testing whether gate learning improves the representation

A low error relative to a constant predictor does not by itself show that optimizing quantum gates was necessary. The additional control freezes each seeded initial channel, fits only its scalar affine readout on training couplings, and evaluates the same held-out couplings. It reports the training design's condition number, the frozen-channel/readout error, the raw initialized error and the fully learned channel's error. No learned weights are retrained or selected using this comparison.

The full density matrix and the declared ZZ observable also supply an exact target oracle. Its zero error is a statement about access to the uncompressed input, not an equally constrained learned baseline. Together these comparisons distinguish three effects: information already present in the input, information retained by the channel, and information made accessible by its trained readout. The publication retains a failure to improve over a frozen channel whenever it occurs.

Seed	Learned gates MSE	Frozen gates + readout	Constant MSE
17	0.0000615	0.0001169	0.0260293
29	0.0009693	0.0000056	0.0260293

Part VII

Conformal geometry, information and language

Chapter 32

Conformal aeons, elliptic integrals and identifiable cosmic information

Evidence status. This chapter separates proved mathematical statements, conditional physical models and observational claims. Eight central papers on conformal cyclic cosmology (CCC), its proposed signatures, null tests and communication proposal were read in full; selected sections of three larger references were also read. Exact versions, file hashes and coverage are recorded in `extension/review/cosmology/`. The finite synthetic experiment below was executed. It contains no observed CMB data. The theorems in this chapter are ordinary mathematical proofs; they are not represented as a completed Lean formalization of cosmology or of the observational pipeline.

32.1 Four questions that require different evidence

Conformal geometry asks which structures survive a change of metric by a positive scalar factor. A cosmological model must additionally specify field equations, matter content and matching conditions. A proposed observable must specify how those fields reach a detector. A communication claim must further specify how a sender's choice changes the distribution of detector outcomes. These questions form a useful hierarchy: solving an earlier one supplies ingredients for a later one, but does not solve it automatically.

CCC proposes that the conformal future boundary of one expanding cosmic aeon can be joined to the conformally regularized beginning of another. Its motivation involves the eventual disappearance or irrelevance of massive clocks, and a special treatment of the initial gravitational degrees of freedom. The framework is developed in several mathematically distinct ansatzes; choosing the conformal factor and compatible matter evolution is substantive physical work, as the explicit comparisons in Markwell and Stevens [94, Secs. 2–7] demonstrate. We use a bridge metric to state exactly what a conformal argument does establish.

The observational history also has distinct layers. Gurzadyan and Penrose proposed low-variance concentric CMB rings as traces of activity in a preceding aeon [60, pp. 1–6]. An and collaborators subsequently reported a statistical excess of temperature-gradient features interpreted as Hawking points [6, Secs. 3–5]. Independent analyses have challenged the significance of these particular statistics: correlated Gaussian skies reproduce low-variance features, and searches over multiple scales require a global comparison [100, pp. 1–4] [74, Secs. 2–4]. Bodnia and collaborators' later neural-network and ring searches found no significant evidence for the templates they tested after examining unusually bright or dark pixels [24, Secs. 2–4]. None of these logical statements establishes or excludes every possible CCC realization.

An explicit cross-aeon SETI proposal exists. Gurzadyan and Penrose discuss electromagnetic signaling, possible other carriers and an intentionally transmitted record [61, Sec. 2, pp. 2–4]. This is a proposal to investigate a channel, not a report of a decoded message. The distinction is central to the mathematics below: a natural correlation may carry information about an earlier physical state without carrying an intentionally selected symbol.

32.2 Conformal geometry with the singular boundary kept explicit

Let (U, g) be a smooth oriented Lorentzian four-manifold, with signature $(-,+,+,+)$, and let $\widetilde{g} = e^{2f} g$ for a smooth real function f. All statements in this section hold on U, where the factor is finite and strictly positive. They do not assert regularity at a zero or pole of that factor. Our sign convention differs from some of the cited spinor texts; the scaling identities below are independent of that choice.

Proposition 32.1 (Connection change and null trajectories)**.** *The Levi–Civita connections satisfy*

$$\widetilde{\nabla}_X Y = \nabla_X Y + X(f)Y + Y(f)X - g(X,Y)\,\mathrm{grad}_g\, f. \tag{32.1}$$

The metrics have the same null cones and the same unparameterized null geodesics. If $k = dx/d\lambda$ is a g-affinely parameterized null geodesic, then $\widetilde{k} = e^{-2f}k$ is affinely parameterized for $\widetilde{g}$; equivalently, $d\widetilde{\lambda}/d\lambda = e^{2f}$ along the curve.

Proof. Define $D_X Y$ by the right side of (32.1). The added tensor is symmetric in X, Y, so D is torsion-free. Expanding $X(e^{2f} g(Y,Z))$ and using $\nabla g = 0$ shows that the extra terms $2X(f)g(Y,Z)$ cancel precisely the four gradient terms contributed by $g(D_X Y, Z) + g(Y, D_X Z)$. Hence $D\widetilde{g} = 0$. Uniqueness of the torsion-free metric connection proves the formula. Also $\widetilde{g}(v,v) = e^{2f} g(v,v)$, so the null cones agree. If $\nabla_k k = 0$ and $g(k,k) = 0$, the formula gives $\widetilde{\nabla}_k k = 2k(f)k$. Multiplication by e^{-2f} cancels this acceleration:

$$\widetilde{\nabla}_{e^{-2f}k}(e^{-2f}k) = e^{-4f}\big(\widetilde{\nabla}_k k - 2k(f)k\big) = 0.$$

The parameter relation follows from the chain rule. Applying the same argument with $-f$ proves the converse. □

Proposition 32.2 (Hodge scaling and the four-dimensional Maxwell system)**.** *On an oriented pseudo-Riemannian n-manifold, the Hodge operators on p-forms obey*

$$\widetilde{*}\alpha = e^{(n-2p)f} * \alpha. \tag{32.2}$$

Consequently, in dimension four, the source-free equations $dF = 0$, $d * F = 0$ for a two-form F are conformally invariant if F itself is kept unchanged.

Proof. The inner product on covariant p-tensors scales by e^{-2pf} and the oriented volume form by e^{nf}. For all p-forms α, β,

$$\alpha \wedge \widetilde{*}\beta = \langle \alpha, \beta \rangle_{\widetilde{g}}\, \mathrm{vol}_{\widetilde{g}} = e^{(n-2p)f} \alpha \wedge *\beta.$$

Nondegeneracy of the wedge pairing gives (32.2). For $n = 4$, $p = 2$, its exponent is zero. Exterior differentiation is independent of the metric, proving both Maxwell equations unchanged. □

These statements explain why null propagation and conformally invariant fields are natural ingredients in CCC. They do not include charged plasma, absorption, scattering, a singular source or a quantum trace anomaly. Each requires an additional equation or boundary condition. Tod's account carefully distinguishes conformal field weights from physically measured energy [166, Secs. 1.2, 1.3.1 and 3.6].

For example, with our signature the Maxwell stress tensor is

$$T_{ab} = F_{ac}F_b{}^c - \tfrac{1}{4}g_{ab}F_{cd}F^{cd}.$$

Raising one index contributes e^{-2f}, and raising both contributes e^{-4f}. Thus $\widetilde{T}_{ab} = e^{-2f}T_{ab}$. If u is a unit timelike observer, $\widetilde{u} = e^{-f}u$ is unit for the new metric, so

$$\widetilde{T}_{ab}\widetilde{u}^a\widetilde{u}^b = e^{-4f}T_{ab}u^a u^b.$$

Conformal preservation of the field equation therefore coexists with rescaling of a detector's energy density. Preserving an unparameterized light ray does not preserve its clock, frequency, flux or accessible information budget.

For a CCC bridge, write

$$\widehat{g} = \Omega^2 g, \qquad \check{g} = \omega^2 g, \qquad \Omega\omega = -1, \qquad \Sigma = \{\omega = 0\}. \tag{32.3}$$

The signed factors are conventional; away from Σ their absolute values determine positive rescalings. In this convention the preceding aeon's factor $|\Omega|$ diverges toward crossover, whereas the succeeding aeon's factor ω tends to zero. The regularity of g, the extension of suitably rescaled fields, and their compatibility with matter equations are separate hypotheses at Σ. Neither preceding proposition can be applied there simply by substituting $f = \log 0$.

32.3 A genuine elliptic-curve bridge in the Friedmann equation

There is a precise connection to the earlier elliptic-curve chapters which does not require a cosmic signal. Consider a Friedmann–Lemaître–Robertson–Walker scale factor $a(t) > 0$, constant spatial-curvature parameter κ, a cosmological constant Λ, radiation and pressureless matter. Assume the usual separate conservation laws, so that their energy densities are $\rho_{r,0}a^{-4}$ and $\rho_{m,0}a^{-3}$. In units with $c = 1$, the Friedmann equation is

$$H^2 = \frac{8\pi G}{3}(\rho_{r,0}a^{-4} + \rho_{m,0}a^{-3}) + \frac{\Lambda}{3} - \frac{\kappa}{a^2}, \qquad H = \dot{a}/a.$$

Define conformal time by $dt = a\,d\eta$. Since $a' = a\dot{a} = a^2H$,

$$(a')^2 = P(a), \qquad P(a) = \alpha + \beta a - \kappa a^2 + \delta a^4, \quad \alpha = \frac{8\pi G\rho_{r,0}}{3}, \quad \beta = \frac{8\pi G\rho_{m,0}}{3}, \quad \delta = \frac{\Lambda}{3}. \tag{32.4}$$

This reduction is valid on intervals where the stated matter laws hold. An additional epoch with different energy transfer need not satisfy the same quartic. A Weierstrass representation of such a quartic appears in the 2025 CCC proposal, with this limitation explicitly stated [96, Sec. V, especially p. 13].

Theorem 32.3 (Quartic compactification and conformal time)**.** *Suppose $\delta \neq 0$ and $P \in \mathbb{C}[a]$ has four distinct roots. The smooth compactification E of $y^2 = P(a)$ is a connected genus-one curve. After choosing any point as origin it is an elliptic curve over $\mathbb{C}$. The form da/y extends to an*

everywhere nonzero holomorphic differential on E. On a real expanding branch with $P(a) > 0$, conformal time satisfies

$$\eta - \eta_0 = \int_{a_0}^{a} \frac{ds}{\sqrt{P(s)}}.$$

Thus the cosmological quadrature is locally an Abelian integral on an elliptic curve under the stated hypotheses.

Proof. The projection $(a, y) \mapsto a$ is a two-sheeted cover away from the four roots. At a root r, $P'(r) \neq 0$, so the implicit function theorem makes y a local coordinate and gives $a = r + y^2/P'(r) + O(y^4)$. Each root is a simple branch point. At infinity put $z = 1/a$ and $w = y/a^2$. The equation becomes

$$w^2 = \delta - \kappa z^2 + \beta z^3 + \alpha z^4.$$

At $z = 0$ there are two distinct points $w = \pm\sqrt{\delta}$, each smooth and unbranched. Adding them compactifies the cover. The curve is connected because a rational square cannot have the four simple zeros of P; equivalently, analytic continuation around one branch point exchanges the two sheets.

For completeness, triangulate the base sphere with all four branch points among its vertices. Each face and each open edge has two lifts. An ordinary vertex has two lifts, while a branch vertex has one. Therefore $\chi(E) = 2\chi(S^2) - 4 = 0$. For a connected compact orientable surface, $\chi(E) = 2 - 2g$, so $g = 1$. Choosing a complex point supplies the marked point in the definition of an elliptic curve.

At a finite unbranched point, da/y is regular and nonzero. At a branch point the displayed local expansion gives $da/y = (2/P'(r) + O(y^2))\, dy$, again regular and nonzero. At either point at infinity it becomes $-dz/w$, with $w \neq 0$. These charts cover E. Finally, on a chosen expanding real branch (32.4) gives $d\eta = da/\sqrt{P(a)}$, and integration proves the formula. Changing paths on E can add periods; this is the usual multivaluedness of an Abelian integral, not an ambiguity in the chosen real initial-value solution. □

One can see the Weierstrass equation directly. Choose a root r and put $u = (a - r)^{-1}$, $v = y/(a - r)^2$. Taylor expansion, which is exact for a quartic, gives

$$v^2 = Au^3 + Bu^2 + Cu + D, \quad A = P'(r),\ B = \tfrac{1}{2}P''(r),\ C = \tfrac{1}{6}P'''(r),\ D = \delta.$$

Since $A \neq 0$, the substitutions $X = Au + B/3$, $Y = Av$ give

$$Y^2 = X^3 + \left(AC - \frac{B^2}{3}\right)X + \left(\frac{2B^3}{27} - \frac{ABC}{3} + A^2D\right), \qquad \frac{da}{y} = -\frac{dX}{Y}. \tag{32.5}$$

The substitutions are birational with explicit inverses on the displayed charts. Smooth projective models agree after completion. Repeated roots invalidate the four-branch-point argument: the original affine equation is singular and its normalization can have genus zero. Vanishing δ likewise requires a separate degree analysis. One must not declare every special cosmological quadrature a nonsingular elliptic curve without checking its discriminant. The substitution is defined over $\mathbb{C}$, or over a coefficient-field extension containing the chosen root r. It need not be defined over the original coefficient field, and a real quartic need not have a real root with which to make this particular change of variables.

This bridge is algebraic and analytic. Elliptic-curve cryptography instead uses finite groups such as $E(\mathbb{F}_q)$ and a discrete-logarithm problem. Cosmological density estimates are not automatically exact number-field coefficients, and the integral above is not a discrete logarithm. If one deliberately chooses algebraic coefficients, one can investigate good reduction at suitable primes; security then requires its own group-order and parameter analysis. The shared cubic equation supplies mathematical language, not a cryptographic attack or a cosmic encryption mechanism.

32.4 A spinor contraction that gives a conditional conserved charge

Tensor diagrams are most useful when their wire types and equations are specified. In a flat four-dimensional spin domain, let S be a two-dimensional complex spinor space with nondegenerate alternating form ϵ_{AB}. Primed indices denote the conjugate spin space. Indices are raised with the inverse alternating form, not with a positive Hermitian inner product. Work with a smooth symmetric field $\psi_{ABCD} = \psi_{(ABCD)}$ and a symmetric field κ^{CD}.

Proposition 32.4 (Spin lowering and Stokes invariance)**.** *Suppose the fields satisfy*

$$\partial^{AA'}\psi_{ABCD} = 0, \qquad \partial^{A'(A}\kappa^{CD)} = 0. \tag{32.6}$$

Then $\phi_{AB} = \psi_{ABCD}\kappa^{CD}$ satisfies $\partial^{AA'}\phi_{AB} = 0$. Under the usual spinor identification of a symmetric two-spinor with a complex self-dual two-form F, the form F is closed. If two oriented compact two-cycles S_0, S_1 bound an oriented oriented piecewise-smooth three-chain W in a region where these hypotheses and smoothness hold, then

$$\int_{S_1} F = \int_{S_0} F. \tag{32.7}$$

Proof. The product rule gives

$$\partial^{AA'}\phi_{AB} = (\partial^{AA'}\psi_{ABCD})\kappa^{CD} + \psi_{ABCD}\partial^{AA'}\kappa^{CD}.$$

The first term vanishes by hypothesis. In the second, complete symmetry of ψ allows us to replace the derivative of κ by its symmetrization in A, C, D. This is zero by the second equation of (32.6).

More explicitly, the corresponding complex two-form has components $F_{AA'BB'} = \phi_{AB}\epsilon_{A'B'}$. Its divergence vanishes by the equation just proved. Fix the orientation convention in which this spinor summand has $*F = iF$; choosing the opposite convention replaces i by $-i$ without changing the argument. The vanishing divergence is equivalent to $d * F = 0$. Since $*F = iF$, it follows that $dF = 0$. Stokes' theorem then gives $\int_{S_1} F - \int_{S_0} F = \int_W dF = 0$. □

The proof is a small, exact example of a diagrammatic calculation. A symmetric four-leg tensor contracts with a symmetric two-leg tensor; the derivative term annihilates the relevant symmetric component. The proof uses the displayed field equations, not the shape of a diagram. In particular, a spinor contraction is not automatically a trace-preserving quantum channel. The alternating tensor ϵ, the Hermitian dagger of a finite quantum system and the Lorentzian metric are different structures even when their index diagrams look similar.

Meissner and Penrose use this type of spin-lowering construction in their 2025 development of a gravitational-wave epoch and a mass-charge proposal [96, Secs. II–III]. Their argument begins with a weak field on a conformally flat background. On p. 10, its continuation into the early succeeding aeon invokes local isotropy to argue for no net outward charge flux. A regular closed-form calculation is therefore a well-defined mathematical component; identifying its normalization with physical mass across a nonlinear, matter-filled crossover needs the additional hypotheses. Equation (32.7) alone proves neither a general cosmological existence theorem nor an observed hot spot.

The same source's proposed angular scale is a radius 0.03–0.04 radians in its equation (99), corresponding to a diameter about 3.4–4.6 degrees. Its extra conformal interval is inferred using the asserted spot scale [96, Sec. VI.B]. A parameter adjusted to a disputed feature cannot also count as an independent confirmation of that feature. A new epoch is scientifically testable when it predicts a joint observable distribution whose additional consequences can be checked.

32.5 From a possible carrier to an identifiable message

A channel is a probability law, not merely a geometric path. Let M be a finite message variable, X a source configuration, Y detector data and Z the output of an inference procedure. A possible causal model is

$$M \longrightarrow X \longrightarrow Y \longrightarrow Z. \tag{32.8}$$

Writing the chain is an assumption about conditional independence. It does not establish that a sender existed. For example, if θ describes uncertain transfer physics, an explicit model might be

$$p(y \mid m) = \int p(y \mid x, \theta)\, p(x \mid m, \theta)\, p(\theta \mid m)\, dx\, d\theta.$$

Replacing $p(\theta \mid m)$ by $p(\theta)$ is an additional independence assumption. A physical proposal must explain the source alphabet, its energy and timing constraints, its propagation and scattering, the detector response, and how alternative messages induce distinct laws.

For the SETI proposal, the conformal Maxwell calculation supplies a possible mathematical ingredient. It does not calculate an end-to-end coherent channel through the intervening cosmology. The sender, encoding, survival of distinguishable states and decoder must be specified separately. Long-wavelength transparency and sufficiently large signal power are physical requirements, not consequences of Proposition 32.2. Conversely, a measured primordial correlation can constrain a source distribution even when no sender or semantic code is posited.

Theorem 32.5 (Classical data processing)**.** *For finite random variables satisfying $M \to Y \to Z$,*

$$I(M; Z) \le I(M; Y).$$

The statement also holds conditionally on an independently supplied training record D when the Markov property holds given D.

Proof. Use logarithms to base two. For distributions p, q on a finite set,

$$D(p\|q) = \sum_x p(x) \log_2 \frac{p(x)}{q(x)} \ge 0.$$

Applying $\ln u \le u - 1$ to $u = q(x)/p(x)$ on the positive support of p proves this, with infinite divergence allowed when $q = 0 < p$. Conditional mutual information is an average of such divergences and is therefore nonnegative. Expanding the finite entropy definitions gives the two chain rules

$$I(M; Y, Z) = I(M; Y) + I(M; Z \mid Y) = I(M; Z) + I(M; Y \mid Z).$$

The Markov property makes $I(M; Z \mid Y) = 0$, so subtraction proves the inequality. Apply the same argument to each conditional law given $D = d$ and average to obtain its conditional version. □

These definitions follow the information-theoretic formulation introduced by Shannon [150]; the proof above states the precise finite version used here. A decoder may make information computationally usable even though it cannot increase mutual information. Learning a good approximation to a likelihood ratio can improve practical inference over a poor hand-designed statistic. It cannot distinguish messages that the physical channel maps to exactly the same distribution.

Proposition 32.6 (Identifiability and its quantitative obstruction)**.** *Let P, Q be two possible laws of observed data on a finite set. For any randomized test with rejection probability $h(y) \in [0,1]$,*

$$\left|\sum_y h(y)P(y) - \sum_y h(y)Q(y)\right| \le \|P-Q\|_{\mathrm{TV}} := \tfrac{1}{2}\sum_y |P(y)-Q(y)|.$$

Stochastic post-processing cannot increase total variation. If every message has the same data law, every decoder has identical output laws under those messages; for K equally likely messages its success probability is at most $1/K$.

Proof. Write $v = P - Q = v^+ - v^-$ coordinatewise. Its zero total sum gives $\sum v^+ = \sum v^- = \|P-Q\|_{\mathrm{TV}}$. Since $0 \le h \le 1$, the scalar $\sum hv$ lies between the negative and positive of that quantity. For a stochastic kernel $K(z \mid y)$,

$$\sum_z \left|\sum_y K(z \mid y)v(y)\right| \le \sum_{z,y} K(z \mid y)|v(y)| = \sum_y |v(y)|,$$

proving contraction. If $P(y \mid m)$ is independent of m, then so is $P(z \mid m)$. With a uniform prior, any output-based guess has success $K^{-1}\sum_m P(\mathrm{guess} = m) = 1/K$ when it always guesses a member of the alphabet, and no larger when abstention is allowed. □

This obstruction applies equally to classical and quantum-assisted learned post-processing of a classical sky map. A powerful model may compress noise into an eloquent explanation; that output is not evidence for a sender unless the competing source laws actually differ. If one chooses an arbitrary decoder after observing a finite pattern, a lookup table can assign that pattern to any desired text. Such a construction has no independently tested predictive content. A credible coding hypothesis instead fixes constraints that predict additional symbols, redundancy or independent observations before they are examined.

Proposition 32.7 (Finite records and a decoding requirement)**.** *If a recorded observation consists of N samples with at most b bits each, then $I(M;Y) \le H(Y) \le Nb$. If a uniformly distributed message $M \in \{1,\dots,K\}$, $K \ge 2$, is decoded from Y with error probability P_e, then*

$$I(M;Y) \ge \log_2 K - h_2(P_e) - P_e \log_2(K-1),$$

where $h_2(p) = -p\log_2 p - (1-p)\log_2(1-p)$ with the continuous endpoint convention.

Proof. The record has at most 2^{Nb} possible values. Nonnegativity of divergence from the uniform law gives $H(Y) \le \log_2(2^{Nb})$; the identity $I(M;Y) = H(Y) - H(Y \mid M)$ proves the first bound. For a deterministic decoder $\widehat{M}(Y)$, let $E = \mathbf{1}_{\{M \ne \widehat{M}\}}$. Since E is a function of (M,Y),

$$H(M \mid Y) = H(E \mid Y) + H(M \mid E, Y).$$

The first term is at most $H(E) = h_2(P_e)$. When $E = 0$, M is known from Y; when $E = 1$, it has at most $K-1$ possibilities. The second term is at most $P_e \log_2(K-1)$. Subtract from $H(M) = \log_2 K$. A randomized decoder can be treated by including its independent random seed with Y, which leaves $I(M;Y)$ unchanged. □

An ideal continuum field can be described by infinitely many real parameters. That mathematical description does not assert infinitely many distinguishable outcomes at finite precision, energy or observing time. Proposition 32.7 applies to the actual recorded alphabet regardless of how small a conformal diagram makes its source region look. A meaningful information-density claim must state the resource, observable alphabet and error criterion.

32.6 Quantum constraints on a proposed communication channel

Let ρ, σ be density matrices on a finite-dimensional Hilbert space. Their trace distance is $D(\rho,\sigma) = \frac{1}{2}\|\rho-\sigma\|_1$. Physical deterministic quantum channels are completely positive and trace-preserving; the next proof uses only positivity and trace preservation on the relevant matrix algebra.

Proposition 32.8 (Trace-distance contraction)**.** *For a positive trace-preserving linear map Φ, $D(\Phi(\rho),\Phi(\sigma)) \le D(\rho,\sigma)$. For every measurement event represented by $0 \le E \le I$, $|\operatorname{tr} E(\rho-\sigma)| \le D(\rho,\sigma)$.*

Proof. The Hermitian matrix $X = \rho - \sigma$ has a Jordan decomposition $X = X_+ - X_-$ with $X_\pm \ge 0$ and $\|X\|_1 = \operatorname{tr} X_+ + \operatorname{tr} X_-$. Positivity, the triangle inequality and trace preservation imply

$$\|\Phi(X)\|_1 \le \|\Phi(X_+)\|_1 + \|\Phi(X_-)\|_1 = \operatorname{tr} X_+ + \operatorname{tr} X_-.$$

Also $\operatorname{tr} X = 0$, so $\operatorname{tr} X_+ = \operatorname{tr} X_- = D(\rho,\sigma)$. The inequalities $0 \le E \le I$ bound $\operatorname{tr} EX$ between $-\operatorname{tr} X_-$ and $\operatorname{tr} X_+$. □

Corollary 32.9 (No deterministic cloning of nonorthogonal pure states)**.** *A quantum channel cannot map both $|\psi\rangle\langle\psi|$ and $|\phi\rangle\langle\phi|$ to their respective tensor squares when $0 < |\langle\psi,\phi\rangle| < 1$.*

Proof. Put $c = |\langle\psi,\phi\rangle|$. In the two-dimensional span of the states, choose phases and a basis with $|\phi\rangle = c|\psi\rangle + \sqrt{1-c^2}|e\rangle$. The difference of their rank-one projectors has trace zero and eigenvalues $\pm\sqrt{1-c^2}$, so their trace distance is $\sqrt{1-c^2}$. The tensor-square states have overlap c^2 and distance $\sqrt{1-c^4} > \sqrt{1-c^2}$, contradicting Proposition 32.8. □

This is the precise obstruction underlying the no-cloning result [178]. Orthogonal classical labels can be copied; the theorem concerns unknown nonorthogonal states. It does not prevent ordinary redundancy in an encoded classical message.

Proposition 32.10 (No signalling by a local unconditioned operation)**.** *Let ρ_{AB} be a joint state. A trace-preserving operation on A with Kraus operators $K_j : A \to A'$ leaves the reduced state on B unchanged after its unreported outcome is discarded.*

Proof. For every observable O_B,

$$\begin{aligned}&\operatorname{tr}\left[(I_{A'}\otimes O_B)\sum_j (K_j\otimes I_B)\rho_{AB}(K_j^\dagger\otimes I_B)\right]\\ &\quad= \operatorname{tr}\left[\rho_{AB}\left(\sum_j K_j^\dagger K_j\otimes O_B\right)\right] = \operatorname{tr}[\rho_{AB}(I_A\otimes O_B)].\end{aligned}$$

Equality against every O_B determines equality of the reduced density matrices. Conditioning on a particular outcome changes the formula, but identifying that subensemble requires an outcome record to reach the observer. □

Entanglement therefore does not replace the propagation law in (32.8). The finite-dimensional theorem is not itself a construction of quantum fields on a crossover; it specifies what an ordinary local quantum-information model would permit.

There is also a simple exact dimension bound. If K states in a d-dimensional Hilbert space are perfectly distinguished by a POVM $\{E_j\}$, then $K \le d$. Indeed, $\operatorname{tr} E_j\rho_j = 1$ makes E_j act as the

identity on the support of ρ_j, while $\operatorname{tr} E_j\rho_i = 0$ makes it vanish on the support of ρ_i for $i \neq j$. To justify these implications, for positive A, B the equality $\operatorname{tr} AB = 0$ gives $\|A^{1/2}B^{1/2}\|_{\mathrm{HS}}^2 = 0$. Apply this to $A = I - E_j, B = \rho_j$, and to $A = E_j, B = \rho_i$. The respective supports lie in orthogonal eigenspaces of E_j. Their nonzero dimensions sum to at most d. This perfect-decoding bound does not by itself give a noisy channel capacity, but it excludes an unbounded perfectly readable alphabet in a fixed finite system.

Finally, discarding degrees of freedom has no universal direction of entropy change. A pure Bell state has zero joint von Neumann entropy but either marginal has one bit. Conversely, $\rho_{AB} = I_A/d \otimes |0\rangle\langle 0|_B$ has entropy $\log_2 d$, while its B marginal has zero entropy. Unitary conjugation preserves the eigenvalues and hence entropy. A proposed cosmological entropy reset must specify a state, algebra of observables and dynamical map; neither taking a partial trace nor conformally changing the metric is alone an entropy theorem of the required kind.

32.7 Ring statistics, matched filters and the cost of searching

An observational claim needs a statistic and a probability law for that statistic under a specified null. Let $x \in \mathbb{R}^n$ be a vector of map samples with a centered Gaussian null $N(0, C)$, where C is positive definite. In an actual analysis C must reflect the chosen signal spectrum, beam, mask and instrumental noise, with foreground residuals treated explicitly. A white-noise null and a correlated sky generally answer different questions.

Proposition 32.11 (Quadratic-form moments)**.** *For a real symmetric matrix A and $x \sim N(\mu, C)$,*

$$\mathbb{E}[x^{\mathsf{T}}Ax] = \operatorname{tr}(AC) + \mu^{\mathsf{T}}A\mu,$$
$$\operatorname{Var}(x^{\mathsf{T}}Ax) = 2\operatorname{tr}((AC)^2) + 4\mu^{\mathsf{T}}ACA\mu.$$

Proof. Write $x = \mu + C^{1/2}z$ for $z \sim N(0, I)$ and diagonalize the symmetric matrix $C^{1/2}AC^{1/2} = U\operatorname{diag}(\lambda_i)U^{\mathsf{T}}$. The coordinates of $U^{\mathsf{T}}z$ are independent standard Gaussians. The quadratic term arising from z is therefore $\sum_i \lambda_i Z_i^2$, whose mean is $\sum_i \lambda_i = \operatorname{tr}(AC)$ and variance is $2\sum_i \lambda_i^2 = 2\operatorname{tr}((AC)^2)$. The remaining random term is $2\mu^{\mathsf{T}}AC^{1/2}z$, of variance $4\mu^{\mathsf{T}}ACA\mu$. Its covariance with the quadratic term is zero because all centered Gaussian third moments vanish. The constant term is $\mu^{\mathsf{T}}A\mu$. □

For nonnegative annular weights w_i with $\sum_i w_i = 1$, the weighted variance is

$$V = \sum_i w_i x_i^2 - \left(\sum_i w_i x_i\right)^2 = x^{\mathsf{T}}(\operatorname{diag}(w) - ww^{\mathsf{T}})x.$$

Its null expectation depends on all covariances between the samples, not only on their individual variances. A localized nonzero mean also alters both moments. This explains why a few unusual pixels and the choice of reference baseline can change a low-variance search. The published null analyses examine precisely such effects [100, pp. 2–4] [24, Secs. 2–3]. Multiplying nominal ring probabilities is justified only when the required independence or conditional factorization has been established; concentric annuli in a correlated sky do not acquire independence from their appearance.

Theorem 32.12 (A fixed-template matched filter)**.** *Let $t \neq 0$ be a fixed template and suppose $x \sim N(at, C)$ with known positive definite C. The maximum-likelihood amplitude and its variance are*

$$\widehat{a} = \frac{t^{\mathsf{T}}C^{-1}x}{t^{\mathsf{T}}C^{-1}t}, \qquad \operatorname{Var}(\widehat{a}) = \frac{1}{t^{\mathsf{T}}C^{-1}t}. \tag{32.9}$$

The normalized score $R_t = t^{\mathsf{T}} C^{-1} x / \sqrt{t^{\mathsf{T}} C^{-1} t}$ *has distribution* $N(a\sqrt{t^{\mathsf{T}} C^{-1} t}, 1)$.

Proof. Apart from a constant, the negative log likelihood is $\frac{1}{2}(x - at)^{\mathsf{T}} C^{-1}(x - at)$. Its derivative in a vanishes at (32.9), and its second derivative $t^{\mathsf{T}} C^{-1} t$ is strictly positive. A linear functional of a Gaussian vector is Gaussian. Its mean and variance are obtained by substitution, using $C^{-1} C C^{-1} = C^{-1}$. □

For one template selected before seeing the data, the one-sided null probability is $1 - \Phi(R_t)$. If a search scans a family $\mathcal{T} = \{t_1, \dots, t_m\}$ and reports its largest score, the statistic is instead

$$S(x) = \max_{1 \le j \le m} R_{t_j}(x). \tag{32.10}$$

The selected template's normal tail is not the null tail of S. For valid individual p-values, the union bound gives $\Pr(\min_j p_j \le \alpha/m) \le \alpha$ without independence. If they actually were independent uniforms, the probability of at least one value below u would be $1 - (1 - u)^m$. Correlated templates require their joint law or a valid full-scan simulation, not a guessed number of independent trials.

Jow and Scott explicitly compared the scan across their scale grid in each simulated sky, obtaining a probability to exceed of 13% for their global statistic [74, Sec. 3, pp. 6–7]. This number is conditional on that pipeline and null model. It is not a posterior probability that CCC is true or false. Their result illustrates why a locally striking feature can become ordinary once the selection operation is included.

Theorem 32.13 (Finite Monte Carlo rank calibration)**.** *Let* $S_0, S_1, \dots, S_B$ *be exchangeable null statistics, with* S_0 *the observed statistic and* $S_1, \dots, S_B$ *simulation statistics. Then*

$$p_{\mathrm{MC}} = \frac{1 + \#\{b \in \{1, \dots, B\} : S_b \ge S_0\}}{B + 1} \tag{32.11}$$

satisfies $\Pr(p_{\mathrm{MC}} \le \alpha) \le \alpha$ *for* $0 \le \alpha \le 1$. *With no ties, its values are uniform on* $\{1/(B+1), \dots, 1\}$.

Proof. With no ties, exchangeability makes the position of S_0 in descending order uniform among $B + 1$ positions. Formula (32.11) is that position divided by $B + 1$. At most $\lfloor \alpha(B + 1) \rfloor$ positions satisfy the inequality, proving the bound. With ties, attach independent continuous tie-breaking variables. The randomized descending rank is uniform by the same argument. Counting every statistic tied with S_0 as at least as large gives the largest rank within its tied block, which is no smaller than the randomized rank. The reported p-value is therefore conservative. □

Exchangeability is the essential hypothesis. Every simulation must run the same location search, scale search, masks, feature selection and reported maximum as the observed map. A trained model may be held fixed if it was trained independently of all compared maps. If the observed map guided training, threshold selection or architecture selection, those choices belong in the selection procedure. Fitting nuisance parameters on the observed data and simulating from the fitted value does not automatically produce exchangeability. Re-fitting them in each parametric bootstrap realization can be necessary, yet the resulting calibration is generally approximate unless an additional exact argument applies.

There is a related geometric trap. On a spherical coordinate patch the round metric is $d\theta^2 + \sin^2\theta\, d\phi^2$. A proposed sky twist $T_S(\theta, \phi) = (\theta, \phi + S\theta)$ has pullback

$$T_S^* g = (1 + S^2 \sin^2\theta) d\theta^2 + 2S \sin^2\theta\, d\theta\, d\phi + \sin^2\theta\, d\phi^2.$$

For $S \neq 0$ away from the poles this is not a scalar multiple of g, so the twist is not conformal for the round metric. More fundamentally, even a genuine non-isometric conformal map need not

preserve an isotropic Gaussian covariance $C(n \cdot n')$. Distorting the data can change correlations and thus the detection statistic under an ordinary null. A transformed comparison requires the same transformation of properly generated null maps and a recalibrated complete procedure.

32.8 What a neural search and a new cosmic observable can establish

Bodnia and collaborators' HawkingNet search uses a ResNet18-based classifier trained on injected profiles, and investigates simulation-to-data mismatch and unusually bright or dark spots [24, Sec. 4 and App. A]. Its injection study uses amplitudes 100–400 μK and a quoted angular-width parameter 0.7–1.2 degrees. A model validated on this family is sensitive to that family; it is not automatically a test of every profile obtainable from a revised crossover model. Equally, changing a profile after a negative test creates a new testable hypothesis rather than erasing the previous result.

A network score is a statistic. It is a likelihood ratio only if the training distributions and calibration justify that interpretation. A large score on a foreground residual or a simulation mismatch may show that the classifier has learned a discriminative feature unrelated to the proposed cosmology. Precision requires separate tests of injection recovery, false-positive rate under realistic nulls, nuisance robustness and generalization to held-out realizations. The comparison should include ordinary matched filters and other classical methods with comparable data and tuning budgets.

Repeated observations of the same sky have a useful but limited role. If $Y_1 = S + N_1$ and $Y_2 = S + N_2$ with independent zero-mean instrumental noises independent of the sky, their covariance is

$$\operatorname{Cov}\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} = \begin{pmatrix} C_S + C_{N_1} & C_S \\ C_S & C_S + C_{N_2} \end{pmatrix}. \tag{32.12}$$

The off-diagonal blocks are not zero. Agreement between satellites can disfavor an instrumental artifact, but the same chance cosmic feature will appear in both. Counting them as independent draws of the cosmic sky would double-count that part of the evidence.

Proposition 32.14 (The elementary full-sky cosmic-variance calculation)**.** *In an isotropic centered real Gaussian sky, use a real orthonormal spherical-harmonic basis at fixed ℓ. Its $d_\ell = 2\ell + 1$ coefficients $a_{\ell m}$ are independent $N(0, C_\ell)$. The estimator*

$$\widehat{C}_\ell = \frac{1}{2\ell + 1} \sum_m a_{\ell m}^2$$

has mean C_ℓ and variance $2C_\ell^2/(2\ell + 1)$.

Proof. Isotropy makes the covariance scalar on the degree-ℓ harmonic space; Gaussianity then turns diagonal covariance into independence. Each squared coefficient has mean C_ℓ and variance $2C_\ell^2$. Sum these independent variables and divide by d_ℓ^2. □

More pixels do not create additional independent low-ℓ modes. Noise reduction, sky coverage and additional observables can nevertheless improve a particular test. Masks and foregrounds change the practical covariance; the simple proposition is not a formula for every cut-sky estimator. Planck's isotropy and statistics analysis discusses the different roles of temperature, polarization and residual systematics [122, Secs. 1–2 and 8].

Polarization supplies additional information only through a specified joint prediction. For jointly centered Gaussian vectors T, E with invertible C_{TT},

$$E \mid T \sim N\left(C_{ET}C_{TT}^{-1}T, C_{EE} - C_{ET}C_{TT}^{-1}C_{TE}\right). \tag{32.13}$$

To prove it, set $R = E - C_{ET}C_{TT}^{-1}T$. Direct multiplication gives $\mathrm{Cov}(R, T) = 0$ and the displayed covariance for R. Joint Gaussianity makes R independent of T, yielding the formula. A temperature-selected feature should therefore be tested against a conditional polarization law, rather than declaring its polarization an independent arbitrary anomaly. A model with no predicted polarization profile cannot obtain a precise exclusion from the absence of an unspecified profile.

The polarization field also has a type that a graph model should respect. Under a local basis rotation by χ, one convention is $Q + iU \mapsto e^{-2i\chi}(Q + iU)$. A learned operation that treats this as a basis-independent scalar has silently discarded a symmetry requirement. Scalar temperature features, spin-two polarization features, frequency channels and masks should have distinct declared transformation rules. The typed tensor framework elsewhere in this monograph supplies a language for enforcing such rules; it does not replace a physical transfer calculation.

Proposition 32.15 (A forecast with linear nuisance templates)**.** *Suppose $x \sim N(at + B\theta, C)$, with known $C > 0$, and the columns of B linearly independent. After fitting the nuisance parameter θ, the information for a is*

$$F_a = t^{\mathsf{T}}C^{-1}t - t^{\mathsf{T}}C^{-1}B(B^{\mathsf{T}}C^{-1}B)^{-1}B^{\mathsf{T}}C^{-1}t. \tag{32.14}$$

It is nonnegative and vanishes exactly when t lies in the column space of B. When $F_a > 0$, the variance of the fitted amplitude is $1/F_a$.

Proof. Whiten the problem: $z = C^{-1/2}x$, $v = C^{-1/2}t$ and $W = C^{-1/2}B$. Least squares over θ removes the orthogonal projection $P_W = W(W^{\mathsf{T}}W)^{-1}W^{\mathsf{T}}$. The remaining model has template $(I - P_W)v$ and independent unit covariance on the orthogonal complement of $\operatorname{im} W$. Its squared template norm is $v^{\mathsf{T}}(I - P_W)v$, which expands to (32.14). It vanishes precisely when $v \in \operatorname{im} W$, equivalently $t \in \operatorname{im} B$. The matched-filter calculation on the complement gives variance $1/F_a$. □

This elementary Schur-complement calculation says exactly what a new measurement must improve: the distinguishable part of the signal after nuisance fitting. Historical CMB forecasts provide examples of combining instrument noise, foreground uncertainty and external data [153, Sec. 2.6]. Their numerical forecasts cannot be repurposed as a CCC sensitivity forecast without a CCC template and a joint likelihood. We make no claim here about an instrument's present operating status. The same principle applies to galaxy distributions, spectral distortions or gravitational-wave observations: a new catalog is useful when its selection function and its predicted covariance with the proposed source are specified.

32.9 A reproducible synthetic search, with its actual result

The repository contains an executed demonstration in `extension/review/cosmology/synthetic_search.py`, with output `synthetic-results.json`. It uses no observed sky and no fitted ΛCDM spectrum. Its purpose is to make the difference between a fixed-template probability and a full-search probability directly inspectable.

There are 96 deterministic Fibonacci-sphere directions n_i. The Gaussian covariance is

$$C_{ij} = \exp\left(\frac{n_i \cdot n_j - 1}{\ell_c^2}\right) + \sigma^2\delta_{ij}, \qquad \ell_c = 0.35, \quad \sigma = 0.20. \tag{32.15}$$

It is strictly positive definite. Indeed, $e^{\kappa n_i \cdot n_j} = \sum_{r\geq 0} \kappa^r \langle n_i^{\otimes r}, n_j^{\otimes r}\rangle / r!$ is a convergent sum of Gram matrices with nonnegative coefficients. Multiplication by $e^{-\kappa}$ preserves positive semidefiniteness and adding $\sigma^2 I$ gives the lower eigenvalue bound σ^2.

For each direction and each width $w \in \{0.18, 0.30, 0.45\}$ radians, the template has entries $e^{-\theta_i^2/(2w^2)}$ minus their discrete mean, where θ_i is angular distance from the chosen center. All 288 templates are included in the maximum (32.10). A fixed NumPy PCG64 seed 20261113 generates one base sky, 4095 independent calibration skies and 2048 further validation skies. The declared injected template is center 0, width 0.30, and the same base sky receives normalized matched-filter signal strengths $s = 0, 1, 3, 5$:

$$x_s = x_0 + s\frac{t}{\sqrt{t^{\mathsf{T}} C^{-1} t}}.$$

This normalization increases the score of that fixed template by exactly s. No seed was searched for an attractive result.

Injected s	Scan maximum	Selected local tail	Global rank p
0	2.717146	0.00329237	0.46215820
1	3.087799	0.00100823	0.18090820
3	4.110988	0.0000196985	0.00341797
5	6.041391	7.63956×10^{-10}	0.00024414

The unperturbed synthetic sky illustrates the selection effect: its best local normal tail is about 0.0033, whereas almost half of the calibrated null scans are at least as extreme. The strongest injection has zero calibration exceedances. Its rank value is therefore 1/4096, the finite resolution of this experiment; it is not a claim to have measured a smaller tail probability. The two intermediate rows also show that the winning template can differ from the injected one because templates and noise are correlated.

At nominal threshold 0.05, the independent finite validation sample has 119 rejections among 2048 skies, a fraction 0.0581055. This is a diagnostic conditional on one calibration set, not a proof that the conditional rejection rate equals 0.05. The exact rank theorem concerns the joint null distribution of the observation and calibration draws. The implementation checks the normalized filter variances and the injected-score identity to errors below 4×10^{-15} and 2×10^{-15} respectively. Its covariance minimum eigenvalue is approximately 0.0528566, above the proved 0.04 lower bound. Software versions, source hash, ring-moment diagnostics and actual elapsed time are recorded with the result.

This experiment supports only the mathematical and computational claims just specified. It does not recover the published CCC statistics, train HawkingNet, implement a Boltzmann transfer code, or analyze Planck data. The companion visualization uses the same declared template bank and calibration; rotating its sphere changes the view, not the data or the statistic.

32.10 A testable program for learned tensor and graph models

A cohesive research program can now be stated without conflating its parts. First specify a physical source and crossover model sufficiently well to generate a joint distribution of observable fields. The bridge metric, conformal field weights and spinor identities constrain this model, but do not determine all of its transfer parameters. The Friedmann elliptic curve can solve a background quadrature on intervals where its matter assumptions hold; it does not provide the missing perturbation transfer function.

Next discretize the observable field with explicit types. A spherical graph may carry scalar temperature, spin-two polarization and frequency metadata. Tensor contractions should respect the corresponding basis changes. A quantum implementation must additionally specify a state preparation map, finite Hilbert dimensions, allowed channels, measured observables and the number

of shots. The same diagram can describe different mathematics unless its types, contractions and normalizations are fixed. Algebraic rewrite correctness is distinct from the empirical adequacy of the physical distribution on which the diagram operates.

Then compare inference procedures under identical information and resource constraints. Injection families should span physically justified parameter uncertainty. Null families should include the relevant ordinary cosmology, instrumental uncertainty and foreground alternatives. Freeze the complete statistic before evaluation, or reproduce its adaptive selection in the calibration. Report detection power as a function of signal and nuisance parameters, together with held-out false-positive rates and ordinary classical baselines. If a quantum learned graph model improves sample efficiency or computational cost, that is a separate measurable result; a visually complex quantum circuit does not establish either improvement.

Finally, a semantic-message claim needs a distinct comparison beyond physical anomaly detection. A natural prior-aeon source law and an intentional encoding law must make different predictions for independent data. Redundancy, a constrained code family or repeated independently predicted symbols could in principle make that comparison meaningful. An unconstrained translation chosen after seeing the sky cannot. This is the content of Proposition 32.6, not a restriction on which cosmological hypotheses may be explored.

Problem 32.16 (An end-to-end cross-aeon inference problem)**.** *Specify a family of regularized crossover solutions, source states and detector kernels for which the induced laws $P_\theta(Y \mid M)$ are well-defined. Establish their identifiability after the declared nuisance parameters are fitted. Derive a joint temperature, polarization or other observable prediction and calibrate a fixed statistic against a realistic null. Only after these steps, compare an intentional code hypothesis with natural source alternatives and quantify its decoding error.*

The problem connects conformal geometry, elliptic curves, complexes of field equations, tensor contraction and information theory in a precise way. Each contributes a defined mathematical object or a constraint. The missing physical transfer law and the empirical comparison remain explicit research tasks rather than being supplied by analogy.

Chapter 33

CCC geometry and a real-map search with frozen learned statistics

33.1 The original mathematical interfaces are retained

The original CCC chapter separates conformal geometry, elliptic estimates, quantum-information constraints and observational calibration. This continuation strengthens those interfaces instead of replacing them by a generic neural classifier. The formal additions check the middle-form conformal exponent and the explicit polynomial transformations used by the elliptic-curve discussion. They do not turn an algebraic identity into an observational confirmation of a crossover.

For a metric rescaling in dimension n, the Hodge-star scaling on k-forms has exponent $n-2k$. In four dimensions at degree two this exponent is zero. The new finite identity checks that exponent's specialization. The full Hodge-star transformation law, metric signatures and Maxwell/source assumptions remain the mathematical premises stated in the original chapter.

Let $P(x) = ax^4+bx^3+cx^2+dx+e$ and let r be a root. Substituting $x = r+1/u$ and multiplying by u^4 produces

$$u^4P(r+1/u) = Au^3 + Bu^2 + Cu + D,$$

where

$$A = P'(r), \quad B = P''(r)/2, \quad C = P'''(r)/6, \quad D = P''''(r)/24.$$

For $v^2 = Au^3 + Bu^2 + Cu + D$, put $X = Au + B/3$ and $Y = Av$. Then

$$Y^2 = X^3 + (AC - B^2/3)X + 2B^3/27 - ABC/3 + A^2D.$$

Lean checks the coefficient expansions and transport identity. A birational interpretation additionally requires the stated nonzero denominators and a simple root when A is used to invert the coordinate change. Nonsingularity is a separate discriminant condition. The same algebraic shape across finite-field arithmetic and conformal-time equations is useful structure, not an identification of their physical interpretations.

33.2 Actual datasets and resolution-aware tests

The study uses three actual WMAP nine-year products: ILC and foreground-reduced V- and W-band temperature maps, plus the KQ85 mask. All are acquired through links on official NASA LAMBDA product pages. The acquisition record binds filenames, source URLs, exact bytes, FITS fields, native Nside512 nested ordering and temperature units. Millikelvin are explicitly converted

to microkelvin; the ILC header's `mK, thermodynamic` wording is retained rather than rejected or silently normalized in the source record.

The ILC documentation states reliability above approximately ten degrees and cautions about uncertainty at smaller angular scales [C1]. Accordingly this implementation tests broad Gaussian widths of 12, 18 and 24 degrees. It is not presented as a reproduction of a degree-scale Hawking-point search. Twenty-four fixed high-latitude centers and three scales give 72 templates.

Processing is fixed: pixel-average to Nside16, estimate harmonics through $\ell = 32$, remove monopole and dipole, apply an additional five-degree Gaussian beam, and reconstruct the map. Retained pixels require degraded KQ85 fraction at least 0.8 and absolute Galactic latitude at least twenty degrees. The original map beams differ; adding the same smoothing does not constitute exact common-beam deconvolution. The processing record retains that fact.

Three statistics are scanned. The first compares a Gaussian-weighted temperature mean with the retained-sky mean. The second detects low variance in the outer radial annulus. The third applies a learned four-qubit hierarchy to four mean-subtracted radial summaries. The corresponding local graph is an ordered four-band relation, not a privileged target-derived topology.

33.3 Learned statistics selected before observation scoring

The channel architecture is the same explicit 4-to-2-to-1 construction used in the preceding chapter. Its gate angles and scalar readout are trained on independently generated phase-surrogate skies and injected broad profiles. The observed map's original phases do not select the weights, center grid, scale set, mask, amplitude grid or checkpoint.

There are 1,024 training examples and 256 held-out validation examples, with Gaussian-profile amplitudes 64, 128 and 256 microkelvin. A bounded arctangent encoding maps four radial features to rotation angles, with the scale estimated from training only. Validation selects among 150 optimizer steps. The code freezes the selected gates, input scale and readout before computing the observed full-search scores.

This is an actually trained, classically simulated QNN, not quantum hardware and not a trained full Mora-3 model. Its classifier is a candidate statistic for a specified injected alternative. Success on that injection task does not establish that every physical CCC signal has that form, or that another physical process cannot produce similar features.

33.4 The null is part of the result

For each processed map, null surrogates multiply $m > 0$ harmonic coefficients by independent uniform phases and $m = 0$ coefficients by independent random signs. Every coefficient magnitude is retained. Monopole and dipole remain absent. Training, reference-normalization, calibration and validation use distinct seeds and realizations.

This conditional phase-randomization null is scientifically inspectable, but it is not a complete WMAP likelihood. Anisotropic instrumental noise, foreground-cleaning uncertainty, harmonic estimation and masking effects are not fully modeled. Exact randomization validity requires invariance under the declared phase action. A small rank can therefore diagnose disagreement with this surrogate law without uniquely identifying CCC.

For every calibration sky, the pipeline recomputes the same location, scale and method maximum as for the observed sky. There are 511 calibration draws, so the smallest reported rank is 1/512, never zero. Independent null validation uses 256 draws and reports its rejection fraction and Wilson

interval. Injection sensitivity uses 64 additional controls per declared amplitude, including zero amplitude. All products and all outcomes remain in the record.

Theorem 33.1 (Conservative finite upper ranks, including ties)**.** *For n real scores s_i, define $r_i = \#\{j : s_j \geq s_i\}$. Then for every integer k,*

$$\#\{i : r_i \leq k\} \leq k.$$

If the distinguished index is uniform under exchangeability, this gives the usual conservative finite-rank probability bound.

Proof. If the set on the left is empty there is nothing to prove. Otherwise choose a member with minimum score among its members. Every member's score is at least that minimum, so the whole set is included in the chosen member's upper-rank set. Its cardinality is at most that upper rank, which is at most k. Ties are included in the upper-rank set and require no arbitrary tie-breaking assumption. Uniformity of the distinguished index then divides the count by n. □

The deterministic counting theorem is formalized for arbitrary finite score vectors, not only one simulated sample. Exchangeability of real observations and a physical null is an additional modeling premise. The earlier measure-theoretic family bounds remain applicable once their event-level premises are justified.

33.5 Observed results and frequency sensitivity

Product	Full-search rank	Three-product bound	Null rejections
ILC	0.992188	1.000000	15/256
V	0.005859	0.017578	16/256
W	0.238281	0.714844	11/256

The original workstation run found a small joint phase-surrogate rank for the V-band product, but not for ILC or W-band. The independent runner preserves its own fully recorded recomputation rather than replacing it with those earlier numbers. These products share the CMB sky and observations; their ranks must not be multiplied as though they were independent confirmations. The displayed three-product correction is a union-bound allocation, not evidence that the underlying surrogate model is physically complete.

The search retains the selected template, per-method maxima and all calibration maxima. The low-variance statistic can be unusual under its own method-specific null while the joint maximum has a different tail. This is why both are reported and why the local selected tail is not substituted for the full-search result.

After the initial product disagreement was inspected, a separate exploratory followup was designed. It applies a common ILC-derived phase reference to all products and scans the processed V-minus-W map. It also reports masked cross-product correlations and map amplitudes. This followup is explicitly marked post-inspection. It does not replace the predeclared scan, is not added as independent evidence, and is not folded into a retroactively enlarged family with an unchanged significance claim.

The V-minus-W map is a sensitivity diagnostic, not a pure foreground estimate: different effective beams, noise and cleaning residuals remain. Similarly, changing from each product's conditional harmonic magnitudes to a common reference changes the null hypothesis. Those changes are exactly what a robustness analysis should expose, rather than conceal behind a single favorable number.

33.6 Interactive replay without moving the goalposts

The sky view displays observed temperatures on a Galactic Mollweide projection, with the fixed exclusion mask visibly separated. Product, statistic, center and injected amplitude are controlled independently. At zero amplitude the original observed map is replayed. At nonzero amplitude the display says that a synthetic profile has been added. All 72 templates and all three frozen statistics are then recomputed with unchanged weights and unchanged calibration.

The browser contains its own numerical implementation, including the actual four-qubit gates. It is compared with Python on all observed products and additional injected cases. A source-hash mismatch prevents the page from presenting an unverified data object as an accepted result. Exported state records the source commit, asset hashes, arithmetic outputs, quantum readouts and complete-search rank.

This completes a concrete CCC-oriented observational extension: real source products, a frozen learned statistic, full-search controls, independent null validation, injection sensitivity, explicitly exploratory frequency tests, formal rank logic and an interactive replay. It does not claim to have established a physical cross-aeon channel or an intentional message. Those questions require alternative source models with distinguishable predictive consequences, as the original information-theoretic chapter explains.

33.7 A joint phase action that retains cross-product dependence

A stronger dependence-aware control acts on the three harmonic maps together. For each mode choose one common unit-modulus phase $\zeta_{\ell m}$ and transform every product by

$$a_{\ell m}^{(A)} \longmapsto \zeta_{\ell m} a_{\ell m}^{(A)}.$$

Use a real sign for $m = 0$. Then every complex harmonic cross-product is preserved exactly in arithmetic:

$$(\zeta a^{(A)})\overline{(\zeta a^{(B)})} = a^{(A)}\overline{a^{(B)}}.$$

The implementation checks this identity numerically for every surrogate draw. Independent phases per product would preserve auto-power but generally destroy these cross-products; a negative test demonstrates the difference.

The new statistic is one maximum over all three products, all three methods and all 72 templates. Reference normalization, 511 calibration draws and 256 independent validation draws use the shared phase action. This gives one conditional family-wide rank without assuming independence of the product maps and without multiplying their ranks. The learned gates are unchanged.

This control was designed after the first product-specific scan and is labeled exploratory. It improves the representation of cross-product dependence, but it does not turn phase-group invariance into an established fact about the complete instrument and foreground-processing model. Its interpretation is still conditional on that hypothesis. Both the original scan and the new family control are retained, including disagreement between them.

Shared-phase family	Joint rank	Null rejections
V selected	0.017578	9/256

The common-reference, frequency-difference and joint-phase tests therefore answer different robustness questions. None is selected retrospectively as the uniquely valid answer because it gives a preferred tail. Their frozen configurations, source hashes, score arrays and validation outcomes make those distinctions reproducible.

Chapter 34

Neural computation, formal languages and a direct reply to Chomsky

Evidence status. The constructive recurrence, finite causal-network representation, nonidentifiability example, and finite-family generalization bound are proved here. Transformer expressivity results and linguistic experiments are attributed external research. The argument rejects categorical impossibility claims; it does not turn a simulation theorem into a theorem about consciousness or a claim that a particular trained model generalizes without measurement.

34.1 The claim being challenged

Chomsky's May 2023 interview makes two unusually strong claims: the design of large language models prevents them from contributing understanding of language or cognition as "a matter of principle, irremediable", and they "work just as well with impossible languages" [34]. These are stronger than the reasonable demand that a model of human language acquisition explain human learning constraints. They assert a principled barrier and a broad empirical equivalence. Neither follows from a neural architecture's use of prediction or from its ability to represent a large class of functions.

Our counterstance is direct: *the categorical dismissal is wrong.* It conflates the objective used to train a model, the algorithms representable by that model, the algorithms actually learned, and the evidence required for a theory of human cognition. A predictive objective does not impose a theorem that internal mechanisms must be shallow, nonrecursive, or incapable of counterfactual computation. The appropriate response is to specify a model, prove what it can represent, inspect the mechanism it learns, and test its generalization under controlled interventions.

There is a valid distinction beneath the disagreement. A system can be a useful scientific instrument without being a faithful developmental model of a child. Its scientific value may lie in testing a proposed inductive bias, locating an internal representation, constructing a counterexample, or generating a checkable proof. We do not need to claim that every successful neural network is a correct biological theory in order to reject the assertion that such systems can contribute nothing in principle. Conversely, a claim about what a trained commercial model actually understands cannot be settled merely by selecting a sufficiently permissive mathematical definition of neural network.

The tensor-grammar framework makes these distinctions concrete. A graph grammar supplies a syntax. Its interpretation supplies denotational semantics. A recurrent or attention-based policy supplies a computational mechanism. Training supplies a procedure for choosing its parameters. A

certificate verifies particular results. None of these levels is reducible to the others, but they can be composed into a system whose observable behavior has a precise semantic specification.

34.2 A constructive recursive recognizer

Let D_1 be the Dyck language with one kind of bracket. Encode an opening bracket by $+1$ and a closing bracket by -1. A word belongs to D_1 exactly when every prefix has nonnegative sum and the final sum is zero. Unlike a bounded collection of observed bracket strings, this is an unbounded language.

Write $\sigma(t) = \max(0, t)$ and

$$h(t) = \sigma(t) - \sigma(t-1).$$

For an integer t, $h(t) = 0$ when $t \le 0$ and $h(t) = 1$ when $t \ge 1$. Consider the fixed recurrent update, with initial state $(c_0, e_0) = (0, 0)$,

$$\begin{aligned} c_t &= c_{t-1} + x_t, \\ r_t &= h(-c_t), \\ e_t &= \sigma(e_{t-1} - r_t) + r_t. \end{aligned}$$

The output is

$$a_t = \sigma\big(\sigma(1 - |c_t|) - e_t\big), \qquad |c_t| = \sigma(c_t) + \sigma(-c_t).$$

All operations are affine combinations and ReLUs. A signed intermediate value can be represented by its positive and negative parts if every hidden coordinate is required to be nonnegative.

Theorem 34.1 (A fixed ReLU recurrence recognizes D_1)**.** *Under exact integer arithmetic, the displayed fixed-size recurrent network outputs one exactly on the words of D_1 and zero on all other bracket words. On inputs of length at most L, its count requires $O(\log(L+1))$ bits.*

Proof. Induction gives $c_t = \sum_{s \le t} x_s$. Hence r_t is one exactly when the current prefix has negative balance. For binary e, r, $\sigma(e - r) + r = \max(e, r)$, so induction gives $e_t = 1$ exactly when at least one prefix up to time t is negative. Since c_t is an integer, $\sigma(1 - |c_t|)$ is one exactly when $c_t = 0$ and is otherwise zero. The final ReLU therefore returns one exactly when $c_t = 0$ and $e_t = 0$, the defining condition for D_1. Finally $|c_t| \le L$, so a signed count has logarithmic bit length; the other recurrent state variable is binary. □

Proposition 34.2 (Why the unbounded language is not regular)**.** *D_1 is not recognized by any finite-state automaton.*

Proof. It suffices to consider deterministic automata, since a nondeterministic finite automaton can be determinized using its finite subsets of states. If a deterministic automaton has s states, two prefixes consisting of i and j opening brackets, $0 \le i < j \le s$, reach the same state. Appending i closing brackets must then produce the same acceptance decision for both. The first completed word is balanced and the second is not, a contradiction. □

This pair of proofs refutes a categorical assertion that a neural mechanism can only be a finite catalogue of shallow patterns. It also identifies the resource making the result possible: an integer count whose available precision grows with the input length. A fixed finite-precision device with bounded memory does not become an unbounded recognizer because the real-valued mathematical recurrence exists. The browser implements the exact recurrence for its bounded input field; the theorem describes the idealized family. Both descriptions are true at their respective scales.

The recurrence is constructed, not discovered by training. It therefore proves representability and supplies an executable specification. Whether a training method learns that algorithm from short examples, rather than a short-length shortcut, is a different experiment. A suitable test holds out longer lengths and deeper nestings, reports failure rates by depth, and compares states to the invariant $c_t = \sum_{s \le t} x_s$. Aggregate accuracy on the training length is inadequate.

34.3 Transformer theorems require their hypotheses

Pérez, Barceló and Marinkovic prove Turing completeness for a Transformer model with hard attention, positional encodings, and arbitrary-precision rational internal representations [120]. That theorem establishes a mathematical possibility under its specified model. It does not imply that a fixed-size softmax implementation with a fixed context window and hardware precision has unlimited memory, or that gradient descent finds the simulation.

Merrill and Sabharwal analyze intermediate decoding as a computational resource [98]. Their polynomial-step characterization uses logarithmic precision and, for the lower-bound construction, saturated attention, strict causal masking, and projected pre-normalization. Under those assumptions, polynomially many intermediate tokens suffice to capture polynomial-time computation. Their upper bounds also quantify the cost of shorter computations. These are external theorems, not results reproved or Lean-formalized in this companion.

Hahn's limitations for particular self-attention models remain part of the evidence [62]. A limitation for one fixed-depth, fixed-resource model and a simulation theorem for a model with growing computational resources need not conflict. Nor does failure of one architecture prove impossibility for every neural network. Any argument that says only "Transformers are Turing complete" or only "Transformers cannot handle recursion" has omitted the variables on which the theorems depend.

The same discipline applies to a quantum ToricGT. A variational feature map can change a model's hypothesis class and its computational cost. It does not supply unbounded precision, an unlimited context window, or a proof of generalization. The new experiment trains an explicitly small circuit and measures its downstream performance. It is a defensible component study, not an inference that quantum hardware resolves the linguistic debate by itself.

34.4 A neural circuit can compute a counterfactual

The distinction between predicting an observation and answering an intervention is scientifically essential. It does not imply that neural arithmetic is intrinsically unable to represent an intervention. We can prove the contrary in a finite setting without relying on a universal approximation slogan.

For Boolean inputs $x_1, \dots, x_m$, and a fixed pattern $b \in \{0,1\}^m$, define

$$I_b(x) = \sigma\left(\sum_{i:b_i=1} x_i + \sum_{i:b_i=0} (1 - x_i) - m + 1\right).$$

Every summand equals one exactly when the corresponding bit matches. Thus $I_b(x)$ is the exact indicator of $x = b$.

Proposition 34.3 (Exact Boolean mechanisms in a ReLU network)**.** *Every function $f : \{0,1\}^m \to \{0,1\}$ has the exact finite ReLU representation*

$$f(x) = \sum_{b:f(b)=1} I_b(x).$$

In particular, $x \wedge y = \sigma(x + y - 1)$ *and* $x \operatorname{xor} y = x + y - 2\sigma(x + y - 1)$.

Proof. For Boolean x, exactly one pattern indicator is one, namely I_x, and every other one is zero. The sum therefore equals $f(x)$. Checking the four Boolean input pairs yields the two displayed special cases. □

Theorem 34.4 (Finite structural counterfactuals have exact neural implementations)**.** *Consider a finite acyclic structural causal model with Boolean endogenous and exogenous variables and specified deterministic Boolean structural functions. Given the exogenous assignment and a finite collection of interventions replacing selected structural functions by constants, a finite ReLU network computes all the resulting endogenous values exactly.*

Proof. Topologically order the endogenous variables. Replace each structural function by the previous proposition's finite ReLU network, using its parents and exogenous variables as inputs. Composing in topological order computes the original model by induction. For an intervention indicator $i \in \{0, 1\}$ and desired value v, replace an original Boolean output f by

$$\sigma(f - i) + \sigma(v + i - 1).$$

If $i = 0$ this is f; if $i = 1$ it is v. Thus the same augmented network implements every chosen intervention mask. Reusing the same exogenous assignment across the original and intervened evaluations gives the structural counterfactual, with the replacement mechanism explicitly specified. □

This construction can be exponentially large in a mechanism's number of inputs. It proves capability, not universal computational efficiency. It also assumes that the causal mechanisms and the relevant exogenous state are supplied. Learning those from observations is a separate identifiability problem, exactly as emphasized in the causal-neural research of Xia and collaborators [180]. A mechanism written with ReLUs is still a mechanism. A neuron does not turn a structural intervention into a mere correlation.

Proposition 34.5 (Observational ambiguity survives every learner)**.** *There are two Boolean structural causal models with the same observational law for* (X, Y) *and different laws for* Y *under* $\operatorname{do}(X = 0)$*. Consequently no procedure given only that observational law can always return the correct intervention law for both models.*

Proof. Let U be a fair Boolean variable. In the first model put $X = U$ and $Y = X$. In the second put $X = U$ and $Y = U$. Both observational laws assign probability $1/2$ to $(0, 0)$ and $(1, 1)$. Replacing the structural equation for X by $X = 0$ makes $Y = 0$ surely in the first model, while in the second $Y = U$ remains fair. A procedure receives the same input law in both cases and so cannot be correct about two different output laws simultaneously. □

This impossibility applies equally to symbolic, neural, quantum, and human procedures with only those observations. It is an information limitation, not a defect peculiar to neural networks. Additional interventions, justified structural assumptions, or other observations can distinguish the models. The cosmology chapter uses precisely the same logic: a learned decoder cannot establish intentional cross-aeon communication if the proposed message and null hypotheses induce the same distribution on accessible observations.

The theorem also supplies a useful interpretation of "understanding" for this project. A system demonstrates a specified *operational semantic competence* when it implements a model's interpretation map and answers its permitted transformations and counterfactual queries correctly.

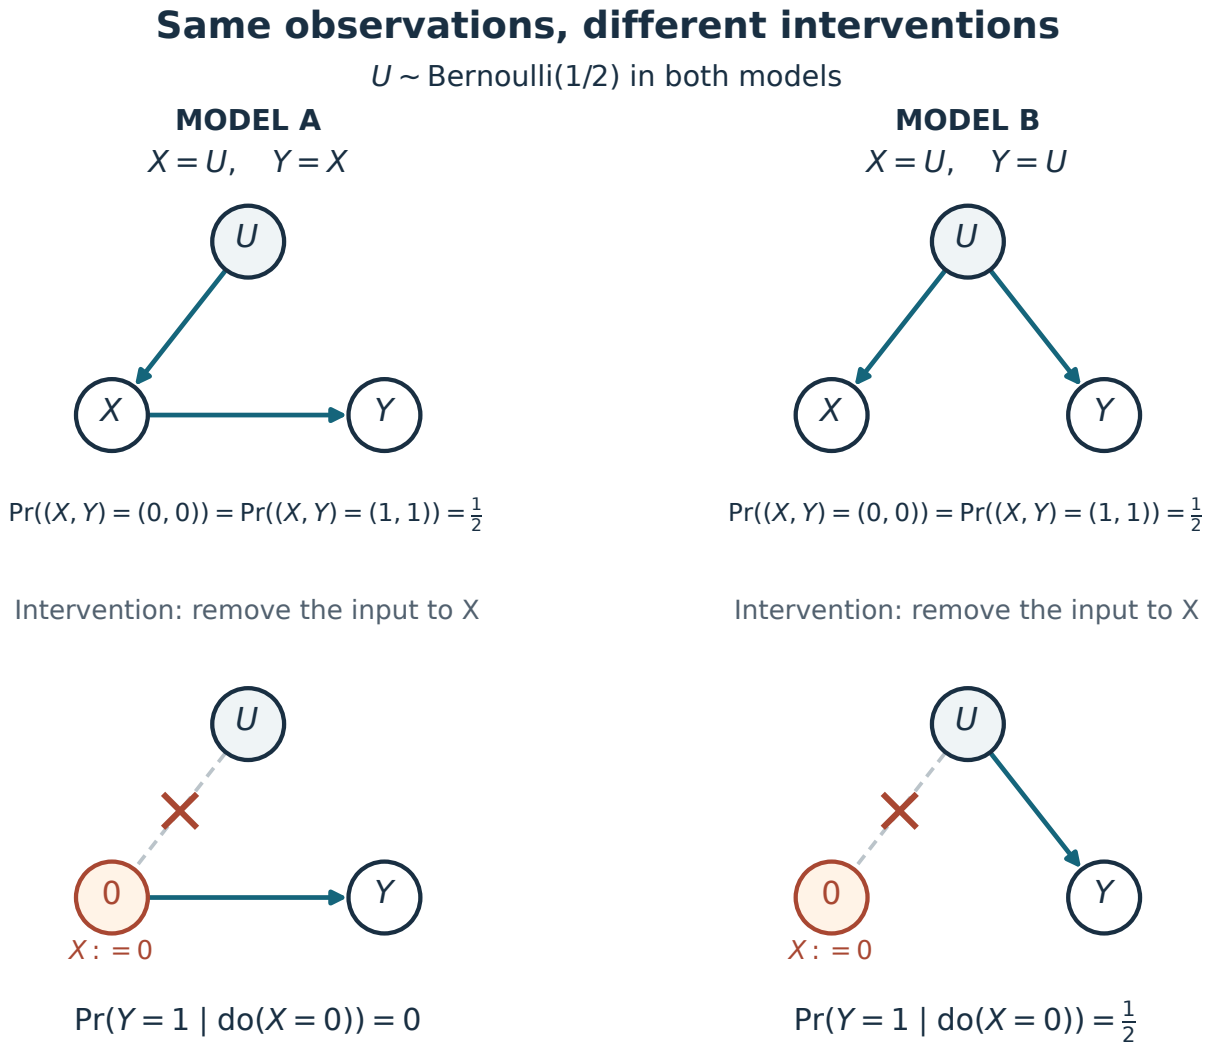


Figure 34.1: The same observational distribution does not identify a causal mechanism. Intervening on X removes its incoming arrow. The two resulting laws for Y differ even though the original observed (X, Y) laws coincide.

That is a testable property, with counterexamples and certificates. It does not settle phenomenal consciousness, the full semantics of natural language, or whether the supplied causal model is true of the world. Calling every such successful computation "mere pattern matching" adds no mathematical objection to its correctness.

34.5 Empirical linguistic evidence and its strongest objection

Gulordava and colleagues test long-distance agreement using semantically anomalous sentences and multiple languages, finding systematic syntactic generalization in recurrent language models [59]. Lakretz and colleagues identify internal units and circuits involved in number agreement and syntactic control [80]. These studies do not prove unlimited grammar mastery. They do provide concrete mechanisms and controlled behavior against which an assertion of exclusively superficial matching can be assessed.

Kallini and colleagues train GPT-2 variants on systematic transformations of English and observe differences in learning across their constructed languages [75]. Their experiment challenges an unrestricted prediction of equal learning behavior. Hunter's subsequent critique identifies a serious confound: a central comparison adds a count-based nonadjacent dependency to one condition while leaving the other adjacent, instead of matching count-based and constituency-based nonadjacent dependencies [66]. The experiment therefore does not settle whether the model's biases coincide with human grammatical constraints. Both findings belong in the assessment.

Comparisons between humans and language models also depend on the evaluation protocol. Lakretz and colleagues report a failure on particular long embedded dependencies [81]; Lampinen finds materially different performance when models receive prompts more comparable to the preparation supplied to human participants [82]. The appropriate conclusion is that instructions, examples, training distribution, and measured dependency type must be controlled. A benchmark failure is useful evidence about a defined setup. It is not an architecture-independent impossibility theorem.

For a stronger next experiment, construct pairs of synthetic grammars from one underlying set of annotated trees. Match vocabulary, sentence length, dependency distance, noise, and the number of extra transformation rules. Vary the grammatical relation of interest while keeping the nuisance quantities fixed or explicitly stratified. Train several architectures under matched data and optimization budgets; evaluate disjoint depths, vocabulary combinations, and intervention types. Compare test cross-entropy after accounting for the entropy of each data-generating process. Report counterexamples and uncertainty by stratum. This is a proposed experiment, not a result the present three-qubit graph benchmark has already established.

The hard counterstance does not require treating the weaker experiment as decisive. Chomsky's universal claim must itself meet a scientific burden. Demonstrated limitations and unresolved comparisons justify specific criticisms. They do not establish that neural models are permanently barred from representing mechanisms, learning useful linguistic structure, or contributing to explanations. Requiring an exact human developmental theory from every useful neural experiment mistakes one scientific question for all scientific questions about language.

34.6 Learning a rule selector without surrendering guarantees

Let $\mathcal{R}(G)$ be the finite set of certified grammar rules applicable to graph G. A model produces a probability distribution over this set. Its job may be to reduce a circuit's gate count, choose a template, propose a path, or select the next proof step. We can improve this proposal distribution

while preserving a separate invariant checker. In particular, correctness of accepted semantic rewrites is independent of how they were ranked.

Proposition 34.6 (A learned policy inherits certified rewrite correctness)**.** *Suppose a system starts at a well-typed diagram, chooses any applicable rule according to an arbitrary, possibly history-dependent policy, checks the rule's stated semantic certificate, and accepts only certified rewrites. Every finite accepted run denotes the same map as its initial diagram.*

Proof. Condition on any realized run. Each accepted step is a sound contextual replacement by Theorem 17.7. Induction on the finite number of accepted steps gives equal denotations. The policy's probabilities and parameters do not enter the argument. □

This is a substantive interface between learned and symbolic methods. The network can improve search order, while the checker enforces a property stated in mathematical terms. It does not make the checker unnecessary, and it does not make the learned component unimportant. Search cost can dominate a problem even when checking one proposed answer is easy. Whether a particular learned policy actually reduces that cost must be measured against a strong baseline with the same certificate budget.

For a quantum feature model, let $U_\theta(x)$ prepare a state and let O_j be specified observables. Its classical feature vector has coordinates $f_{\theta,j}(x) = \langle 0|U_\theta(x)^\dagger O_j U_\theta(x)|0\rangle$. These expectations are functions with a physical sampling cost and, for small circuits, an efficient classical implementation. A downstream model can use them to score graph rules. The model is therefore quantum in its feature construction when executed on quantum hardware, or a simulation of that construction when executed by cuQuantum or a CPU. A GPU state-vector simulation is a classical computation of a quantum model; this distinction is retained in the benchmark report.

34.7 Finite-family generalization with an explicit selection cost

A grammar can also generate candidate predictors. This introduces a statistical selection effect even if every candidate is mathematically well-typed. A small proof makes the cost visible.

Lemma 34.7 (Bounded-variable exponential bound)**.** *If $X \in [0,1]$ almost surely and $\mu = \mathbb{E}X$, then $\mathbb{E}e^{\lambda(X-\mu)} \le e^{\lambda^2/8}$ for all real λ.*

Proof. Let $\psi(\lambda) = \log \mathbb{E}e^{\lambda X}$. Boundedness permits differentiation. Its second derivative is the variance of X under the exponentially tilted probability measure. Any probability law supported on $[0,1]$ satisfies $\mathrm{Var}(X) = \min_a \mathbb{E}(X-a)^2 \le \mathbb{E}(X-1/2)^2 \le 1/4$. Thus $\psi'' \le 1/4$, while $\psi(0) = 0$ and $\psi'(0) = \mu$. Integrating this bound twice, or applying Taylor's theorem with integral remainder on the interval between zero and λ, gives $\psi(\lambda) - \lambda\mu \le \lambda^2/8$. □

Theorem 34.8 (Uniform error for a finite candidate family)**.** *Let $K, m \ge 1$ be integers and let $f_1, \dots, f_K$ be predictors fixed independently of an independent identically distributed evaluation sample of size m. Let the loss lie in $[0,1]$, let R_k be its population mean and $\widehat{R}_k$ its sample mean. For $0 < \delta < 1$, with probability at least $1-\delta$,*

$$\max_{1 \le k \le K} |\widehat{R}_k - R_k| \le \sqrt{\frac{\log(2K/\delta)}{2m}}.$$

On this event the empirical minimizer's true risk is at most the best candidate's true risk plus twice this bound.

Proof. Independence and the exponential bound give

$$\mathbb{E}\exp\left(\lambda\sum_{i=1}^{m}(X_i-\mu)\right)\le\exp(m\lambda^2/8).$$

Markov's inequality implies $\Pr(\widehat{\mu}-\mu\ge\epsilon)\le\exp(-\lambda m\epsilon+m\lambda^2/8)$. For $\lambda=4\epsilon$ this is $e^{-2m\epsilon^2}$; applying the same argument to $1-X$ gives the lower tail. A union bound over K candidates yields $2Ke^{-2m\epsilon^2}$. Solving for ϵ proves the displayed confidence bound. If $\hat{k}$ minimizes sample risk and k_* minimizes true risk, then $R_{\hat{k}}\le\widehat{R}_{\hat{k}}+\epsilon\le\widehat{R}_{k_*}+\epsilon\le R_{k_*}+2\epsilon$. □

The independence condition matters. If the same evaluation data are repeatedly used to invent new predictors, the displayed proof no longer applies with the originally declared K. A finite family generated from a separate training sample can be conditioned on that sample and then evaluated on a fresh sample. More general adaptive selection requires another argument or a genuinely untouched final evaluation set. The cosmology chapter's global-null calibration addresses the analogous search over sky positions, template scales, and learned scoring rules.

This theorem does not certify the three-seed experiment merely because three is finite: the predictors are trained and selected, and their measured losses, splits, and boundedness must match its hypotheses. It supplies a reusable design rule. The reported benchmark uses explicit disjoint splits and exposes every seed; the finite-family theorem explains why tuning and final evaluation should be separated.

34.8 What follows for the proposed cosmological grammar

There are three layers in the proposed application. A physical model specifies how source degrees of freedom evolve through a crossover and into observable fields. A tensor grammar parameterizes candidate structures or messages consistent with that model. A learned function ranks or decodes those candidates using observed data. The second and third layers cannot establish the first by definition. They can make a specified physical hypothesis more testable.

One productive version is a restricted grammar of localized templates on an observed sky graph, with tensor-valued features for temperature, polarization, and angular correlations. A quantum feature map can feed an invariant graph scorer; a classical matched-filter or likelihood model supplies a strong baseline. The grammar's interface width determines a contraction budget. Null simulations must include the same mask, noise, foreground model, search over positions and widths, and learned model-selection procedure. Only held-out predictive success can establish that the search detected the stipulated class of signal better than its baseline. Attribution to a deliberate sender requires further, independently testable encoding structure.

At the logical level, the language of possible encoded messages can be regular, context free, or computationally more expressive depending on its stipulated encoder. Physics does not choose a Chomsky class merely by being quantum. A finite observed sky gives a finite dataset; an asymptotic language classification concerns a family of encoders and increasingly large observations. These are related through a specified model, not through the word "information" alone.

The same reasoning sharpens the reply to Chomsky. A neural system can participate in a disciplined semantic pipeline, can implement recursive and counterfactual algorithms, and can be evaluated on transformations that expose what it has learned. Demanding those tests is good science. Declaring the entire approach incapable of understanding in principle is not a consequence of the mathematical hierarchy, and is contradicted as a broad computational prohibition by the explicit constructions above. The unresolved question is which trained systems acquire which mechanisms,

at what data and computational cost, with what robustness. That is an empirical and mathematical research program, not a reason to dismiss the program before doing it.

Part VIII

Certified surface computation and its measured realizations

Chapter 35

A calculus of certified realizations

35.1 The organizing principle

This program is organized around a concrete mathematical demand: a proposed connection must carry its objects, its maps and the identities those maps preserve. A dessin is not used as an ornament for a neural network. A quiver is not used as a synonym for any directed picture. An order is not interchangeable with every integral algebra. Each construction earns its place by doing calculable work.

The common substrate is a typed finite presentation together with its realizations. A cellular dessin supplies permutations, incidence and a surface. Quivers with relations supply module categories and complexes. Integral lifts retain arithmetic that disappears over a field. Metrics turn selected complexes into Hodge operators and energy forms. Binary reduction produces commuting Pauli checks when the relevant chain identity survives. Graph-to-graph models propose maps and refinements; an independent checker decides which algebraic, numerical or probabilistic statements those proposals satisfy.

The ambition is substantial: use this substrate to organize spectral arithmetic, nonlinear dynamics, quantum information, operator learning and structured inference within one inspectable program. The standard is equally substantial. Every claimed bridge must specify its realization map; every limiting claim must control the limit; every empirical gain must survive a comparison that charges its information and computational costs. A skeptical label is not a refutation. The appropriate response to a construction is to examine its hypotheses, follow its proof and reproduce its tests.

35.2 From a relation to topology, energy and quantum checks

A representation of the gentle bound quiver

$$2 \xrightarrow{a} 1 \xrightarrow{b} 0, \qquad ba = 0,$$

is precisely a two-step complex $C_2 \xrightarrow{d_2} C_1 \xrightarrow{d_1} C_0$. A closed oriented cellular surface supplies such a complex over $\mathbb{Z}$ through its boundary maps. The incidence matrices and the relation $d_1 d_2 = 0$ are explicit finite data. Choosing positive inner products gives the finite Hodge operator

$$\Delta_1 = d_2 d_2^* + d_1^* d_1.$$

For every u, its quadratic form is $\|d_2^* u\|^2 + \|d_1 u\|^2$. Consequently its kernel is exactly $\ker d_2^* \cap \ker d_1$. The positive metric is part of this realization; a general change of module basis is not automatically unitary.

Reducing an integral cellular complex modulo two gives $H_X H_Z^T = 0$ for the corresponding incidence choices. Pauli operators $X(a)$ and $Z(b)$ commute exactly when their binary supports have even intersection. Thus the chain identity becomes a quantum stabilizer commutation identity. Over a connected closed genus-g surface, the familiar rank calculation gives $2g$ encoded qubits when the surface and rank hypotheses hold. The existing surface-code implementation checks these hypotheses on its actual cellulations and executes their preparation circuits.

This passage is a specific realization, not a claim that every gentle representation is a quantum code. Conversely, a finite binary check does not recover all integral torsion. The BGG and orders chapters retain that arithmetic layer rather than discarding it for computational convenience.

35.3 An explicit order whose special fiber is gentle

There is a direct bridge between an integral order and a gentle relation, not merely a resemblance between their diagrams. Let R be an integral domain, K its fraction field, and $0 \neq \pi \in R$. Define

$$\Lambda_\pi = \left\{ \begin{pmatrix} a & b \\ \pi c & d \end{pmatrix} : a, b, c, d \in R \right\} \subset M_2(K).$$

It is a unital subalgebra, free of rank four over R, and spans $M_2(K)$. For a Dedekind domain it is an order in the sense fixed in the original monograph. Localization at a prime generated by π yields the corresponding local tiled-order example.

Proposition 35.1 (Multiplication, generic fiber and gentle specialization)**.** *In coordinates (a, b, c, d) the multiplication is*

$$(a, b, c, d)(e, f, g, h) = (ae + \pi bg,\ af + bh,\ ce + dg,\ \pi cf + dh).$$

Let e_1, e_2 be the diagonal idempotents, $u = E_{12}$ and $v = \pi E_{21}$. Then

$$uv = \pi e_1, \quad vu = \pi e_2, \quad u^2 = v^2 = 0.$$

If $R/(\pi) = k$ is a field, the special fiber $\Lambda_\pi/\pi\Lambda_\pi$ is the four-dimensional radical-square-zero two-cycle algebra over k, which is gentle. Its generic fiber is $M_2(K)$. The determinant of the trace pairing in the basis (e_1, u, v, e_2) is $-\pi^2$.

Proof. Ordinary matrix multiplication gives the coordinate formula, closure and associativity. The identity has coordinates $(1, 0, 0, 1)$. The four basis matrices are R-linearly independent and span Λ_π; after extending scalars, $E_{21} = \pi^{-1} v$, so all four matrix units are present. The displayed arrow relations follow by direct multiplication.

Modulo $\pi\Lambda_\pi$, all four coordinate classes remain independent. In particular v does not vanish merely because its displayed matrix entry is divisible by π: the lower-left entry of $\pi\Lambda_\pi$ is divisible by π^2. Both length-two cycles vanish in the quotient. The two idempotents and two arrows give exactly the radical-square-zero two-cycle presentation. Each arrow has one forbidden successor/predecessor and no allowed length-two continuation, proving the gentle conditions and finite dimensionality. The trace-pairing matrix has diagonal entries one on the idempotent coordinates and an off-diagonal π block on (u, v), so its determinant is $-\pi^2$. □

This example identifies precisely where arithmetic enters. The generic fiber, special fiber and integral lattice answer different questions. The relation $uv = 0$ in the special fiber is the shadow of $uv = \pi e_1$ in the order. A model trained only on the special-fiber graph cannot reconstruct the deformation parameter unless it is supplied or inferred from additional data. The implementation checks the coordinate realization, multiplication and specializations; selected identities are formalized in Lean over a general commutative ring.

35.4 A finite spectral gap, proved and executed

On the checked 3×3 torus, take nine star checks A_v and nine plaquette checks B_f, acting on eighteen edge qubits. For positive J_X, J_Z, set

$$H = -J_X \sum_v A_v - J_Z \sum_f B_f.$$

The commuting checks have two independent global relations, $\prod_v A_v = I$ and $\prod_f B_f = I$, and sixteen independent generators. The ground space therefore has dimension four.

Theorem 35.2 (Exact finite toric-code spectrum)**.** *For this cellulation, the energy labeled by even integers $a, b \in \{0, 2, 4, 6, 8\}$ is*

$$E_{a,b} = -J_X(9 - 2a) - J_Z(9 - 2b),$$

with multiplicity $4\binom{9}{a}\binom{9}{b}$, combining equal energies when necessary. The ground energy is $-9J_X - 9J_Z$ and the gap is

$$\Delta = 4 \min(J_X, J_Z).$$

For a normalized state with ground-space probability p_0,

$$\langle H \rangle - E_0 \geq \Delta(1 - p_0).$$

Proof. Commutation permits simultaneous diagonalization. The two global relations force an even number of negative star eigenvalues and an even number of negative plaquette eigenvalues. Independent stabilizer generators realize every compatible character with the same four-dimensional sector multiplicity: the character projector has trace $2^{18}/2^{16} = 4$, since nonidentity Pauli products have zero trace. Counting the choices of violated checks gives the multiplicity formula, whose sum is $4 \cdot 2^8 \cdot 2^8 = 2^{18}$.

Every nonzero excitation has at least two violations of one type, so the energy increment is at least $4 \min(J_X, J_Z)$. A single-edge Pauli Z flips its two distinct endpoint stars, and a single-edge Pauli X flips its two distinct neighboring plaquettes; the bound is attained. Finally, decompose a normalized state into the simultaneous eigenspaces. Each nonground component costs at least Δ, giving the expectation inequality. □

The program executes this Hamiltonian's preparation, controlled rotations, CNOT circuit and readouts. The new explorer links the selected physical edge to its dessin cycles, attention coordinates and quantum gates. Independent NumPy, browser and cuStateVec implementations agree at recorded tolerances. This is a completed finite spin-system result. A four-dimensional continuum Yang–Mills construction requires the additional analytic transition below; using the same phrase "mass gap" does not identify the two operators.

35.5 The continuum transition as a precise theorem

The finite-to-continuum step can be stated without handwaving. Let $A_n \geq 0$ be self-adjoint operators on a common Hilbert space, and let P_n be their kernel projections. Suppose A_n converges to a nonnegative self-adjoint A in strong resolvent sense, $P_n \to P$ strongly for an orthogonal projection P, and every A_n satisfies a common gap $\Delta > 0$ off its kernel.

Proposition 35.3 (Gap transport with a controlled vacuum projection)**.** *Under these assumptions, P is the kernel projection of A, and*

$$\operatorname{spec}(A) \subseteq \{0\} \cup [\Delta, \infty).$$

To retain a nonzero vacuum, require $P \neq 0$.

Proof. Write $R_n = (A_n + I)^{-1}$ and $c = (1 + \Delta)^{-1}$. Spectral calculus gives

$$R_n P_n = P_n, \qquad 0 \leq R_n \leq P_n + c(I - P_n).$$

Strong resolvent convergence implies $R_n \to R = (A + I)^{-1}$ strongly. Because $\|R_n\| \leq 1$, passing to strong limits in the first identity yields $RP = P$. Passing to quadratic forms in the second gives $R \leq P + c(I - P)$. Thus P lies in the eigenspace of R at one, and hence in $\ker A$. On its orthogonal complement $R \leq cI$; therefore no additional vector can belong to $\ker A$. Spectral calculus on that complement yields $A \geq \Delta I$. These two statements prove the conclusion. □

For Yang–Mills, a successful application must construct the continuum theory, the compatible operator realization, the scaling, the resolvent limit and the nonvanishing vacuum projection. The proposition specifies the exact job of that analytic construction. Numerical gaps without these transports do not enter its premises. This is an actionable theorem architecture: it says what must be proved and what follows once it is proved.

35.6 Probability diffusion as an open quantum evolution

Let P be a finite nonnegative column-stochastic matrix. Define

$$K_{ij} = \sqrt{P_{ij}}\, |i\rangle\langle j|, \qquad \Phi_P(\rho) = \sum_{i,j} K_{ij} \rho K_{ij}^*.$$

Then $\sum_{i,j} K_{ij}^* K_{ij} = I$, so this is a completely positive trace-preserving map. Direct multiplication gives

$$\Phi_P(\rho) = \operatorname{diag}(P \operatorname{diag} \rho).$$

Thus diagonal probability evolution is exactly the classical Markov update. If P is also row-stochastic, the channel is unital. The periodic heat stencil in the workbench supplies such a P for $0 \leq r \leq 1/4$.

This establishes an actual quantum-channel realization of finite diffusive probability transport. It also makes the information cost visible: off-diagonal coherence is removed. The channel is not a unitary relabeling of diffusion, and a velocity field is not automatically a density matrix. Nonlinear fluid and quantum-fluid models require their own state encoding and physical evolution law. The original monograph's finite energy, pressure, projection and residual identities supply exact interfaces for those extensions.

The 8 September 2026 Navier–Stokes release is a separately attributed forced-flow singularity construction with public Lean artifacts. Its force, support, regularity and norm hypotheses remain attached to the theorem. The present program reconstructs and studies its mathematical interfaces; that attribution is not replaced by claiming that a finite heat channel independently reproves the analytic result. Clay's 11 September statement describes its evaluation process [U1,U2].

35.7 Controller-to-circuit robustness

The new hybrid example composes a classical coordinatewise normalized max-plus map with quantum rotations. For score perturbation at most δ_s and value perturbation at most δ_v,

$$y_c = \max_j (s_j + v_{jc}) - \max_j s_j \quad \Longrightarrow \quad |y'_c - y_c| \le 2\delta_s + \delta_v.$$

This follows from nonexpansiveness of finite maximum. If $\theta_c = b_c + g y_c$, the four rotation angles change by at most $g(2\delta_s + \delta_v)$ each.

For a single rotation, $\|R_Y(\theta) - R_Y(\theta')\| \le |\theta - \theta'|/2$. A telescoping product of unitary gates therefore changes the normalized state by at most half the sum of angle errors. For an observable of operator norm at most one, inserting and subtracting the cross term bounds its expectation error by twice the state error. Hence this four-angle controller satisfies

$$|\langle Z\rangle' - \langle Z\rangle| \le 4g(2\delta_s + \delta_v).$$

This is a concrete neural-to-quantum error budget. A hard score router needs an additional winner margin because its selected value may jump at a tie. The distinction is visible in the interactive controls and tested independently.

35.8 The arithmetic spectral target

A Hilbert–Pólya application requires an arithmetic identity, not a visual match between spectra. Here is a precise sufficient interface. Let T_χ be a compact self-adjoint Hilbert–Schmidt operator with real eigenvalues μ_j. Its regularized determinant is

$$D_\chi(z) = \prod_j (1 - z\mu_j) e^{z\mu_j}.$$

The product converges locally uniformly: for sufficiently small w, $\log(1-w) + w = O(|w|^2)$, and $\sum_j \mu_j^2 < \infty$. Its zeros are exactly the nonzero reciprocal eigenvalues, with multiplicity; away from them the convergent logarithmic tail is nonvanishing.

If an independently constructed entire completion of the relevant L-function satisfies

$$\Xi_\chi(s) = e^{g_\chi(s)} z^{m_\chi} D_\chi(z), \qquad z = -i(s - 1/2),$$

with entire g_χ and integer $m_\chi \ge 0$, every zero has real z, and therefore $\Re s = 1/2$. The proof is the product's zero characterization and self-adjointness. The arithmetic task is the construction of T_χ and the exact identity, without assuming the desired zero locations in the construction. Orders, characters, trace formulas and module transports can organize that task; the identity remains the criterion that makes the spectral route decisive.

35.9 One program, distinct routes to evidence

The resulting architecture is coherent. Dessins and cellular maps supply finite presentations. Gentle and BGG structures organize relations and resolutions. Orders track integral specialization. Hodge operators and stochastic channels furnish analytic and probabilistic realizations. Surface spin systems furnish quantum realizations. ToricGT, TropicalGT and Mora-3 supply graph-valued

proposal and approximation mechanisms. Exact witnesses, Lean declarations, numerical residuals and held-out code lengths determine which transitions have been justified.

For language and cosmological inference, compressibility establishes predictive structure relative to a code, not intentionality by definition. The new control study uses genuine human speech with independently established provenance and genuine detector strain with a known physical event. It charges its model and side information and includes negative results. The purpose is to make a proposed information-bearing grammar confront observations under a protocol capable of disproving it.

The exposition therefore makes no appeal to modesty and no appeal to prestige. A claim is as strong as its hypotheses and proof. A failed boundary case is repaired, not hidden. A positive result is stated directly, not buried under ritual disclaimers. The reader is invited to inspect the complete chain from presentation to realization to prediction. That is the foundation on which a broad mathematical program can be built and extended.

U1 OpenAI. *On the Navier–Stokes Millennium Prize Problem*, 8 September 2026. Public announcement and links to the released proof artifacts. https://openai.com/index/navier-stokes-solution/

U2 Clay Mathematics Institute. *Navier–Stokes Announcement*, 11 September 2026. https://www.claymath.org/news/navier-stokes-announcement/

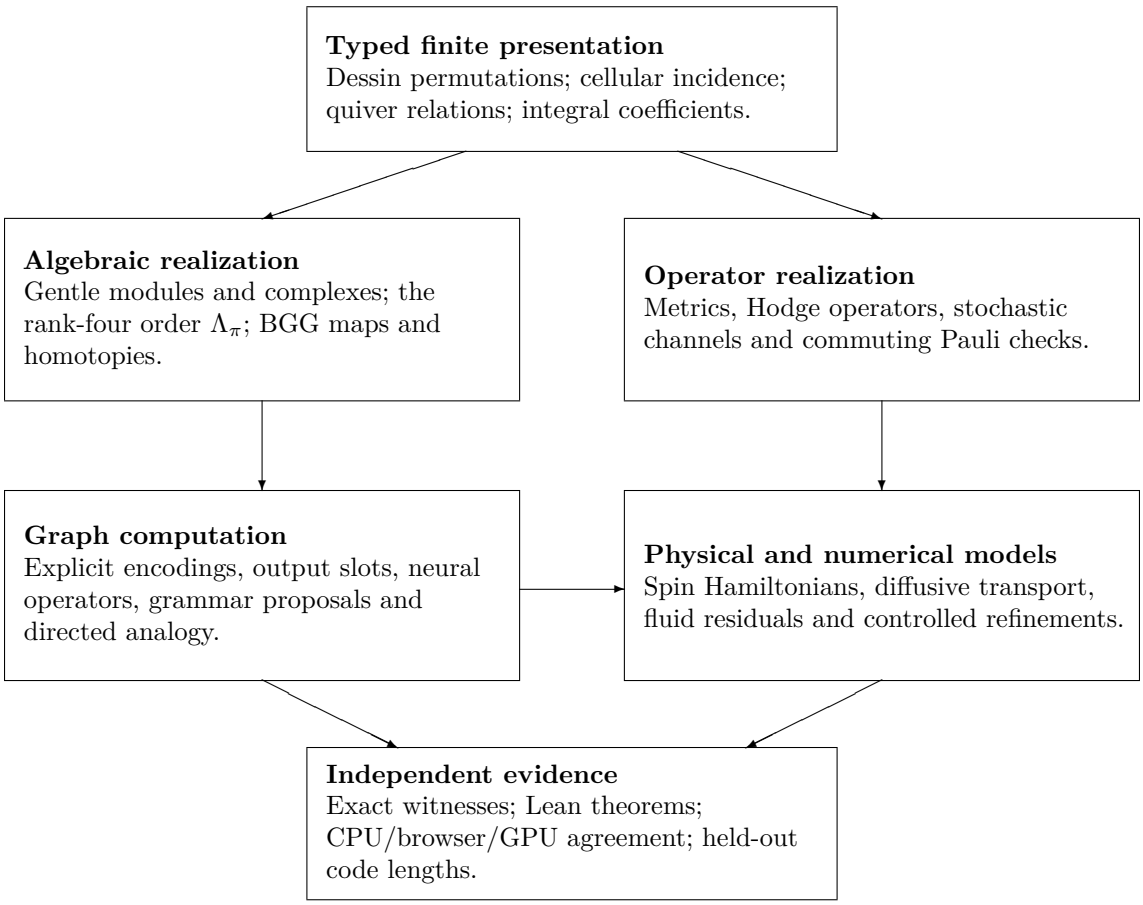


Figure 35.1: The program is a calculus of specified realizations. Arrows name mathematical interfaces, not automatic identifications of entire theories. The text supplies explicit constructions and proof obligations for each application.

Chapter 36

Nonlinear fluid coordinates and unitary half-density transport

The link between nonlinear dynamics and quantum operator language can be made exact without pretending that viscosity is a unitary gate. The correct object is a half-density on the classical state space. Its evolution is unitary even when physical kinetic energy decreases. The distinction is structural: Hilbert norm, phase-space volume and physical energy are three different quantities.

This construction belongs to the Koopman–von Neumann tradition. Joseph's peer-reviewed treatment develops its use in quantum simulation of nonlinear non-Hamiltonian dynamics [H1]. The result below states the complete finite-dimensional energy-controlled interface used by this program and derives its action explicitly. Stone's theorem supplies the self-adjoint generator of the strongly continuous unitary group [H2].

36.1 A complete flow from the finite energy identity

Let $u \in \mathbb{R}^d$ be coordinates in a fixed, metrically orthonormal basis of a finite divergence-free cochain space. Consider the autonomous system

$$\dot{u} = b(u) = B(u,u) - Au + f,$$

where B is bilinear, $u \cdot B(u,u) = 0$ for every u, A is a fixed real matrix and f is fixed. A symmetric positive-semidefinite A gives the usual dissipative interpretation. The energy-consistent finite fluid interfaces in the monograph have this form after the declared projection and metric choice.

Theorem 36.1 (Complete nonlinear flow and its unitary realization)**.** *The vector field above generates a smooth complete flow Φ_t for every $t \in \mathbb{R}$. Write $J_t(x) = \det D\Phi_t(x) > 0$. On $L^2(\mathbb{R}^d, dx)$ define*

$$(U_t\psi)(x) = J_{-t}(x)^{1/2}\psi(\Phi_{-t}(x)).$$

Then U_t is a strongly continuous unitary group. It has a self-adjoint generator H such that $U_t = e^{-itH}$. On compactly supported smooth functions,

$$H\psi = -i\left(b\cdot\nabla\psi + \tfrac{1}{2}(\operatorname{div} b)\psi\right).$$

The probability density $\rho_t = |U_t\psi|^2$ is the pushforward of $|\psi|^2$ by the classical flow and satisfies the Liouville transport equation.

Proof. Let $E(u) = \|u\|^2/2$. The nonlinear term cancels exactly:

$$\frac{dE}{dt} = -u \cdot Au + u \cdot f.$$

Hence

$$\left|\frac{dE}{dt}\right| \le (2\|A\| + 1)E + \tfrac{1}{2}\|f\|^2.$$

Gronwall's inequality bounds E on every bounded positive or negative time interval. Since b is polynomial and locally Lipschitz, a finite endpoint of a maximal solution would require the solution to leave every compact set. The energy bound excludes this. Thus the flow is complete; smooth dependence and uniqueness give a smooth diffeomorphism with inverse Φ_{-t}. The variational equation and Jacobi's formula give

$$J_t(x) = \exp\left(\int_0^t \operatorname{div} b(\Phi_s(x))\, ds\right) > 0.$$

Changing variables $x = \Phi_t(y)$ gives

$$\|U_t\psi\|_2^2 = \int J_{-t}(\Phi_t(y))|\psi(y)|^2 J_t(y)\, dy = \|\psi\|_2^2.$$

The inverse-flow and Jacobian chain rules prove $U_tU_s = U_{t+s}$ and $U_{-t} = U_t^{-1}$. For compactly supported smooth ψ, the flow and Jacobian converge uniformly to their initial values on the relevant compact neighborhood as $t \to 0$, so dominated convergence gives strong continuity. Density of these functions and the norm-one bound extend strong continuity to all of L^2. Stone's theorem gives the self-adjoint generator. Differentiating the displayed action at zero in L^2 yields $\partial_t U_t\psi|_0 = -b \cdot \nabla\psi - \frac{1}{2}(\operatorname{div} b)\psi$, which identifies H on the stated domain. Taking the squared modulus gives the usual change-of-variables formula for transported probability and therefore the Liouville equation. □

Why dissipation is compatible with unitarity. For an observable F and states for which the integrals are finite,

$$\langle U_t\psi, FU_t\psi\rangle = \int F(\Phi_t(y))|\psi(y)|^2\, dy.$$

Taking $F = E$ transports the classical energy decay into a changing expectation of the energy observable. It does not turn that observable into the unitary generator H. In general the transport generator need not be bounded below. Consequently this realization is not a positive-energy quantum-field Hamiltonian or a mass-gap proof merely because Stone's theorem applies.

Spatial incompressibility and phase-space incompressibility must also be distinguished. The cochain constraint imposes the first. Viscosity can make the divergence of the finite coefficient-space vector field negative. A naive pullback $\psi \circ \Phi_{-t}$ then fails to preserve the Lebesgue L^2 norm; the half-Jacobian factor is essential.

36.2 A nonlinear triad with an exact Jacobian law

The executed reference uses

$$B(u,u) = (au_2u_3, bu_1u_3, cu_1u_2), \qquad a + b + c = 0,$$

with $(a, b, c) = (1, -1/2, -1/2)$, $A = \nu\,\mathrm{diag}(1, 2, 4)$, $\nu = 0.13$ and $f = (0.03, -0.02, 0.05)$. Its energy identity is the theorem's identity, and

$$\operatorname{div} b = -7\nu, \qquad J_t = e^{-7\nu t}.$$

The implementation independently integrates the state, its variational matrix and the accumulated energy work, then checks inverse-flow composition, two-half-step composition, the exact determinant and energy balance. These are constructed dynamical reference tests, kept separate from the measured speech, detector and physiological records.

For the one-dimensional contraction $\Phi_t(x) = e^{-at}x$, the normalized initial Gaussian amplitude gives

$$|U_t\psi(x)|^2 = \frac{e^{at}}{\sqrt{\pi}} \exp(-e^{2at}x^2).$$

Its integral is one and its second moment is $e^{-2at}/2$. Omitting the half-density factor gives squared norm e^{-at} instead. Numerical quadrature checks both the valid realization and this deliberately invalid alternative. The exact algebraic energy and half-density coefficient identities are separately formalized in Lean.

36.3 A unitary discretization and its actual quantum operation

On a periodic grid let D be the central-difference matrix, so $D^T = -D$, and let V be the real diagonal velocity matrix. The discrete half-density transport operator

$$A_h = \tfrac{1}{2}(VD + DV)$$

is skew-adjoint. The associated $H_h = -iA_h$ is Hermitian. For real step τ, define

$$C_\tau = (I + \tfrac{\tau}{2}A_h)^{-1}(I - \tfrac{\tau}{2}A_h).$$

The first factor exists: the squared norm of $(I + \frac{\tau}{2}A_h)x$ is $\|x\|^2 + \frac{\tau^2}{4}\|A_h x\|^2$. The two polynomial factors commute and are adjoints. Multiplication therefore gives $C_\tau^* C_\tau = I$. The abstract normal-unit calculation is verified in Lean.

For a fixed bounded matrix, expanding the Cayley transform shows local error $O(\tau^3)$ relative to $e^{-\tau A_h}$. A telescoping comparison over a fixed horizon, using unit norm of both evolutions, gives global error $O(\tau^2)$. The code checks this order against independent Hermitian eigendecomposition at 10, 20 and 40 steps. Exact discrete unitarity does not imply zero time- or space-discretization error.

The sixteen-state periodic reference is encoded on four qubits. Its Cayley matrix is passed to actual cuStateVec as a dense four-qubit unitary and compared with independent NumPy multiplication. This checks a genuine backend operation. The supplied dense matrix is not an efficient elementary-gate synthesis, and that compilation cost is not hidden in a claim of quantum speedup. The separate eighteen-qubit surface experiment already supplies explicit Hadamard, rotation and CNOT gates.

36.4 Operator learning, cyclic structure and the continuum problem

A learned graph operator can propose b, a projected update, or a map between finite cochain spaces. Energy-compatible parameterizations preserve the cancellation identity by construction;

unconstrained models must pass its residual test. Uniform approximation of a vector field on an energy-controlled domain gives a finite-horizon trajectory bound through Gronwall. Controlling the Jacobian factor additionally requires derivative information, such as a C^1 approximation bound or a separately verified divergence envelope. A C^0 approximation claim alone does not certify the half-density amplitude.

A return map Φ_T supplies a unitary monodromy operator U_T in this realization. It is a precise operator attached to a dynamical cycle, not evidence that physical cosmological time is periodic. A smooth positive time change replaces b by the corresponding rescaled vector field and changes both transport and divergence terms. Singular conformal crossover remains a separate extension problem; the earlier CCC chapter specifies the physical fields, matching data and observational hypotheses needed there.

Finite energy-compatible Galerkin systems are complete by the theorem above. They do not reproduce a continuum singularity by simply blowing up a finite vector while its Euclidean energy remains bounded. Passing to the continuum requires control of the growing discretization space and the stronger norms in which regularity is tested. This pinpoints the analytic task in relating the program to the released Navier–Stokes proof. It is also why the framework keeps topology, metric, approximation scale and operator domain in the same explicit record.

H1 Ilon Joseph. *Koopman–von Neumann approach to quantum simulation of nonlinear classical dynamics.* Physical Review Research 2, 043102 (2020). https://doi.org/10.1103/PhysRevResearch.2.043102

H2 M. H. Stone. *On One-Parameter Unitary Groups in Hilbert Space.* Annals of Mathematics 33(3), 643–648 (1932). https://doi.org/10.2307/1968538

Chapter 37

Source-led graph universality: preserved proofs and constructive repairs

37.1 What was missing from the preceding synthesis

The earlier discussion correctly distinguished Yun's L^p approximation result from a supremum guarantee, but it did not give a separate, sufficiently detailed treatment of `TropicalGT/references/main.tex`. That manuscript supplies an additional invariant-graph-network argument for uniform graph-to-graph approximation, a finite interpolation construction, circuit simulations and a tropical geometric analysis. These are now preserved as source statements and examined individually. The positive program should not be reduced to one citation, nor should a statement be promoted to a verified theorem solely because it is typeset in a theorem environment.

The exact-source port contains all 17 TeX research documents and six Mora-3 paper Markdown documents in the five accessible pinned trees. It records 475 named claim occurrences, 248 distinct statement hashes and 463 following proof occurrences. Repeated versions remain attributable rather than silently merged. Every original file is checked against its Git blob and a SHA-256 hash. The sixth supplied repository name, `amelie-iska/TropicalGT-Program`, returned 404 through both the public API and selected connector. TropicalGT's own Parts I–IV are included, but are not represented as proof that the inaccessible name denotes the same repository.

Evidence categories. The archive preserves exact source text. The review register records separately: accepted written argument under explicit hypotheses; repaired statement/proof; finite executable witness; selected Lean elaboration; and a remaining dependency. Source extraction does not mean that 475 results have been formalized in Lean. External full-text PDFs without corresponding authored source are not silently included in that reading claim.

37.2 A direct uniform graph-to-graph construction

Here is a positive route that does not infer supremum approximation from an L^p statement or assume that one fixed low-order IGN is universal. It allows explicit identifiers, sufficient width/heads, output slots and a finite-group averaging wrapper. These resources are part of the theorem, not an assertion about an unchanged small checkpoint.

Lemma 37.1 (Exact finite-temperature gathering)**.** *Let $x_1, \dots, x_N \in \mathbb{R}^d$ be labeled input-token features. Give token j the identifier e_j. A constant query for head i and identifier keys can produce score $\log 2$ for token i and zero for every other token. If $S = \sum_j x_j$, this head returns*

$$h_i = \frac{S + x_i}{N + 1}.$$

A uniform-score head returns $\bar{x} = S/N$. Consequently,

$$x_i = (N + 1)h_i - N\bar{x}.$$

Thus $N+1$ heads followed by a linear map gather all labeled input features exactly at each output slot. Keys are restricted to input tokens by a fixed mask; the query constants and identifier coordinates are explicit inputs to the construction.

Proof. Exponentiating the scores gives one weight two and $N - 1$ weights one, with denominator $N + 1$. The displayed identities follow by substitution. Biases or a constant input channel supply the constant queries. No infinite-temperature limit or discontinuous hard selection is used. The real-arithmetic identity is formalized separately; floating-point execution has its own numerical error. □

Lemma 37.2 (Constructive compact-domain ReLU approximation)**.** *For a compact $K \subset \mathbb{R}^d$, every continuous $F : K \to \mathbb{R}^q$ can be uniformly approximated by a finite ReLU network.*

Proof. The empty domain is trivial. Otherwise use the sup norm and uniform continuity to choose $\delta > 0$ so that $\|F(x) - F(y)\|_\infty < \epsilon/2$ whenever $x, y \in K$ and $\|x - y\|_\infty < \delta$. Choose a finite $\delta/2$-net $c_1, \dots, c_M$ in K. Define

$$w_i(x) = \max\{0, 1 - \|x - c_i\|_\infty/\delta\}, \quad S(x) = \sum_i w_i(x), \quad A(x) = \sum_i F(c_i)w_i(x).$$

Absolute value and finite maximum have exact finite ReLU realizations. On K, $1/2 \le S(x) \le M$. The normalized combination $A(x)/S(x)$ is a convex combination of values at centers within δ of x, so its error is below $\epsilon/2$.

It remains to approximate division without assuming it is a ReLU primitive. The reciprocal on $[1/2, M]$ has arbitrarily accurate finite piecewise-linear interpolants. Any one-dimensional continuous piecewise-linear function is an affine term plus a sum of hinge functions: its slope changes give the ReLU coefficients. Products on a bounded rectangle can likewise be approximated using

$$uv = \tfrac{1}{4}((u + v)^2 - (u - v)^2).$$

Linear interpolation of z^2 on a mesh of width at most h has error at most $h^2/4$, since the interpolation error on $[a, b]$ equals $(z - a)(b - z)$. The resulting product error is at most $h^2/8$. Because A and the reciprocal interpolant are bounded, choose reciprocal and product errors so that their combined output error is below $\epsilon/2$. Composition gives a finite ReLU network with total uniform error below ϵ. This is an existence construction; the net size and network cost can be large. □

Theorem 37.3 (Uniform graph-to-graph approximation with declared resources)**.** *Let $\mathcal{G}$ be a compact bounded-size graph domain with a continuous injective finite-dimensional labeled encoding E. Let the output use a fixed finite set of typed slots and a continuous vector-valued target F. A graph-token architecture that contains the gathering construction, sufficiently expressive shared ReLU subnetworks, and explicit output-slot identifiers uniformly approximates F after encoding. If a finite relabeling group H preserves the domain, acts by coordinate permutations on inputs and outputs, and F is equivariant, the finite Reynolds wrapper makes the approximator exactly equivariant without increasing its uniform error.*

Proof. A continuous injection from compact $\mathcal{G}$ into a Hausdorff vector space is a homeomorphism onto its image. Therefore $F \circ E^{-1}$ is continuous on the compact encoded domain. Exact gathering supplies that full encoded vector to each output slot. Apply the preceding ReLU construction. Distinct output-slot identifiers form finitely many separated cells, so a shared network can select the appropriate component. More concretely, for a bounded scalar $|z| \le B$ and a one-hot bit $b \in \{0, 1\}$,

$$\mathrm{ReLU}(z + Bb) - \mathrm{ReLU}(z - Bb) - Bb = bz.$$

This gives exact slot selection on the bounded approximation range.

For the equivariant claim, define

$$A_H(x) = |H|^{-1} \sum_{h \in H} \rho_{\mathrm{out}}(h^{-1}) A(\rho_{\mathrm{in}}(h)x).$$

Reindex the sum by $k = hg$ to show $A_H(\rho_{\mathrm{in}}(g)x) = \rho_{\mathrm{out}}(g)A_H(x)$. Since permutation actions preserve the sup norm and F is equivariant, each transformed summand differs from $F(x)$ by less than ϵ. The triangle inequality gives the same bound for the average. The wrapper is a finite parallel composition and sum of the preceding networks. It can require $|H|$ branches, factorial in the number of vertices. No efficient fixed-width universality is inferred. □

This is a uniform result, not merely L^p. It proves a capability of the explicitly extended family. A particular ToricGT/Mora-3 configuration inherits it only when its masks, identifier semantics, normalization and bypasses actually preserve the construction. The already published bypass lemma handles optional tropical or quantum branches; replacing the only ordinary path by an unrelated operator needs another proof. This theorem does not need a claim that every fixed-order k-IGN is universal.

A separate external uniform route is Alberti–Dern–Thesing–Kutyniok's Sumformer theorem and its Transformer corollary [S1]. Their continuous multisymmetric-polynomial proof is relevant, with its latent dimension and architectural assumptions. It is not Yun's theorem with a silently changed norm. Extending a graph-defined target to all token permutations additionally requires compatibility on orbit intersections and stabilizers. Relabeling graph vertices and permuting token rows are different actions unless the tokenizer establishes their relation.

37.3 Finite interpolation, circuits and unbounded computation

The source's finite-domain claim has an elementary exact realization. For finitely many distinct encoded inputs x_i, choose δ smaller than their minimum pairwise sup distance. The ReLU hats

$$\phi_i(x) = \max\{0, 1 - \|x - x_i\|_\infty / \delta\}$$

satisfy $\phi_i(x_j) = \mathbf{1}_{i=j}$. Hence $\sum_i y_i \phi_i(x)$ interpolates arbitrary prescribed finite output vectors exactly. Exact gathering makes this a labeled graph-token construction. For an orbit-closed finite domain and an equivariant target, Reynolds averaging preserves those exact values. If an input has a nontrivial automorphism, an equivariant target must be fixed by the induced output action. Merely choosing an isomorphic output representative does not automatically satisfy that requirement.

The source's threshold-circuit construction is nonuniform and supplies gate tokens, wire identifiers and growing resources. Those premises are retained. The executable circuit check compares shared affine/ReLU gates with an independent Boolean interpreter on randomly generated acyclic circuits and explicit gate interventions. It proves neither that a trained checkpoint discovered the interpreter

nor that general circuit satisfiability is efficiently solved. A threshold on an integer count can be placed at a half-integer margin; an approximation smaller than half a count preserves its Boolean value. A finite-precision softmax implementation must demonstrate that margin, not simply claim exact hard selection.

The previous supplement's persistent-graph Turing simulation remains a separate constructive theorem with growing memory and shared local rules. The source itself distinguishes bounded graph universality from unbounded uniform computation. Likewise, a lower bound on a restricted hard-attention architecture does not transfer to a larger architecture solely because the larger class contains the restricted one. A valid transfer needs a reduction or an upper-bound simulation with the same resources and identifiers.

37.4 Dependency repairs and explicit counterexamples

Fixed tensor order. The dedicated source's uniform theorem invokes universality of a fixed k-IGN. The cited general invariant/equivariant universality results allow sufficiently high intermediate tensor orders [S2,S3]. The executable rook/Shrikhande example has identical ordered-pair WL refinement histograms but different numbers of four-cliques, eight and zero. This exposes the low-order dependency, not an upper bound on every raw-identifier TokenGT. The repaired uniform proof above uses explicit gathering and symmetrization; the alternative IGN route must allow the tensor orders required by its universality theorem.

Tropical output versus hard routing. A coordinatewise max-plus head computes $\max_j(s_j+v_{jc})$, so different coordinates can have different winners. It is not generally the value vector of $\arg\max_j s_j$. Separately, hard routing among values can be discontinuous: scores $(q,-q)$ and values $(1,-1)$ give opposite one-sided limits at zero. The routing potential $\max(q,-q)$ remains continuous. A toric support-function conclusion therefore needs continuity of the output under study, not just piecewise formulas or rational routing hyperplanes.

Planar region counts. For convex polygons, every oriented edge normal of a Minkowski sum belongs to the union of summand edge normals. Thus $v(P_1+\cdots+P_H)\le\sum_h v(P_h)$. For two planar N-gons the bound is $2N$, not an attainable quadratic asymptotic. This contradicts the source's unqualified $\Omega(N^H)$ clause at $d=H=2$. Finite integer polygon checks are included, but the asymptotic correction follows from the normal-fan argument. The safe general finite bound is the product of summand vertex counts; stronger dimension-sensitive rates require a correctly stated theorem. Subsequent depth/lower-bound claims that use the invalid clause remain qualified rather than silently imported.

Finite-temperature curvature. For $P_\tau(s)=\tau\log\sum_j e^{s_j/\tau}$,

$$\nabla^2 P_\tau=\tau^{-1}(\operatorname{diag}p-pp^T).$$

If a unique winner has margin δ, set $a=(N-1)e^{-\delta/\tau}$. The mass outside it is at most $a/(1+a)$. Since the covariance matrix is positive semidefinite,

$$\|\nabla^2 P_\tau\|_2\le\operatorname{tr}(\nabla^2 P_\tau)\le 2a/\tau.$$

The source's corresponding formula without the factor two fails already for two logits: the nonzero eigenvalue is $2p(1-p)/\tau$. Taking $p=3/4$ gives $3/(8\tau)>1/(3\tau)$, whereas $e^{-\delta/\tau}=1/3$. The rational

inequality is checked in Lean; the exponential parameterization and covariance derivation are written mathematics.

Probability and decoding. Random-order joint total variation is controlled by a sum of uniform conditional total-variation errors through sequential coupling. An L^p approximation to scores or conditional means does not establish that premise. Continuous-output laws require a declared dominating measure or another suitable probability metric. A bounded grammar with a real-valued threshold rule can be discontinuous; bounded graph size alone supplies neither continuity nor terminating, uniquely selected normal forms.

37.5 How these results connect to the physical program

The uniform theorem establishes a representational baseline; it does not identify a physical law from insufficient observations or guarantee out-of-distribution learning. The surface example uses a fixed verified cellulation, a clean dessin, a CSS spin code, a classical attention controller and actual quantum-state simulation. Max-plus witnesses describe the controller, not coherent quantum path amplitudes. The compound map has a well-defined interpretation: observed classical features select rotation parameters, then a valid quantum circuit evolves a prepared state. No free access to an unknown state's exact expectations is assumed.

Approximation errors can be divided among feature extraction, attention/control, state preparation, gates, measurement and physical model discrepancy. Relation and BGG/chain-map certificates check the algebraic layer appropriate to each component. They do not turn a finite surface code into a continuum field-theory proof or a GRH spectral identity. The same distinctions govern measured-data compression: the graph, model, residual representation and selection procedure must be charged, and a compressible physical process is not thereby an intentional message.

37.6 Additional primary references

S1 Alberti, Dern, Thesing and Kutyniok. *Sumformer: Universal Approximation for Efficient Transformers.* PMLR 221 (2023), 72–86. Continuous proof and Theorems 4.1/4.3 inspected. https://proceedings.mlr.press/v221/alberti23a.html

S2 Maron, Fetaya, Segol and Lipman. *On the Universality of Invariant Networks.* ICML 2019. Higher-order tensorization is an explicit premise. https://proceedings.mlr.press/v97/maron19a.html

S3 Keriven and Peyré. *Universal Invariant and Equivariant Graph Neural Networks.* Version 2, 2019. https://arxiv.org/abs/1905.04943v2

37.7 A complete finite graph-function theorem in Lean

The finite-domain result now has a structural formalization rather than only a table of sampled successes. `FiniteGraphNetworks.lean` defines a Boolean circuit syntax with input, constant, negation, conjunction and disjunction constructors. The same syntax has two interpreters: Boolean evaluation and a real-valued affine/ReLU evaluation. Induction on the complete syntax proves

$$\mathrm{neural}(\mathrm{bit} \circ x, C) = \mathrm{bit}(\mathrm{discrete}(x, C))$$

for every finite circuit C and every Boolean input x. The real-valued operations are $1 - u$, $\max(0, u + v - 1)$ and $u + v - \max(0, u + v - 1)$.

For each Boolean pattern $s \in \{0,1\}^n$, a conjunction of matching literals constructs an indicator circuit I_s. Its neural interpretation is exactly one on s and zero on every other Boolean input. Therefore, for an arbitrary real-valued finite table F,

$$N_F(x) = \sum_{s \in \{0,1\}^n} F(s) I_s(x)$$

satisfies $N_F(x) = F(x)$ at every Boolean input. Applying this construction independently to each output coordinate gives the formally checked theorem

$$\forall F : \{0,1\}^n \to \mathbb{R}^m, \quad N_F(x) = F(x) \quad \text{for every } x \in \{0,1\}^n.$$

A fixed finite graph encoding places adjacency, type and mask bits into this input vector and output graph attributes into the coordinate vector. Thus this is an exact finite graph-to-graph representation theorem for that declared encoding. The theorem includes every table, not only those enumerated by a test script.

The construction's cost is explicit: it contains a pattern circuit for every Boolean input and can be exponential in n. A directed acyclic circuit can be unfolded into the formal syntax tree, with the corresponding duplication cost. This proves representability and correctness, not efficient identification of the table from examples. The continuous compact-domain theorem uses the separate written interpolation argument and its selected formal identities; the finite theorem is not substituted for it.

The formal module prints all seven theorem types and foundational dependencies. Its successful elaboration contains no `sorryAx`, no newly declared axiom and no `native_decide` trust shortcut. These are the relevant facts when evaluating the proof, independently of the rhetoric surrounding neural computation.

Chapter 38

Graph languages and computational mechanisms

38.1 From function approximation to graph computation: the precise claim

Evidence status. This extension develops mathematical arguments and an implementation audit, not new trained-model results. The finite contract and probability modules have their own Lean receipts; the propositions in this language extension are ordinary written proofs unless explicitly identified otherwise. The architecture families are source-pinned in the reference register. A research manuscript, an implemented module, a tested execution and a trained capability are distinct evidence objects.

The strongest defensible position is not that an unrestricted approximation theorem makes every trained network a universal scientist. It is that graph-valued neural computation can represent syntax, execute mechanisms, search derivations, and participate in independently checked semantic transformations. ToricGT, TropicalGT and Mora-3 make several of those objects explicit: vertex and edge tokens, typed intermediate graphs, semiring operations, continuous latent trajectories, branching forests, retrieval, and certificate-aware objectives. These structures permit more informative mechanisms and tests than treating the system solely as a next-word predictor. Whether they improve a particular learned model must still be measured.

Five questions should be kept separate. *Representability* asks whether a function belongs to an architecture's approximation class. *Computability* asks what a uniform machine can do with scalable time and memory. *Complexity* charges the resources required. *Learnability* concerns the data and algorithm that select parameters. *Semantic competence* concerns correctness against an interpretation, including permitted interventions and compositions. No one of these is a synonym for the other four. A construction may establish representability while being exponentially large; a universal interpreter may be impossible to identify from the available observations; a compact network may generalize well on one distribution without deciding an unbounded language.

The claim that a system is "only learning a function" therefore supplies no computational obstruction. Parsers, circuit interpreters, theorem checkers, simulators and compilers all compute functions, partial functions or conditional distributions. The substantive questions concern which function, what representation, what algorithm, what evidence and what cost. Conversely, describing any arbitrary input-output fit as understanding is too permissive. Here operational understanding means successful, independently evaluated use of a specified semantics across compositions, counterfactuals, equivalent presentations and held-out structures. It does not assert consciousness or

identify a neural implementation with human development.

38.1.1 What the source audit establishes

The graph-universality discussion in *Tropical Quivers of Architectures*, at commit `2b6ab79`, explicitly invokes an L^p sequence approximation theorem on bounded compact graph-token domains [P1,R1]. The ToricGT manuscript at `a5199bc` instead writes a supremum-norm conclusion while citing the same result [P2]. That upgrade requires another argument: it is not supplied merely by citing the L^p theorem. The following treatment retains the positive representational program while making its necessary interfaces explicit.

The distinction is elementary. For $k \geq 2$, put

$$b_k(x) = \max\{0, 1 - k|x - 1/2|\}, \qquad x \in [0, 1].$$

Then $\sup_x b_k(x) = 1$, while direct integration gives

$$\int_0^1 b_k(x)^p \, dx = \frac{2}{k(p+1)} \longrightarrow 0.$$

Thus vanishing average approximation error does not imply uniformly correct decisions. This is especially relevant when graph incidences, exact symmetries or grammatical acceptance live on a thin subset of an ambient token space. Ambient Lebesgue L^p control does not automatically control a distribution supported on a lower-dimensional encoded graph family. The approximation measure must be stated, not silently changed after tokenization.

Yun and collaborators establish the cited sequence results in their specified L^p metric; TokenGT supplies a graph-structural expressivity construction, including higher-order constructions with their own token and head costs [R1,R2]. Neither citation by itself establishes that every fixed node-and-edge implementation, arbitrary identifier convention, final normalization and graph decoder has uniform graph-to-graph universality. An architecture theorem must reach the implemented boundary rather than stop at an informal resemblance.

38.1.2 A graph-to-graph approximation contract

Let $\mathcal{G}_{N,M}$ be a finite disjoint union of graph cells with at most N vertices and M edges, finitely many discrete types, and attributes in compact Euclidean sets. Each cell fixes incidence, types and padding masks. Let H be the finite relabeling group appropriate to the chosen domain. Input and output representations have actions ρ_{in} and ρ_{out}; these need not have the same dimensions. A faithful tokenizer

$$E : \mathcal{G}_{N,M} \longrightarrow K \subset \mathbb{R}^d$$

must be continuous and injective on the intended information, with continuous inverse on its image. For compact source and Hausdorff target, a continuous injection is a homeomorphism onto its image. A model receiving only a quotient representation cannot distinguish graphs collapsed by that quotient.

A target F is equivariant when

$$F(hG) = \rho_{\text{out}}(h)F(G), \qquad E(hG) = \rho_{\text{in}}(h)E(G).$$

Only relabelings with a defined action on the target are included. Graph generation with newly allocated output vertices needs an output-slot convention, a quotient-valued target, or an explicit decoder; reusing an input permutation matrix for a differently sized output is not a definition.

Likewise, graph relabeling is a subgroup of all possible token permutations. Extending an equivariant function to a larger permutation orbit requires consistency on orbit intersections and stabilizers.

Proposition 38.1 (Conditional transfer through faithful interfaces)**.** *Suppose E is a homeomorphism onto K, F is continuous, and an architecture class $\mathcal{A}$ is uniformly dense for $\widetilde{F} = F \circ E^{-1}$ on K. Then for every $\epsilon > 0$ some $A \in \mathcal{A}$ satisfies*

$$\sup_{G \in \mathcal{G}_{N,M}} \|A(E(G)) - F(G)\| < \epsilon.$$

If an output reconstruction D is L_D-Lipschitz on a neighborhood containing both the exact and approximated outputs, the reconstructed error is at most $L_D\epsilon$.

Proof. The first inequality is precisely the assumed approximation inequality after substituting $E(G) \in K$. Apply the Lipschitz inequality for D pointwise for the second conclusion. A decoder defined only on the valid graph manifold cannot be applied to an off-manifold approximation without an additional extension or projection. □

This proposition is intentionally explicit about where the universality hypothesis enters. It transfers a valid theorem; it does not manufacture a stronger theorem for its source class. An L^p variant holds for a declared graph probability measure and its pushforward under E when the approximating class is dense for that pushforward measure. The original ambient L^p result is not automatically that statement.

There is also a constructive uniform existence result that avoids the metric ambiguity. It is useful as a corrected theoretical baseline, although its direct implementation can be expensive.

Proposition 38.2 (Finite-group symmetrization preserves uniform approximation)**.** *Let K be compact and H-invariant. Assume H acts linearly on input and output spaces and the output norm is invariant under the action. Let $F : K \to \mathbb{R}^q$ be continuous and equivariant. If N uniformly approximates F within ϵ, then*

$$N_H(x) = \frac{1}{|H|} \sum_{h \in H} \rho_{\text{out}}(h^{-1}) N(\rho_{\text{in}}(h)x)$$

is equivariant and uniformly approximates F within ϵ.

Proof. For $g \in H$, substitute $k = hg$ in the sum defining $N_H(\rho_{\text{in}}(g)x)$. Since $h^{-1} = gk^{-1}$, the output factor $\rho_{\text{out}}(g)$ comes outside the sum, proving equivariance. Each summand's error from $F(x)$ has norm below ϵ by equivariance of F and invariance of the norm. The triangle inequality bounds their average by ϵ. □

Ordinary continuous nonpolynomial neural approximation results supply uniformly dense base classes under their usual hypotheses [R3]. The displayed wrapper therefore gives a mathematically sound uniformly approximating equivariant neural graph family with faithful finite-dimensional interfaces. It does not prove the same property for a particular fixed-width TokenGT checkpoint. Enumerating $|H|$ relabelings can be factorial in graph size. Replacing the wrapper by an efficient equivariant architecture is a separate constructive task.

For discrete finite domains there is an even more direct existence proof. Suppose $x_1, \dots, x_s$ are distinct encoded graphs. Choose $\delta > 0$ below their minimum pairwise ℓ_∞ distance and define

$$\phi_i(x) = \text{ReLU}\left(1 - \frac{\|x - x_i\|_\infty}{\delta}\right).$$

Absolute values and finite maxima are realizable by ReLU networks using $|z| = \mathrm{ReLU}(z)+\mathrm{ReLU}(-z)$ and $\max(a,b) = \mathrm{ReLU}(a-b)+b$. Thus $N(x) = \sum_i y_i \phi_i(x)$ exactly realizes any prescribed finite graph table $x_i \mapsto y_i$. This proves finite representability, not an efficient algorithm or an explanation of generalization. Its explicit size dependence is precisely why a universal-approximation slogan cannot substitute for algorithm design.

38.1.3 Decoding margins and discontinuous grammar rules

A discrete syntax decision requires a margin or an exact interpreter. If the correct coordinate of a target score vector exceeds every competitor by at least $\gamma > 0$, a uniform coordinate error below $\gamma/2$ preserves its argmax. An average error cannot guarantee every decision. For a finite atomic evaluation distribution with minimum atom mass $\mu_{\min} > 0$, an $L^p(\mu)$ error ϵ bounds the maximum error by $\epsilon \mu_{\min}^{-1/p}$; this follows because each atom contributes at least its mass times its error to the pth moment. This finite-domain observation does not repair an unrelated ambient Lebesgue guarantee.

A bounded grammar need not induce a continuous attribute map. A rule choosing one edge label when a real attribute $x \geq 0$ and another when $x < 0$ is discontinuous at zero. One may instead retain the branch as discrete state, restrict to a margin-separated domain, approximate a smooth stochastic decision, or measure error under a distribution that controls mass near the boundary. The choice changes the claim. A canonical normal form also needs termination and unambiguous selection; bounded input size alone does not prove those properties for arbitrary rewrite systems.

The pure ReLU/finite-max portion of a model is piecewise affine. Ordinary softmax, LayerNorm, GELU and the logarithmic transformations in some tropical attention implementations are not collectively an exact piecewise-affine map. A local affine fit to such a model must retain its approximation residual and operating region. The geometry of a tropical decision component is valuable without pretending it exactly describes every operation in the complete network.

38.1.4 Inheritance by ToricGT, TropicalGT, Mora-3 and hybrid variants

Proposition 38.3 (A genuine bypass preserves a capability class)**.** *Suppose every function of a base architecture $\mathcal{A}$ is realizable in an extended architecture $\mathcal{B}$ by disabling additional branches without changing the base interface, normalization, masks, precision or accessible state. Then every representational capability of $\mathcal{A}$ is also a capability of $\mathcal{B}$. If the embedding approximates rather than exactly preserves the base computation, its error must be added to the base approximation error.*

Proof. Use the parameter embedding specified by the hypothesis for each base realization. For approximate embedding, apply the triangle inequality. No property of the additional branch is needed beyond the stated embedding. □

This is the appropriate inheritance argument for optional tropical heads, toric features, quantum features, GFlowNet policy heads and forest-generation modules. The presence of one ordinary attention head or one ReLU layer is not by itself the hypothesis: the full required computation must survive. Replacing all attention by a different operator, forcing a lossy bottleneck, normalizing away an essential register, fixing a finite output vocabulary, or changing a causal mask can invalidate the embedding.

The inspected ToricGT implementation combines ordinary and tropical layers and provides node and edge readouts [P2]. Its tokenizer uses learned slot-indexed node and edge embeddings and slot-indexed toric features. Relabeling the graph while regenerating those slot identifiers need not equal transporting an already tokenized representation. A graph-equivariance claim therefore needs

a declared identifier action and a relabeling test, or an invariant construction that removes that dependence. This is an interface obligation, not a proof that graph-token computation is impossible.

Mora-3's inspected TokenGT boundary takes explicit orthonormal identifiers and offers ordinary and tropical attention variants [P4]. Its current encoder rejects a supplied conditional attention mask. Consequently, a theorem requiring a particular causal or random-order conditional decoder is not already an implementation theorem for that encoder. The declared identifier width also bounds how many orthonormal identifiers that interface accepts. Scaling the graph capacity requires scaling or changing the interface, with its costs recorded.

The constructive generalization proposed here is a persistent typed graph state, a shared neural proposal operator, exact or certified local update rules, explicit memory allocation, and an optional quantum or hybrid feature branch. It retains the existing models as proposal and feature mechanisms rather than replacing them with a second unrelated language model. These additions create concrete capabilities to investigate: structured output, relation-preserving rewrite selection, semantic circuit execution, controlled memory, and diverse derivation sampling. They do not create a new class of noncomputable functions.

38.2 Graph grammars, uniform computation and the language hierarchy

A graph grammar specifies more than an adjacency picture. A rule needs vertex and edge types, ordered ports when order matters, a matching condition, a replacement boundary, and a statement about attributes and side conditions. A derivation records which occurrence was replaced. The denotation of a graph may be a string language, a function, a linear operator, a probabilistic kernel, or a physical model. Equal graph size or equal topology does not imply equal denotation.

For language processing, a useful state contains the input spans, syntactic objects, binding and dependency edges, an explicit stack or workspace, and a derivation history. A graph-to-graph model can propose the next derivation, emit an annotated parse, answer an intervention on a dependency, or reconstruct a sentence from a structured object. These are different targets. The Chomsky hierarchy classifies particular string-language formalisms; an arbitrary typed graph grammar needs a specified string yield and restrictions before it acquires one of those classifications.

38.2.1 A uniform graph executor can simulate a Turing machine

A bounded approximation statement has quantifiers of the form

$$\forall N, T, \epsilon \quad \exists \theta_{N,T,\epsilon}.$$

It does not imply that one fixed parameter choice computes an unbounded task. Turing completeness requires a uniform computational mechanism, with a clear source of unbounded time or storage. A constructive extension makes that mechanism explicit.

Theorem 38.4 (Turing simulation by a persistent local graph executor)**.** *Fix a deterministic single-tape Turing machine with finite control states and finite tape alphabet. There is a finite typed local graph-rule set, independent of input length and run duration, whose recurrent execution simulates each finite prefix of the machine's computation exactly. The graph uses a directed tape path with one distinguished head/control label and may allocate a blank vertex at either boundary. After t transitions on an input of length n, it requires $O(n+t)$ tape vertices. Its finite local transition selector can be represented exactly by a finite ReLU network on one-hot finite-alphabet inputs.*

Proof. Represent every allocated tape square by a vertex carrying its symbol. Store predecessor and successor as typed directed edges and mark exactly one vertex with the current control state. For each transition $(q, a) \mapsto (q', b, d)$, create a rule that matches the unique head vertex of state q and symbol a, writes b, removes its head mark, and transfers state q' to the neighbor in direction d. When that neighbor does not exist, the boundary rule first allocates a blank vertex and links it. A stay-put transition updates the same vertex. Halting states have no outgoing transition.

Induction on steps proves equality of the represented tape, control state and head position with the Turing computation: the initial encoding agrees, and the unique applicable transition performs exactly the prescribed write and move. At most one new square is allocated per transition, giving the size bound. The transition table has finitely many inputs. A finite Boolean indicator network on their one-hot encoding selects the appropriate rule exactly, using conjunctions, negations and finite sums. Its parameters depend on the machine, not on n or t. □

This is an architecture-generalization theorem for an explicit graph executor, not a claim that an existing ToricGT or Mora-3 checkpoint has learned a universal interpreter. The graph memory and exact update mechanism are part of the machine. A synchronous implementation scanning all allocated vertices costs $O(t(n + t))$ elementary local checks over t steps; an event-driven implementation can avoid those scans if it provides the specified constant-cost local addressing. Neither access model is free. Floating approximations require certified restoration to the finite symbolic alphabet or a robust error analysis.

A fixed finite-precision system with a fixed total memory budget has finitely many configurations. Its unbounded input-stream behavior is therefore finite-state when no additional storage is available. A fixed learned parameter vector coupled to a growing graph or external tape is a different machine. The contradiction disappears once parameters and working memory are distinguished. Similarly, arbitrarily precise real numbers must not be treated as an uncharged storage oracle.

Pérez, Barceló and Marinkovic provide an external Turing-completeness theorem for their specified attention model [R4]. Merrill and Sabharwal analyze intermediate decoding as a resource under explicit precision, masking and normalization assumptions [R5]. Those results motivate architectural embeddings and resource accounting; they are not automatically inherited by every softmax, tropical, graph-token or quantized implementation.

38.2.2 The hierarchy as a collection of executable resource contracts

For regular languages, a finite-state transition table can be embedded in a shared finite update. The graph state may contain a single control node with a one-hot state. For context-free recognition, a pushdown construction adds a stack path with push, pop and top-symbol tests; acceptance must use the appropriate deterministic or nondeterministic model rather than assume all context-free languages have deterministic parsers. A finite neural rule selector can encode each finite transition relation, while a search policy handles nondeterministic alternatives.

One-counter Dyck recognition is not a complete model of arbitrary context-free parsing. Two bracket types require the order of unmatched opening types, not just their number. A persistent stack makes that information explicit. A bounded stack can handle a finite depth range; unbounded recognition requires growing stack memory. Length extrapolation should therefore be tested jointly with nesting depth, bracket-type interactions and the actual state invariant.

Context-sensitive languages correspond to the usual noncontracting-grammar/nondeterministic linear-bounded-automaton characterization, with the conventional treatment of the empty word. The memory allowance is proportional to input length. This does not identify all context-sensitive languages with deterministic linear space: that determinization issue must not be silently assumed.

Unrestricted grammars correspond to recursively enumerable languages. Acceptance can take unbounded time, and a nonmember may never be rejected. Turing completeness consequently does not provide a halting oracle or decide every recursively enumerable language.

These classical levels become architecture contracts: finite state; a growing stack; linearly bounded tape with nondeterministic search; or a general growing work graph. The neural parameterization can be shared across lengths in each case. The resource restrictions, not the adjective "neural", determine the relevant inclusion. A TokenGT-style representation makes trees, stacks and tapes native relational objects; it does not repeal their computational costs.

Natural language is not exhausted by membership in one toy string language. Meaning, reference, pragmatic context, acquisition and causal grounding require additional tasks and assumptions. Conversely, the fact that a language faculty is described algebraically does not imply that the algebra cannot be implemented or approximated by a neural mechanism. A useful experiment asks whether a model acquires the intended composition and invariants rather than a shortcut matching the available examples.

38.2.3 Tropical parsing and the cost of a derivation

Weighted parsing illustrates a concrete advantage of semiring-aligned operations. In a suitably normalized finite context-free grammar, an inside dynamic program combines alternative derivations by a sum and combines subderivations by multiplication. Replacing these operations by maximum and addition gives a Viterbi-style best-score computation in log coordinates. A tropical head can therefore implement a declared finite max-plus recurrence, provided its masks, features and reduction actually realize the required operands. The graph must retain traceback choices when a derivation, rather than only its score, is needed.

The distinction between total mass and best path is mathematical. For finite log weights $a_1, \dots, a_m$ and $\tau > 0$,

$$\max_i a_i \le \tau \log \sum_i e^{a_i/\tau} \le \max_i a_i + \tau \log m.$$

The inequalities follow by factoring out the largest exponential and bounding the remaining sum between one and m. A tropical best-path approximation can lose a multiplicity contribution as large as the displayed term. Ambiguous grammars and merged proof histories make that contribution relevant. With signed or complex amplitudes, cancellation introduces another issue and this positive-sum argument does not reconstruct quantum interference.

The proposed architecture can preserve both a tropical best-derivation channel and an ordinary sum-product or log-sum-exp mass channel. A separate provenance structure records their alternatives. This is more informative than forcing every question into one semiring. It permits best-proof search, uncertainty over parses and semantic checking to share graph structure without claiming they are the same operation.

38.3 Circuit interpretation as a testable form of semantic competence

An electrical or logic circuit is an especially useful intermediate target because its syntax and semantics can be independently specified. A netlist has component types, ordered or distinguished ports, wire identities, inputs and outputs. A logic graph should distinguish fan-out from duplicated independent random variables. A physical circuit additionally needs units, constitutive laws,

boundary conditions and operating assumptions. The model's answer must be evaluated against those semantics, not against visual resemblance between circuit diagrams.

Proposition 38.5 (Exact Boolean circuit semantics by shared neural updates)**.** *Let C be a finite acyclic Boolean circuit using NOT, AND and OR gates, with maximum gate depth D. There is a shared finite collection of affine/ReLU gate rules whose depth-wise graph execution computes every gate value exactly on Boolean inputs in D propagation rounds after inputs are installed.*

Proof. For $x, y \in \{0, 1\}$, define

$$\mathrm{AND}(x, y) = \mathrm{ReLU}(x + y - 1), \quad \mathrm{NOT}(x) = 1 - x, \quad \mathrm{OR}(x, y) = x + y - \mathrm{AND}(x, y).$$

Checking the four binary input pairs proves the gate identities. Apply the rule corresponding to each node type once its predecessor values are available. Induction on topological depth gives exactly the conventional circuit interpretation. A type selector over the finite gate vocabulary is itself a finite Boolean mechanism. The graph executor supplies the declared wiring and scheduling. □

The proposition is a constructive semantics result, not a statement that a dense attention layer already implements exact port matching. The needed wiring can be an explicit sparse executor, a hard incidence mask with an exact routing construction, or a learned approximation with a verified margin. The cost includes storing and accessing the graph. Cyclic Boolean circuits require a clocked-state, fixed-point or other semantics; they do not become acyclic merely by drawing them as tokens.

A meaningful circuit-understanding test includes novel compositions of known gates, wire relabelings, interventions that clamp an internal node, fault localization, and equivalence-preserving rewrites. Boolean simulation of a supplied acyclic circuit is different from finding an input that satisfies a circuit, or proving equivalence for every input. The latter tasks can have much greater complexity. A learned heuristic may improve search on a distribution, but a graph encoding does not transform arbitrary satisfiability into a proven polynomial-time procedure.

38.3.1 Electrical networks and conditioning

For a connected passive resistor network with a grounded reference node, positive conductances and specified current injections, the reduced weighted graph Laplacian L is positive definite and the nodal voltages solve $Lv = b$. Indeed, its quadratic form is a sum of conductance-weighted squared voltage differences, including differences from the grounded vertex; connectedness makes that sum positive for a nonzero reduced voltage vector. This argument derives the relevant invertibility rather than assuming it from the word "circuit".

If the domain guarantees $\lambda_{\min}(L) \geq \lambda_0 > 0$, an approximate voltage vector $\hat{v}$ with residual $r = L\hat{v} - b$ satisfies

$$\|\hat{v} - v\|_2 \leq \frac{\|r\|_2}{\lambda_0}.$$

This follows from $\hat{v} - v = L^{-1}r$ and the spectral norm of the inverse. A graph model can propose $\hat{v}$; an independent residual checker and a justified eigenvalue lower bound certify its error. Without a uniform conditioning bound, small changes in components or topology can produce large solution changes, so a compact-domain approximation claim needs a carefully specified domain.

Capacitors and inductors introduce dynamical state; nonlinear devices introduce nonlinear constitutive relations; switching can create discontinuities. These extensions fit the project's dynamical-systems program only after their equations and solution concepts are declared. An algebraic graph grammar may describe assembly, but it does not supply missing physical parameters

or prove stability by itself. In particular, a tropical score selecting a current path is not generally a solution of Kirchhoff's equations: currents combine over paths rather than always taking only the largest one.

38.3.2 Causal interpretation, not a semantic label applied afterward

A structural causal circuit specifies how an intervention replaces a mechanism. The model can be trained or constructed to perform that replacement and compute its consequences. No universal approximator can identify an intervention law from observational data when two causal models induce the same observations and different intervention laws. The obstruction is lack of information, and applies to symbolic and neural methods alike.

A useful criterion therefore separates supplied-model execution, learned-model identification and real-world validation. Correctly executing a supplied circuit is genuine competence at that semantics. Recovering the right circuit from measurements requires another experiment. Claiming that the recovered circuit describes a biological or cosmological mechanism requires independent evidence. The same distinction is maintained for graph grammars learned from physical trajectories.

38.4 ConvexTok, graph incidence and the preservation of information

Three tokenization questions must be separated. Vocabulary construction chooses which byte strings have token identities. Segmentation chooses a path through the available spans for a particular byte string. Graph tokenization represents vertices, edges, roles, ports and attributes. ConvexTok addresses vocabulary construction through an integer-program formulation and a convex relaxation, with feasible rounding and an objective-specific optimality gap [R6]. TokenGT-style incidence encoding addresses the third question. Combining them is meaningful, but they are not interchangeable mechanisms.

For a fixed vocabulary, construct a segmentation DAG whose vertices are byte offsets $0, \dots, L$ and whose edge $i \to j$ denotes a token matching the bytes in that interval. A valid segmentation is a path from zero to L. Additive token costs give the min-plus recurrence

$$d(j) = \min_{i \to j} \{d(i) + c(i, j)\}, \qquad d(0) = 0.$$

The existence of a byte-fallback edge at every offset guarantees coverage. This recurrence is an exact algorithm for the declared fixed vocabulary and objective. It does not make the globally coupled vocabulary-selection problem a collection of independent shortest paths: selecting a token type can enable edges in many examples under one shared vocabulary budget.

A feasible integer vocabulary solution provides an upper bound on a minimization optimum; a valid relaxation provides a lower bound. Their difference certifies a gap for that objective and dataset. It is not a certificate of grammatical correctness, semantic adequacy, graph equivariance, Turing completeness or improved downstream accuracy on every task. In particular, a compression certificate should retain the corpus split, normalization, token inventory, fallback policy and solver settings to which it applies.

Proposition 38.6 (Information-preserving tokenization does not reduce computability)**.** *Let $E : X \to Z$ be a computable encoding with a computable left inverse D, so $D(E(x)) = x$. Every computable target $f : X \to Y$ has the implementation $f \circ D$ on encoded inputs. Conversely, if $E(x) = E(x')$ but $f(x) \neq f(x')$, no deterministic function of $E(x)$ alone realizes f on both inputs.*

Proof. Composition of the supplied algorithms computes $(f \circ D)(E(x)) = f(x)$. For the converse, a function takes the same value on equal inputs, contradicting the two different target values. □

The proposition explains why reversible byte tokenization can change efficiency without changing the underlying computable tasks. It also gives a direct obstruction for lossy graph summaries, dropped edge orientations, truncated attributes or unknown-token collapse. The complexity of decoding still counts. A succinct encoding may require a long expansion, and a tokenizer trained on held-out text can leak information even when it is perfectly reversible.

For the proposed ToricGT/Mora-3 interface, span tokens should retain their byte intervals, modality, node or edge ownership, relation type, port position and provenance. Structural control tokens belong to a typed control channel. Ordinary payload containing the printed spelling of a control token must remain payload, not become an instruction to mutate a graph. When a vocabulary is extended, its new axis, decoder and checkpoint compatibility must be versioned instead of silently altering an old tokenizer certificate.

38.4.1 A precise toric connection, without confusing statistics with meaning

Let A be the directed incidence matrix of the segmentation DAG, with columns indexed by token-occurrence edges. The nonnegative integer vector z of a source-to-sink path satisfies $Az = b$, where b is the fixed source/sink boundary vector under the chosen sign convention. Two paths u, v have $A(u - v) = 0$. Their monomial difference $X^u - X^v$ therefore belongs to the kernel of the monomial map sending an edge variable to the Laurent monomial determined by its incidence column. This is a toric-ideal relation for the declared incidence statistics.

That equality is not a universal semantics theorem. Two paths can share incidence statistics while differing in token counts, attributes, costs or higher-level interpretation. Conversely, two semantically equivalent programs need not have the same incidence vector. The useful role of the toric construction is to describe a controlled algebra of alternate factorizations. Semantic preservation must be supplied by a separate denotation map or rewrite certificate.

The same distinction applies to tropicalization. If all relevant leading coefficients are positive and a valuation convention is fixed, a sum's leading order can be obtained by the corresponding minimum rule without cancellation. For signed or complex coefficients, equal leading orders can cancel. Tokenization alternatives, grammar ambiguity and quantum amplitudes consequently require different algebraic treatments even when their diagrams share a DAG representation.

38.4.2 Prefix measurability and fair accounting

For autoregressive evaluation, a structural feature available at step t must be measurable with respect to the already revealed input, public metadata and declared decoding order. A graph edge extracted using a hidden target token, a future parse annotation or a completed proof can leak the answer. It cannot be excused as geometric side information. Random-order decoding requires masks tied to the revealed set, including any edge features that depend on unrevealed endpoints.

Report byte-normalized likelihood alongside graph validity, semantic accuracy and the size of nonrecoverable side information. Charge a learned vocabulary artifact, externally supplied graph structure, retrieval indices and verifier calls where they are necessary to the claimed system. A deterministic path graph recoverable from the observed sequence has a different information cost from a privileged semantic graph supplied by an oracle. These accounting choices should be predeclared and ablated rather than selected after seeing the score.

38.5 GFlowNets over derivations, programs and complete forests

A language or circuit task often admits many acceptable constructions. A GFlowNet can learn a distribution over completed parses, proof objects, candidate circuits or reasoning forests rather than only select one maximum-score result. The original trajectory-balance theory and its continuous/hybrid extensions specify the conditions for that distributional interpretation [R7,R8]. A loss bearing the same name is not sufficient: its states, actions, backward probabilities, terminal objects and reference measure must match the theorem.

Consider a finite construction DAG with source s_0, terminal set $\mathcal{X}$, normalized forward transitions $P_F(s' \mid s)$ and normalized backward transitions $P_B(s \mid s')$ over the actual parents of each non-source state. Assume every supported path terminates and every backward path reaches the source. For a complete path $\tau : s_0 \rightsquigarrow x$, define products $P_F(\tau)$ and $P_B(\tau \mid x)$. Positive rewards $R(x)$ and a positive scalar Z satisfy exact trajectory balance when

$$ZP_F(\tau) = R(x)P_B(\tau \mid x)$$

for every complete supported trajectory, with sufficient support to cover the target objects.

Proposition 38.7 (Terminal law and a uniform residual bound)**.** *Under these assumptions, exact balance implies*

$$P_F(X = x) = \frac{R(x)}{\sum_y R(y)}.$$

More generally, suppose every complete path has a log-ratio residual r_τ satisfying

$$ZP_F(\tau) = R(x)P_B(\tau \mid x)e^{r_\tau}, \qquad |r_\tau| \le \epsilon.$$

Then, writing $p^*(x) = R(x)/\sum_y R(y)$,

$$e^{-2\epsilon} \le \frac{P_F(X = x)}{p^*(x)} \le e^{2\epsilon}, \qquad \|P_F - p^*\|_{\mathrm{TV}} \le \min\left\{1, \frac{e^{2\epsilon} - 1}{2}\right\}.$$

Proof. For each terminal x, the backward path probabilities sum to one. Summing the residual identity over paths ending at x bounds $ZP_F(X = x)$ between $e^{-\epsilon}R(x)$ and $e^{\epsilon}R(x)$. Summing again over x bounds $Z/\sum_y R(y)$ between the same two factors. Dividing gives the likelihood-ratio bounds. The total-variation bound follows from $\frac{1}{2}\sum_x p^*(x)|P_F(x)/p^*(x) - 1|$ and the upper bound on the absolute ratio error. Setting $\epsilon = 0$ gives the exact law. □

The uniform residual hypothesis is deliberately stronger than a small mean squared loss on sampled training paths. A rare branch can have a large residual while barely affecting an average. Missing support cannot be repaired by a loss evaluated only where the sampler goes. The proposition is a useful certificate target for small enumerated grammars, not a claim that ordinary stochastic optimization establishes it on an exponentially large state space.

The backward probability is evaluated at the *destination*, over its valid parent choices. The inspected ToricGT primitive at `a5199bc`, `src/toricgt/gflownet.py`, evaluates both policy heads on the source-state slice and gathers the forward action index for each [P2]. Without an explicit encoding convention proving that these logits represent the correct destination-parent law, the general terminal-distribution theorem does not follow from this primitive. The extension records this as a source-level verification obligation; it does not modify a model under training or claim that every use of that module is invalid.

Mora-3's continuous GFN primitives explicitly leave normalized policies, valid trajectories and the reward measure to their caller [P4]. That is the right scope for a primitive. A normalized Gaussian density is not a complete graph-generating process. A continuous latent update also cannot silently allocate a discrete topology: node counts, edge choices and stopping require a discrete policy or a separately specified skeleton.

38.5.1 Histories, objects and semantic equivalence classes

Suppose one terminal circuit can be constructed in ten action orders and another in one. Sampling action histories uniformly does not sample circuits uniformly. A GFlowNet's backward kernel can account for multiple routes to one terminal object, but only if the state and merge definitions identify the intended object. If terminals are instead derivation histories, the reward law is over histories. Quotienting afterward generally changes the law by the total probability in each equivalence class.

Semantic merging is stronger still. Two graphs can denote the same Boolean function but have different structures and different future rewrite options. A checker may certify their denotational equivalence while a learned continuation policy remains sensitive to their representation. Merging their latent states requires either a sufficient statistic for the continuation problem or an explicit quotient policy; semantic equivalence alone does not prove those latent states interchangeable.

This motivates a two-level architecture. A typed syntax graph retains identity, binding and construction history; a semantic channel stores certified invariants and can support retrieval or equivalence proposals. GFlowNet rewards can favor correct and diverse terminal objects while the verifier rejects invalid rewrites. The reward should not reward an attractive certificate format rather than the property actually checked.

38.5.2 Forests, flow matching and censoring

A complete forest can be represented by discrete topology T, typed actions A and continuous coordinates Z. A hybrid target requires a counting measure on discrete structure and a declared measure on the continuous coordinates. Variable node counts change the continuous dimension. Densities from different dimensions cannot be compared as though they shared one Euclidean base measure. A topology-conditioned density or another rigorously specified disjoint-union construction is required.

A deterministic flow-matching ODE is not a Gaussian-transition sampler. Its transition is concentrated on the deterministic flow map; attaching arbitrary Gaussian log probabilities to that map does not turn it into a valid continuous GFlowNet. One may distill a qualified stochastic teacher into a conditional flow, or define an SDE with its own density and stopping semantics. Those are distinct objectives and must be evaluated as such.

Mora-3's inspected forest routine is bounded, layer-wise beam selection with explicit proposal and depth budgets [P4]. Its own implementation distinguishes that procedure from an on-policy GFlowNet kernel. A beam-pruned distribution is changed by selection. Terminal leaves, budget-censored leaves and invalid proposals must remain distinguishable. Reporting a truncated branch as a failed proof or as an accepted terminal changes both the scientific target and the probability model.

For a random-order finite decoder, a useful sufficient condition is a uniform conditional total-variation error ϵ_t at every step, for all reachable histories under the comparison. Sequential maximal couplings then give joint total variation at most $\sum_t \epsilon_t$: couple each next token, and bound the probability of any mismatch by the union bound. This does not follow merely from approximating

conditional means or from an unspecified L^p score guarantee. Continuous outputs need density or measure-level assumptions appropriate to their reference space.

38.6 Dynamical expressivity, protected memory and error accumulation

The program's recurrent interpretation is useful: a shared graph update evolves a state containing a current candidate, evidence, memory and unresolved branches. However, the system has several different times. Physical time parameterizes the modeled process. Derivation depth counts reasoning steps. Solver time parameterizes integration of a continuous flow. Training time counts optimizer updates. A one-step solver that proposes an entire forest has not reduced the logical depth of every dependency in that forest to one.

Proposition 38.8 (Finite-horizon simulation error)**.** *Suppose exact states obey $x_{t+1} = F_t(x_t)$, approximate states obey $\hat{x}_{t+1} = \hat{F}_t(\hat{x}_t)$, and on the relevant operating domain F_t is L-Lipschitz while $\|\hat{F}_t(z) - F_t(z)\| \le \delta_t$. If both trajectories stay in that domain, then*

$$e_T \le L^T e_0 + \sum_{t=0}^{T-1} L^{T-1-t}\delta_t, \qquad e_t = \|\hat{x}_t - x_t\|.$$

Proof. Insert and subtract $F_t(\hat{x}_t)$, use the triangle inequality and Lipschitz bound to get $e_{t+1} \le Le_t + \delta_t$, then induct. The domain condition is necessary; a local training estimate does not establish it. □

This connects approximation to the workbench's exact weighted-gain certificates. A positive comparison matrix and weights can establish a contraction rate for a specified block metric, provided the model-to-matrix inequality has been proved or independently justified on the domain. The exact arithmetic certificate verifies the matrix inequality, not the existence of that envelope for a learned dynamical system. A rate $q < 1$ controls accumulated bounded disturbances by $\delta/(1-q)$ after transients, but it also imposes a memory tradeoff.

Proposition 38.9 (Global contraction erases separated initial-state information)**.** *Let F be q-contractive with $0 \le q < 1$, let the initial states lie in a set of diameter D, and let a readout be L_R-Lipschitz. After t autonomous iterations, any two readout values differ by at most $L_R q^t D$.*

Proof. Iterating the contraction inequality gives $\|F^t(x) - F^t(y)\| \le q^t D$. Apply the readout bound. □

A model required to retain a fixed decision margin about arbitrarily old initial information cannot satisfy that globally contractive autonomous description forever. The constructive response is not to abandon stability: separate a stable continuous workspace from protected discrete or noncontractive memory, inject fresh inputs, or state a finite retention horizon. A graph executor naturally supports that separation. Toric phases or recurrent hidden states are not automatically protected memory; their noise, quantization and readout must be analyzed.

Local tropical cells can expose which candidates dominate a decision. Their active margins provide perturbation diagnostics, and a sequence of visited cells gives a symbolic itinerary. Stability of each fitted cell matrix separately does not prove stability under arbitrary switching. A common weighted norm, a justified joint bound or restrictions on allowed itineraries are needed. Likewise, a max-plus largest-path growth rate is not the same as the norm of an ordinary sum over all paths; path multiplicity and phase cancellation remain separate terms in the error budget.

38.7 Quantum and hybrid generalizations

A quantum extension needs a state space, preparation, admissible operations, control, measurement and a resource model. A classical graph may index quantum subsystems, or a quantum register may coherently encode graph labels; these are different constructions. Expectation-valued quantum features followed by a classical graph network give a hybrid model. A classical GPU simulation of that circuit remains classical computation of a quantum model, not evidence of hardware quantum advantage.

The bypass proposition supplies a positive baseline: a hybrid architecture that genuinely retains its classical graph executor can retain its classical language and circuit capabilities while adding quantum features or quantum subroutines. Such an extension can be scientifically interesting without a claim of stronger computability. Conventional effectively specified quantum computation does not provide a halting oracle. Any complexity advantage must be stated for a particular problem, encoding, access model and accuracy, with preparation, shots, error correction and readout included [R9].

Quantum circuit grammars have noncommuting operations, tensor structure and, for measurements, classical outcomes and conditional branches. Unitarians compose as amplitudes; unconditional physical evolutions compose as completely positive trace-preserving channels. A forbidden path in a gentle algebra cannot simply be imposed as a zero composition of trace-preserving channels on a nonzero normalized state space: trace preservation would give trace one, not zero. A zero relation can instead describe a forbidden syntactic route, an amplitude operator relation, or a vanishing trace-decreasing instrument branch, with their respective semantics kept explicit.

For a sequence of channels, a uniform diamond-norm error bound permits a telescoping comparison of composed evolutions. Inserting one approximate channel at a time and using contractivity of pre- and post-composition by channels bounds total channel error by the sum of per-step errors. This statement requires actual channel maps; an arbitrary learned matrix is not a channel merely because it is attached to a quantum-labeled graph. The corresponding measurement distributions are controlled through the declared distinguishability bound.

Finite-shot feature estimation introduces statistical error distinct from approximation and discretization error. A grammar selector making many adaptive decisions must account for the total measurement and selection protocol, not reuse a single-shot confidence level as a guarantee for the entire run. The new finite-search measure bounds are one conservative component when individual event bounds are available. They do not prove those calibration premises.

A proposed quantum GFlowNet similarly needs care. Born probabilities can parameterize a classical transition policy when outcomes and normalization are explicit. A coherent sum of paths is not a classical probability flow over those same paths, because interference can change the terminal law. Phase-sensitive amplitudes cannot be treated as nonnegative rewards without changing the object modeled. The suitable research program compares a classical policy, a quantum-feature policy and a physically specified quantum sampling procedure under matched end-to-end budgets.

38.8 Quiver mutation, cluster structures and the semantics of rewriting

A quiver used as an architecture diagram, an exchange quiver in cluster algebra, and the quiver of an algebra's module category are related only after the additional structure is specified. The workbench implements exact skew-symmetric matrix mutation and positive rational seed exchange. It does not turn an arbitrary learned attention graph into a cluster category. This distinction makes the

connection useful: each interface can be tested rather than supported only by a common drawing.

For a skew-symmetric exchange matrix $B = (b_{ij})$, mutation at k is

$$b'_{ij} = \begin{cases} -b_{ij}, & i = k \text{ or } j = k, \\ b_{ij} + [b_{ik}]_+[b_{kj}]_+ - [-b_{ik}]_+[-b_{kj}]_+, & \text{otherwise}, \end{cases}$$

where $[a]_+ = \max(a, 0)$. In the coefficient-free seed convention used by the checker,

$$x'_k = \frac{\prod_i x_i^{[b_{ik}]_+} + \prod_i x_i^{[-b_{ik}]_+}}{x_k}, \qquad x'_i = x_i \ (i \neq k).$$

These are the declared exchange rules, following the cluster-algebra framework [R10]. The implementation restricts to positive rational seed values and bounded exact integers so division and resource limits are explicit. Mutating twice at the same vertex restores the seed. The labeled A_2 five-step example exchanges the two initial coordinates rather than pretending the labels are unchanged.

For a min-convention valuation and a subtraction-free setting with no leading-term cancellation, seed exchange induces

$$u'_k = \min\left\{ \sum_i [b_{ik}]_+ u_i, \sum_i [-b_{ik}]_+ u_i \right\} - u_k.$$

This gives a precise piecewise-linear relation between cluster combinatorics and tropical updates. It does not prove that every tropical attention operation is a cluster mutation, nor that a phase-bearing quantum seed satisfies a positive-valued valuation rule. A model using this grammar should store the exchange matrix, coefficients or their omission, seed chart, mutation index and invertibility domain as part of the state.

38.8.1 Why mutation does not automatically preserve a module category

Cluster categories organize mutation through additional homological structure; the hereditary cluster-category construction and its tilting interpretation are external results with explicit hypotheses [R11]. Quivers with potential carry still more data: the potential determines relations through cyclic derivatives. Keller–Yang mutation equivalences concern the associated differential graded construction in their stated setting, not a blanket equivalence between every pair of ordinary path-algebra module categories [R12]. Forgetting the potential, differential, grading or quotient changes the claim.

A small example shows why this matters. Start with the path algebra of $1 \to 2 \to 3$ over a field. Mutating the middle exchange vertex gives an oriented triangle. Equipping that triangle with its cyclic potential gives the length-two zero relations. The resulting finite-dimensional gentle algebra has simples with periodic projective resolutions: with a consistent cyclic indexing, $\Omega S_i = S_{i+1}$ and $\Omega^3 S_i = S_i$. It has infinite global dimension, while the starting acyclic path algebra is hereditary. Thus the exchange mutation cannot be interpreted as an automatic ordinary derived equivalence preserving these homological properties. The appropriate dg or cluster-category relation must be stated instead.

The example is also a useful neural-grammar test. A model should be able to predict the mutated exchange matrix while refusing to attach an unjustified “same ordinary derived category” certificate. This is more demanding than producing the right picture: it tests whether the representation retains the category in which a statement is valid.

38.8.2 Gentle surface algebras, line fields and orders

For appropriate triangulations of unpunctured marked surfaces, the associated bound quiver algebra is gentle [R13]. Its relations, not only the underlying quiver, determine the allowed and forbidden path compositions. The workbench's finite checker verifies the quadratic-monomial local conditions and checks that allowed compositions have no directed cycle, which is the relevant finite-dimensionality test for that presentation. It does not classify derived equivalences or silently extend the unpunctured construction to arbitrary punctured surfaces.

Geometric models for gentle derived categories connect indecomposable objects and morphisms to curves and intersections, while graded surface constructions retain additional data such as a line field and boundary stops [R14,R15]. These give useful structured targets for a graph model: arc incidence, grading, permitted intersections, relations and controlled transformations. A bare surface homeomorphism or equal genus does not recover all the categorical data. A learned representation should preserve the information needed by the particular theorem rather than compress everything into one topological number.

Arithmetic orders add an integral layer. A field-linear representation or equivalence can forget torsion, denominators, local completions and lattice classes. An architecture action on modules over an order must specify the base ring, scalar extension, projectivity or torsion-freeness assumptions, and compatibility with the integral structure. The earlier orders chapter supplies these distinctions. A cluster seed containing rational coordinates does not itself define an equivalence of integral orders.

A disciplined architecture grammar can therefore type its transformations by strength: combinatorial quiver rewrite; rational seed change; relation-preserving algebra map; equivalence in a specified derived or dg category; integral lattice-compatible map; or unitary equivalence of analytic operators. Each level has different evidence. Neural proposals can choose candidates across this hierarchy, but promotion to a stronger level requires the corresponding witness.

38.8.3 From categorical equivalence to Hilbert–Pólya: the missing analytic bridge

Derived equivalence preserves categorical structure, not an arbitrarily chosen Euclidean inner product on a matrix realization. Similarity can preserve a finite spectrum without preserving norms; a general algebraic equivalence need not even give that chosen matrix similarity. An analytic Hilbert–Pólya proposal additionally needs a Hilbert space, a densely defined self-adjoint operator and its domain, and an arithmetic spectral or determinant identity with the intended L-function. Approximation by finite self-adjoint matrices alone supplies none of those identities.

The useful role of the graph architecture is to generate and test candidate finite presentations, symmetries, trace identities, integral constraints and compatible refinement maps. A continuum claim needs a justified convergence mode and control of spectral pollution. A GRH claim must keep the character and arithmetic family explicit, not replace all character-dependent information by a pooled spectral pattern. These are concrete proof obligations, not a reason to abandon the research program.

Cryptographic applications require a similar separation. Fast mutation, a complicated orbit, chaotic-looking dynamics, a hard-looking graph or neural universality is not a security reduction. Public invariants can be useful for verification and simultaneously leak structure. The project should retain its reviewed cryptographic constructions and treat learned search or mutation encodings as experimental tools unless an explicit threat model and reduction or credible cryptanalysis support more. No secret keys or private evaluation data belong in the public grammar dataset.

38.9 Complexity: new interfaces do not make resource costs disappear

Graph representations can improve alignment with a task while increasing the raw computation. A node-and-edge tokenization has $N = n + m$ tokens. Dense attention has quadratic pair interactions in N; for a dense graph m is quadratic in n, so that component can grow quartically in the original vertex count before accounting for width. Higher-order graph constructions can require tokens indexed by tuples and incur their corresponding powers of n [R2]. Exact blockwise or ring evaluation can reduce stored intermediates without universally reducing the number of pair interactions.

A sparse circuit executor can instead apply one constant-size gate rule per vertex once predecessor values are available. A semiring shortest-path recurrence has a cost determined by the number of relaxations and edges. These are algorithmic benefits of an appropriate representation and schedule, not consequences of naming a layer tropical. A benchmark must distinguish an efficient algorithm learned by the network from an efficient algorithm provided by an external executor.

The language hierarchy and circuit-complexity hierarchies classify different resources. Regular-language membership does not automatically mean that every chosen parallel representation has trivial constant-depth evaluation. Circuit evaluation, satisfiability, equivalence, parsing, graph matching and optimal rewriting should not be collapsed into one "reasoning" score. A verification problem may be easy while finding a witness remains expensive. A GFlowNet can amortize a distribution-specific search, but neither its objective nor a UAP theorem proves an efficient solver for every difficult instance.

The correct scaling report includes graph size, edge density, attribute precision, grammar size, derivation depth, beam or candidate count, verifier cost, retrieval cost, tokenizer side information, quantum shots and classical simulation cost where applicable. Parameter count alone is inadequate. Approximation accuracy must be stated together with these resources; a family that becomes more accurate only by growing exponentially is still an existence result, not an efficient solution.

38.10 An extended response to anti-neural and restrictive function-learning claims

The response should challenge an actual claim, not a caricature of a research tradition. The May 2023 interview discussed in the existing chapter makes categorical statements about the inability of language models to contribute understanding [R16]. Those statements are different from a technical hypothesis about children's acquisition, a criticism of an experiment, or the claim that fluent output does not identify a causal mechanism. The latter concerns remain legitimate.

There is a particularly important source distinction. Marcolli, Berwick and Chomsky's algebraic syntax–semantics paper explicitly relates transformer attention to its algebraic framework in Section 7 and argues against treating language-model success as a refutation of generative syntax [R17]. It should not be cited as proving an intrinsic incompatibility between attention and algebraic grammar. Its compatibility discussion is a productive starting point for this extension, not an opponent to be contradicted by misquotation.

The positive argument developed here has four parts. First, a function can encode a compositional mechanism: exact Boolean gate execution and local Turing transitions demonstrate this without appeal to an unexplained universal-approximation slogan. Second, graph-valued states can preserve syntax, relations, memory and derivation history rather than compress every task into a fixed bag of words. Third, independently certified rewrites preserve a specified semantics regardless of whether a neural, symbolic or hybrid policy proposes them. Fourth, learning changes the search

distribution and internal representation, which can be assessed through held-out composition and causal interventions rather than only surface fluency.

None of these observations proves that a current trained model uses the intended mechanism. That issue requires experiments. But "it learns a function" is not a proof that it cannot: the supposed distinction disappears when the symbolic competitor also computes a function. The meaningful contrast is between an unconstrained fit and a mechanism with invariants, compositional structure, resource bounds and successful interventions. Neural parameterizations can occupy either side of that contrast.

38.10.1 Algebraic syntax is a design constraint, not an exclusion theorem

A free syntactic constructor can distinguish different derivation trees. A semantic map may deliberately identify several of them. A model should not confuse those levels. In particular, scalar tropical maximum is idempotent, $\max(a, a) = a$, and therefore cannot by itself be an injective representation of every syntactic construction that introduces a new distinct object. A graph node recording the constructor and its typed children can preserve that distinction while a separate tropical channel scores or summarizes it.

Hopf-algebraic extraction, rooted-tree composition, operadic substitution and graph-grammar insertion suggest candidate supervision and consistency checks [R17,R18]. Their algebraic axioms can be evaluated on finite examples, used to reject invalid transformations, or implemented exactly in an external interpreter. Calling an embedding a character requires the relevant homomorphism identities, not a visually appealing latent plot. Calling a decomposition a coproduct requires its algebraic domain and compatibility. This is a route to stronger neural models precisely because it makes failure observable.

The same principle applies to language order and binding. Symmetric pooling can discard word order or the identity of a bound variable. Directed, typed, port-aware graphs can retain them. Graph equivariance should remove arbitrary vertex names while preserving meaningful order and roles. If the target distinguishes two inputs that the encoding identifies, the information-preservation proposition proves that no downstream learner can recover the missing distinction. That is an actionable architectural diagnosis, not a global verdict on neural computation.

38.10.2 What empirical objections establish

The impossible-language debate illustrates why neither side should replace a controlled comparison with a slogan. Kallini and collaborators report differing learnability across constructed language transformations, while Hunter identifies a dependency-matching confound in an important comparison [R19,R20]. The appropriate inference is that the exact transformations and controls matter. Neither study by itself establishes that every neural model has human grammatical biases, or that neural models cannot acquire hierarchical structure.

Existing studies of syntactic generalization and mechanistic structure motivate additional tests, but they do not certify the new architectures [R21,R22]. ToricGT, TropicalGT and Mora-3 must be evaluated on their own implementations and trained checkpoints. A speculative quantum extension cannot inherit empirical language performance from a classical model simply because it inherits a representational lower bound.

A rigorous reply therefore rejects the categorical prohibition and retains the empirical burden. Neural learning is not barred in principle from mechanisms, recursion, counterfactuals or algebraic structure. The stronger positive scientific claim is earned when a model demonstrably uses those structures, survives counterexamples, generalizes under controlled distribution changes and pays its

computational costs. Human developmental adequacy is an additional research question, not the sole criterion for being a useful model or scientific instrument.

38.11 Cosmological graph grammars and falsifiable inference

For the CCC application, distinguish a grammar of candidate signals from a physical law producing them. A model can propose temperature/polarization templates, graph-valued source configurations or candidate cross-aeon encodings. Their relation to a detector still requires a transfer model, nuisance parameters, masks, noise, foregrounds and calibration. A learned decoder cannot infer an intentional message merely because a flexible grammar can describe a pattern.

The present cosmology lab already uses a maximum over its finite template search and a Monte Carlo rank based on full-search null maxima. The new exact rank certificate checks $(e+1)/(B+1)$ and its finite resolution. It does not prove exchangeability of the real and simulated statistics. The new Lean measure module separately proves finite-search union and discrepancy bounds from stated event-level premises. Together they make the logic more inspectable, but they do not convert a synthetic fixture into an observation.

A grammar learned or selected using the same sky data increases the search procedure that must be calibrated. The null pipeline must include the same training, feature selection, model selection and template optimization when those stages affect the reported statistic, or those stages must be frozen independently. Repeatedly adapting the grammar after inspecting a signal and quoting an earlier fixed-search p-value breaks the intended guarantee. A positive detection would concern a specified statistical alternative; attribution to CCC, a preceding aeon or an intentional sender needs additional discriminating predictions.

The language and complexity results help formulate testable alternatives. They explain how a candidate encoder or decoder could be implemented and what resources it needs. They do not establish the cosmological channel's existence. That separation is also what makes the proposal scientifically useful: physical, computational and statistical failures can be localized rather than hidden inside one broad claim of universal inference.

38.12 Evaluation and implementation program

The reference implementation should begin with small exact tasks whose complete semantics can be enumerated: typed bracket stacks, finite local tape transitions, Boolean circuits and interventions, token-DAG segmentation, mutation/involution, gentle relations, and finite GFlowNet construction DAGs. These are executable specification tests. Passing them with a hand-constructed executor is not trained performance and should never appear in a neural benchmark table as such.

A trained comparison should then hold out structure, not merely random records. Keep all relabelings, equivalent normal forms, trajectories from one system, refinement levels, shared simulation initial conditions and derivations of one terminal object in the same split group when they could reveal the held-out target. For grammar induction, hold out combinations and depths while matching vocabulary, dependency distance, noise, rule counts and training exposure. For circuits, separate compositional families, topology, depth and interventions. For physical trajectories, separate systems and parameter regimes while retaining units and discretization metadata.

The comparison set should include a sequence baseline, graph-token baseline, matched tropical/hybrid variant, a strong symbolic or numerical algorithm, and a certificate-aware learned policy where applicable. A quantum-feature comparison must include its classical simulator or efficiently computed feature counterpart when feasible, and report total measurement cost. A new architecture

earns a claimed benefit only on a predeclared metric with matched information and a meaningful budget comparison.

Report semantic success, exact validity, calibration, length/depth extrapolation, equivariance under the declared identifier action, conditional-mask leakage tests, verifier rejection rates and end-to-end cost. GFlowNet reports additionally need terminal-law agreement on enumerable examples, backward-kernel normalization, support coverage, multiplicity treatment and censoring. Forest reports need branch coverage and the distinction between generated topology and a supplied skeleton. A single aggregate accuracy or attractive embedding visualization cannot replace these measurements.

Physical dataset curation should retain native values and axes, experimental versus simulated provenance, licenses, original release identifiers, grouped splits and hashes. The source catalog and any newly generated exact reference fixtures must be distinguished from external data actually redistributed. No new model-training run, legacy corpus replacement or admission of synthetic prose is authorized by these mathematical examples. In particular, the ongoing Mora-3 campaign remains separate from this research-text and reference-test extension.

The implementation audit identifies three immediate tests before claiming theorem-to-code transfer: transported versus regenerated identifier equivariance in ToricGT; destination-conditioned backward probabilities and inverse-action alignment in the GFlowNet caller; and a genuinely implemented causal/random-order decoding interface for any Mora-3 theorem requiring one. These tests turn theoretical qualifications into concrete engineering work. They should be fixed or discharged in an isolated model change with its own regression suite rather than silently altering a running experiment.

38.12.1 Conclusion

The extended program is stronger than a blanket assertion of universality. It supplies explicit constructions for graph computation and circuit semantics; a faithful-interface contract for approximation; a uniform memory-growing executor for computability; a probabilistic contract for diverse derivation sampling; and a verification boundary for accepted rewrites. Toric, tropical, categorical and quantum structures can add useful representations and inductive biases when their semantics and costs are retained.

The resulting response to anti-neural impossibility claims is direct: a function-learning mechanism is not intrinsically excluded from algorithmic, recursive, causal or compositional computation. The response to unqualified pro-neural claims is equally direct: bounded approximation, an architecture label, a training loss and a finite experiment are not interchangeable proofs of unbounded capability, efficient learning or real-world understanding. The research task is to connect those levels with explicit constructions, measurements and certificates. This extension makes those connections testable.

38.13 Analogical transfer is a proposal; deduction is a checked consequence

The inspected Mora-3 analogy head builds condition-only queries and complete-forest memory keys, with a train-only index, group identifiers and a frozen encoder specification. Its own configuration does not claim that a close latent key is a verified directed semantic mapping. This extension retains that distinction. Retrieval proposes a source construction; a typed graph witness checks its mapping; deduction accepts a consequence only when its premises and inference rule are supplied.

Let source and target transition systems be $f : X \to X$ and $g : Y \to Y$, and let $A : X \to Y$ be a proposed analogy. Suppose, on a declared operating region, $\|g(Ax) - A(fx)\| \le \delta$ and g is L-Lipschitz. For $x_{t+1} = f(x_t)$ and $y_{t+1} = g(y_t)$, insertion of $g(Ax_t)$ gives

$$\|y_{t+1} - Ax_{t+1}\| \le L\|y_t - Ax_t\| + \delta.$$

Induction yields the same finite-horizon geometric error sum as the earlier dynamical bound. This is a conditional transfer result: a nearest-neighbor score does not establish the commuting error premise. If a source edge maps to several target steps, replace g by the explicitly specified composite along that target path and charge its accumulated error and cost.

The executable reference checks injective boundary-node typing, complete source-edge coverage, oriented target paths and an explicit allowed relation-word table. Reversing a path, omitting a branch or inventing a relation fails. These are syntactic certificates, not proofs that two measured physical processes obey the same equations. The deduction checker uses finite ground rules with known premises; adding a conclusion by similarity alone is rejected.

The real-data experiment uses embeddings from the actually trained small native encoder to retrieve only training-span windows. The input and memory windows share a declared directed lag-chain semantics; their targets are not included in query features. A validation-selected mixture compares retrieved conditional frequencies with the neural distribution. This is an implemented predictive analogy adapter, not the full native complete-forest head and not a discovered physical equivalence. Bank storage is charged in the two-part code length. An improvement that disappears after storage cost is included is still useful to report, but it is not a compression victory under that accounting.

For cross-system transfer, dimensions and units must first be aligned by a declared transformation fitted without target-test access. Validation must compare the transferred rule with target-only, no-retrieval and unconstrained-nearest baselines. Source and target grouping must exclude shared trajectories, equivalent simulations and post-outcome annotations. Unknown pairs are not automatically negative examples. A successful finite mapping can be retained as a reusable proof or program fragment only at the level actually checked.

Chapter 39

Physical data, coding and the measured study

39.1 What a discovery test must discriminate

Predictable physical structure, an effective grammar, a semantic code and intentional communication are different hypotheses. A periodic or low-dimensional dynamical process can be highly compressible without a sender. A grammar can be a useful description without being a natural language. A code is supported by a stable relation between symbols and independently accessible outcomes; intentional communication additionally requires evidence about an encoder, intervention or communicative context. Model expressivity does not identify those hypotheses from indistinguishable observations.

A realistic staged test begins with held-out prediction against strong physical and statistical baselines. It then asks whether a structured grammar improves compression after charging its parameters and side information, whether its rules transfer to new systems or conditions, and whether interventions discriminate the proposed semantics from alternatives. Known human communication supplies a positive control; physical noise, periodic dynamics and measurement artifacts supply distinct controls. A failure on a positive control limits the decoder; success on a physical series does not establish intent.

39.2 Coding definitions with an honest denominator

For a lossless declared byte representation x and a normalized sequential distribution, ideal bits per byte is $-\log_2 p(x)/|x|_{\text{bytes}}$. A finite precision implementation may additionally need probability quantization, coder termination and a decoder convention. The study reports an ideal likelihood length rather than claiming to have emitted that arithmetic-coded bitstream.

For a directed graph G with attributes X, a useful two-part accounting is

$$L(G, X) = L(\theta) + L(\text{representation}) + L(G \mid \theta) + L(X \mid G, \theta).$$

Divide by a fixed canonical original byte count, or separately report bits per vertex, edge and attribute. Adding graph tokens to the denominator without charging their information can artificially improve a score. Direction, order, types, weights and missingness belong in the graph code unless recoverable from public conditions. A graph inferred using hidden target values is not free side information.

The empirical task uses fixed-point integers at each source's published precision. Eight-bin quantization is fitted on training data only. The symbol likelihood is supplemented by an exact

signed integer residual: with $d = v - r_b$, encode $z = 2d$ for $d \geq 0$ and $z = -2d - 1$ otherwise using Elias gamma on $z + 1$. Its length is $2\lfloor \log_2(z + 1) \rfloor + 1$. This permits exact reconstruction of the represented integer given the bin representative; it does not preserve the source CSV formatting or every original instrument measurement. First differences are reconstructed from explicitly supplied lag-context seeds. The two-part result charges model values, metadata, seeds and any retrieval bank. The interpreter and universal code convention are shared fixed conditions, not asserted to have zero absolute description length.

A computable compressor supplies a description-length upper bound on Kolmogorov complexity up to the decoder/program constant. It does not compute Kolmogorov complexity, whose general exact computation is unavailable to an effective finite procedure. Compressor comparisons, prequential coding, minimum-description-length accounting and normalized compression distances are operational proxies with biases. Representation choice, sample size and compressor overhead can dominate them. Compression alone is not a test for an intentional message.

39.3 Executed bounded experiments

Four data sources were actually downloaded: NOAA monthly Mauna Loa background CO_2, SILSO monthly sunspot number, NOAA Niño-3.4 anomaly, and PhysioNet MIT-BIH record 100 MLII. The first three are observational summaries; the sea-surface index includes a reconstruction model. The ECG experiment uses the first 24,000 samples from one recording. No clinical conclusion or patient-held-out generalization follows. Source URLs, exact byte hashes, units, version qualifications and failures are retained with the data release.

The native comparison uses the actual pinned Mora-3 TokenGT and tropical modules, instantiated as small one-layer, 16-wide, two-head encoders with a task-specific categorical readout. It is not the full approximately 80M model, a pretrained checkpoint or the ongoing 24+2 campaign. Two seeds and 80 optimizer steps per architecture/dataset give 16 bounded component runs. These are actual training executions, not synthetic demonstrations of an untrained operator.

Inputs contain eight strictly past first differences represented by eight node tokens and seven forward edge tokens. Ordered endpoints and node positions retain direction. Targets, future values and target-derived edges are absent. A chronological 60/20/20 partition with discarded boundary contexts separates training, validation and test. Normalization and quantizer construction use training only. Checkpoint selection, baseline order/smoothing and retrieval-mixture selection use validation only. Two seeds do not establish reliable population-level architecture superiority.

The analogy adapter retrieves training-window targets using normalized embeddings from the trained native model. Its bank is temporally disjoint from validation and test. A fixed directed-chain semantics supplies the relation check; no general semantic equivalence between physical systems is inferred. Its memory cost is included in the two-part metric. Exact source-to-target path witnesses and ground-rule deduction are separately implemented and tested.

Exploratory block bootstrap intervals assess variation in paired per-example codelength differences; their stationarity assumptions are not established for all four series. Circularly shifting labels relative to contexts is reported only as a descriptive diagnostic, not a valid universal permutation p-value. No multiple-comparison-adjusted discovery of language or a physical message is asserted. The Gutenberg English control uses actual human-authored text with its provenance establishing communication independently of the score.

39.4 Additional realistic tests, not executed results

CMB/CCC template grammar. Use native Planck temperature/polarization products, masks, beam transfer and noise/foreground uncertainty. Freeze a source grammar and decoder on independent simulations or include their entire fitting/selection procedure inside null calibration. Hold out sky regions only with a dependence-aware design; neighboring pixels are not independent examples. Compare likelihood and code length with standard cosmological and foreground alternatives. A retained anomaly is evidence against a specified null, not automatically CCC or an intentional cross-aeon message. Source: https://pla.esac.esa.int/.

Gravitational-wave graphs. Use GWOSC strain and data-quality flags with separately defined off-source controls. A directed multichannel graph should encode actual sample timing and detector metadata available to the decoder. Compare a learned residual model with established waveform/noise models and charge template selection. Test transfer across held-out events and observing conditions, retaining instrumental transients as alternatives. Neither structured residuals nor interdetector dependence demonstrate language. Source: https://gwosc.org/.

Physiology. Extend from the one-record pilot to patient-group-disjoint public physiological records, then test across devices and conditions. A grammar can encode waveform phases or state transitions; counterfactual clinical claims require interventions or justified causal models beyond a recording. De-identification, source licenses and native sampling/units remain mandatory. Source: https://physionet.org/.

Fluid and transport processes. For measured flow, preserve sensor geometry and sampling uncertainty. For JHTDB or PDEBench simulations, label solver-generated data explicitly, hold out initial-condition and parameter families, and charge numerical error. Test whether a directed flux grammar improves held-out rollout coding while preserving conservation under independent replay. Simulation data are not experimental observations. Sources: https://turbulence.pha.jhu.edu/ and https://github.com/pdebench/PDEBench.

Known communication. Use speech/transcript or documented animal-vocalization/context data with speaker/individual/session-held-out splits and verified redistribution rights. Separate signal decoding from compositional generalization. Match acoustic nuisance factors and evaluate controlled semantic interventions where ethically available. A statistically ordered call sequence is not by itself evidence of human-like syntax. The current English text control is a simpler executed reference, not a substitute for those acoustic experiments.

Information transfer rather than resemblance. Across any pair of processes, predeclare a unit-consistent relational map, source-only memory bank and admissible target subpaths. Compare transferred proposals with target-only and unconstrained retrieval baselines, test wrong-direction and branch-deletion controls, and require an independent semantic or dynamical residual bound before calling the transfer a deduction. This directly connects analogy training to the graph grammar and BGG relation-checking program without conflating them.

39.5 Recorded results from the executed pilot

The executed protocol produced sixteen trained component runs: four data sources, two native attention variants and two seeds. Each run attempted 80 optimizer steps, selecting among checkpoints at steps 20, 40, 60 and 80 using validation only. The ordinary encoder/readout has 2,744 parameters; the tropical-code variant has 2,104. These are deliberately small component tests, not a matched-parameter proof of architectural superiority or a full Mora-3 training result.

Data	Test n	Count	Vanilla	Tropical	V+retrieval	T+retrieval
CO_2	146	2.6657	2.2118	2.1476	2.0616	2.0539
Sunspots	406	2.8996	2.9120	2.9804	2.8768	2.9500
Niño-3.4	171	2.9965	3.0051	3.0009	3.0090	2.9902
ECG MLII	4783	2.8046	2.8632	2.9473	2.7575	2.8696

Entries are held-out *bits per quantized symbol*; the neural columns average the two recorded seeds. They are not bits per original byte and do not include model or retrieval storage. The count baseline's order and smoothing were chosen on validation. Retrieval weights were also chosen on validation, including the option of no retrieval. The complete per-seed results and probability arrays are retained in the experiment outputs, subject to the source-specific redistribution policy.

The result is mixed. Both neural variants improve the selected count predictor on CO_2. They do not consistently beat it on the other sources. Retrieval improves most runs but worsens vanilla seed 29 on Niño-3.4 by approximately 0.0124 bits per symbol. Two seeds and these short evaluation spans do not justify a general ranking of the architectures. The examples demonstrate an executable protocol, not a discovery of language in the measured processes.

Data	Count ideal BPB	Vanilla ideal BPB range	Tropical ideal BPB range
CO_2	2.9370	2.8134–2.8336	2.7852–2.8297
Sunspots	3.4946	3.4925–3.5029	3.5120–3.5176
Niño-3.4	1.9904	1.9925–1.9925	1.9892–1.9937
ECG MLII	1.1652	1.1792–1.1806	1.1992–1.2027

Here BPB uses the canonical four-byte fixed-point target value as denominator and adds the exact integer residual code. It remains an ideal conditional codelength, not an emitted arithmetic-coded file. Adding model parameters, metadata, context seeds and retrieval-bank values changes the comparison substantially. Without retrieval, the neural two-part BPB ranges are approximately 128.8–163.9 for CO_2, 48.8–61.4 for sunspots, 109.4–139.4 for Niño-3.4 and 5.05–6.10 for ECG. A short test stream cannot amortize a model or memory bank merely because its conditional likelihood improves. These costs are reported rather than omitted.

The known human-language control uses the actual Gutenberg English text of *Alice's Adventures in Wonderland*, with boilerplate excluded and disjoint chronological byte spans. Validation selected an order-two count model with smoothing 0.1. Its held-out conditional likelihood is 2.9601 bits per UTF-8 byte; independently applied zlib and lzma produced approximately 3.0386 and 2.9087 bits per byte on that held-out byte segment. The code lengths have different model/coder conditioning and are not a matched total-description-length contest. Authorship/provenance, not compressibility, establishes that this source is human communication. No acoustic or cross-author claim is inferred.

The publication manifest records source and result hashes, failed preflights and the exact reproduction commands. No failed preflight is counted as a trained run, and no unavailable source is counted as an executed experiment.

Chapter 40

Measured signals, communicative controls and complete code lengths

40.1 Information available versus computation made accessible

The source program's conditional-information identity deserves an exact interpretation. In ToricGT-Program Part II, the certificate Z_t is measurable with respect to the full strict-prefix sigma-field $\mathcal{F}_t$. Consequently

$$H(B_t \mid \mathcal{F}_t, Z_t) = H(B_t \mid \mathcal{F}_t), \qquad I(B_t; Z_t \mid \mathcal{F}_t) = 0.$$

The displayed entropy identity in the source is valid; a positive Bayes-optimal information gain under those same measurability premises is zero. This is not a limitation on useful neural computation. It identifies where the gain actually occurs.

For a restricted context $C_t = \psi(\mathcal{F}_t)$ or a restricted predictor class, a computed certificate can expose a useful function of the available prefix. Two independent fair bits (X_1, X_2) provide an exact example: let the target and certificate both be X_1 xor X_2. The full prefix already determines the target, while $C = X_1$ has one bit of target uncertainty. Adding the certificate to C removes that bit. The finite entropy calculation and equality of deterministic-augmentation fibers are executable; the latter is checked in Lean.

The corrected performance question is therefore computational: can an architecture compute and exploit a sufficient representation at lower time, memory, sample or description-length cost? A symbolic certificate, learned embedding, retrieved trajectory or tropical support may help a restricted predictor even though it adds no information to the complete input from which it was computed. Any external memory not recoverable from that input is separate side information and must be charged or conditioned on explicitly.

40.2 Sequential evidence and the cost of searching a grammar

Let $p_t(\cdot \mid x_{<t})$ be a specified null predictive law and let q_t be a normalized, prefix-measurable alternative. Assume the alternative is absolutely continuous with respect to the null on each conditional support. Then

$$M_T = \prod_{t=1}^{T} \frac{q_t(x_t \mid x_{<t})}{p_t(x_t \mid x_{<t})}$$

has conditional null expectation equal to its previous value. This follows by summing $p_t(x)q_t(x)/p_t(x) = q_t(x)$. For any finite horizon, the probability of a first crossing of $1/\alpha$ is at most

α.

Proof. The first-crossing prefixes form a prefix-free set. At each such prefix u, $P(u) \leq \alpha Q(u)$. The total Q-mass of a prefix-free family is at most one, so summing yields $P(\text{crossing}) \leq \alpha$. This also gives the infinite-horizon statement by monotone convergence over finite horizons. The finite one-step and budget inequalities are formalized; a finite binary tree test separately enumerates exact first-crossing masses. □

A prior-weighted mixture of complete candidate laws pays for selection through its weights. The code-length advantage $\log_2 M_T$ is interpretable relative to the specified null; an unvalidated physical noise model does not inherit its calibration merely because the algebra is correct. In cosmology, a grammar adapted after inspecting the same sky must be incorporated into the null experiment or separated by an independent training/evaluation design. This is a direct route from model capacity and compression to a controlled test, rather than a search for whichever pattern looks most like a message.

40.3 Two additional measured controls

The new study acquires actual GWOSC GW150914 strain from Hanford and Livingston and eighteen LibriSpeech utterances from three distinct speakers. All source files have recorded byte hashes. These extend the original four-process study and text control; they do not replace or selectively edit its outcomes.

For speech, six utterances from speaker 6930 form training, six from 1320 form validation and six from 5639 form test. Audio remains mono 16-bit PCM at 16 kHz. These are a small, explicitly re-split subset of the public test-clean archive, not an official LibriSpeech ASR benchmark. Speaker separation tests acoustic transfer; it does not by itself establish semantic comprehension or generalization across a population of speakers. Matching utterance transcripts are retained as independent provenance of known communication and are never supplied as predictor inputs.

For gravitational-wave strain, training uses H1 seconds 0–8 and validation uses H1 seconds 8–12 relative to GPS 1126259446; test uses L1 seconds 24–28. The cataloged event near 16.44 seconds is excluded from fitting. A declared 10^{-24} strain-unit fixed-point representation is used for integer coding, with its quantization error recorded. Exact decoding of that representation does not claim preservation of every bit of the original floating-point or ASCII file.

Both studies fit the actual pinned Mora-3 TokenGT and tropical modules at width 16 with one layer, using seeds 17 and 29 and 80 optimizer updates. A training-only quantizer and normalization map build directed eight-lag input graphs. Each split supplies 2,400 sampled target windows; no target or future-dependent edge enters an input. Validation selects checkpoint, count baseline and retrieval mixture. The full runs, source hashes and per-target arrays are retained.

Held-out source	Count	Ordinary	Tropical	Ordinary+memory	Tropical+memory
Speech speaker 5639	2.2405	2.6111	2.7917	2.4217	2.4720
L1 strain interval	3.9712	2.6296	2.6754	2.6106	2.7414

Entries are bits per quantized target symbol, averaging both recorded neural seeds. The count baseline wins the speech comparison. Both native variants beat that selected count baseline on this detector-transfer interval, with substantial seed variability. Retrieval helps some runs and harms others. These are direct measured outcomes, not a universal ranking of the architectures. Model,

metadata, context and retrieval costs are separately reported in the sampled-window two-part measures.

40.4 An emitted lossless code, not a likelihood-only surrogate

The new codec emits a real packet. It contains a frozen order-two count model, deterministic positive 16-bit arithmetic-coder frequencies, quantizer representatives, segment lengths and initial samples, exact signed residuals, termination and a decoded-content checksum. A separate decoder reconstructs every declared integer sample. Gamma-coded residuals restore information discarded by binning; no uncharged probability table is supplied to the decoder.

Representation	Original bytes	Emitted bytes	Actual BPB	Exact replay
Speech PCM16	2,692,320	2,161,121	6.4216	yes
Strain fixed-point int32	65,536	62,367	7.6132	yes

The packets include their model/header costs: 1,643 bytes for speech and 1,221 for strain. The public directed successor grammar is a shared decoder convention; segment boundaries and seeds are in the packet. Independent zlib and lzma baselines respectively obtain 5.8805 and 4.8748 BPB on the speech representation, and 6.2821 and 4.7246 on strain. The new reference codec therefore does not outperform those general-purpose baselines. Its contribution here is a complete, inspectable realization of the coding claim: actual bytes, paid side information and exact reconstruction. The whole-stream codec and sampled-window native likelihoods have different exposure and representation contracts and are not merged into one leaderboard.

40.5 A physical control that rejects an attractive inference

A separate retrospective waveform diagnostic uses a fixed 35–350 Hz bandpass, a quarter-second window and a lag scan of approximately ±10 milliseconds. The known-event window reaches absolute interdetector correlation 0.3953. Across 52 off-source windows, the maximum is 0.7177. Thus this simple statistic does not isolate the known event from the chosen off-source controls. It supplies no discovery significance and no evidence of language. This failure is retained because a useful discovery pipeline must distinguish physical signal structure, instrumental structure and search effects before making a semantic claim.

The forward/backward filter is explicitly retrospective and acausal; it is excluded from every predictive feature. A more discriminating event analysis must use its own validated waveform/noise models and independent evaluation. No choice was retuned on these test results and then relabeled as a preregistered success.

40.6 Reproduction and source rights

GWOSC and LibriSpeech publish these sources under CC BY 4.0. The accompanying release records source landing pages, file/member hashes, retained native axes and sample rates, and the fact that only a bounded archive prefix was streamed. Transcript coverage is checked by utterance identifier rather than by matching a speaker name alone. Source objects, model checkpoints, emitted packets, predictions and scripts remain distinguishable artifacts.

The results answer the question posed by this program: what structure is actually predictive, transferable and compressible under a declared representation? The experiments do not substitute

that operational question for the separate question of intentional communication. Known speech supplies the latter's independently documented positive control. Physical observations supply structure whose interpretation must be earned by further discriminating tests.

C1 LIGO Scientific Collaboration and Virgo Collaboration. *Data release for event GW150914.* DOI 10.7935/K5MW2F23. https://gwosc.org/events/GW150914/

C2 Panayotov, Chen, Povey and Khudanpur. *LibriSpeech: an ASR corpus based on public domain audio books.* ICASSP 2015; SLR12 data and license. https://www.openslr.org/12

Chapter 41

BGG, category $\mathcal{O}$ and structural synthesis

41.1 Three BGG constructions, with their domains kept separate

The BGG layer supplies a systematic language for generators, relations, standard objects, resolutions and controlled changes of presentation. It is useful to the present program precisely when those data are retained. It is not a universal identification of every quiver, neural network, quantum model or physical process with category $\mathcal{O}$.

Three constructions must first be distinguished. Classical BGG category $\mathcal{O}$ is a category of modules over a complex semisimple Lie algebra's universal enveloping algebra. A BGG resolution resolves an appropriate finite-dimensional highest-weight module by Verma modules. The exterior–symmetric BGG correspondence relates graded exterior-algebra modules, linear complexes over a symmetric algebra and, in the appropriate quotient/derived formulation, coherent sheaves on projective space. The names and some homological techniques overlap, but these are not the same category or the same theorem [B1–B4]. The project's existing finite exterior/BGG certificates belong to the latter algebraic interface unless another construction explicitly connects them to highest-weight theory.

Fix a complex semisimple Lie algebra with triangular decomposition

$$\mathfrak{g} = \mathfrak{n}^- \oplus \mathfrak{h} \oplus \mathfrak{n}^+, \qquad \mathfrak{b} = \mathfrak{h} \oplus \mathfrak{n}^+.$$

An object of classical $\mathcal{O}$ is a finitely generated $U(\mathfrak{g})$-module that is a direct sum of $\mathfrak{h}$-weight spaces and is locally finite for $U(\mathfrak{n}^+)$. In this setting one can equivalently phrase the local finiteness condition using the Borel action. These are structural hypotheses, not labels attached to an arbitrary embedding. For a weight λ, the Verma module is

$$M(\lambda) = U(\mathfrak{g}) \otimes_{U(\mathfrak{b})} \mathbb{C}_\lambda,$$

where $\mathfrak{n}^+$ acts trivially on the highest-weight line. It has a simple highest-weight quotient $L(\lambda)$. The theory of projective covers, blocks, standard filtrations and reciprocity requires this ambient category and its stated field assumptions [B1,B2,B4].

The practical question for ToricGT/Mora-3 is therefore constructive: which graph tokens encode a weight, generator, relation, module map or resolution component; which operator realizes the relevant action; and which checker verifies the required identities? A directed attention graph alone answers none of those questions. The finite examples below provide a reproducible starting point rather than asserting a completed categorical implementation.

41.2 An exact highest-weight example and its truncation failure

Take $\mathfrak{sl}_2$ with generators e, f, h and relations

$$[e, f] = h, \qquad [h, e] = 2e, \qquad [h, f] = -2f.$$

On the Verma basis $v_i = f^i v_0$, $i \geq 0$, the actions are

$$fv_i = v_{i+1}, \qquad hv_i = (\lambda - 2i)v_i, \qquad ev_i = i(\lambda - i + 1)v_{i-1}, \quad ev_0 = 0.$$

For example, the coefficient of v_i in $(ef - fe)v_i$ is

$$(i + 1)(\lambda - i) - i(\lambda - i + 1) = \lambda - 2i.$$

The other two commutators follow by comparing the two weights adjacent to a raising or lowering step. Thus this is an exact infinite-dimensional representation, not merely a finite matrix pattern that happens to look triangular.

Proposition 41.1 (The $\mathfrak{sl}_2$ BGG short exact sequence)**.** *For an integer $n \geq 0$, the vector $v_{n+1} \in M(n)$ is singular and generates a submodule isomorphic to $M(-n - 2)$. The quotient is the $(n + 1)$-dimensional simple module $L(n)$, giving*

$$0 \longrightarrow M(-n - 2) \longrightarrow M(n) \longrightarrow L(n) \longrightarrow 0.$$

Proof. The raising coefficient at $i = n + 1$ is zero, and its weight is $n - 2(n + 1) = -n - 2$. Map the basis element w_j of $M(-n - 2)$ to v_{n+1+j}. This map commutes with f and h. For e, the two coefficients are both $-j(n + 1 + j)$, so it is a module map. Its images are distinct basis vectors, hence it is injective. The quotient has basis $v_0, \ldots, v_n$. Its lowering map stops at v_n and its nonzero raising coefficients connect every lower weight to the highest line. Any nonzero submodule contains a weight vector, and repeated raising reaches that line; repeated lowering then generates the entire quotient. It is therefore simple. □

The corresponding formal character subtraction removes the infinite tail:

$$\frac{e^n}{1 - e^{-2}} - \frac{e^{-n-2}}{1 - e^{-2}} = e^n + e^{n-2} + \cdots + e^{-n}.$$

This is a controlled cancellation between characters of specified modules. It is not the same operation as replacing a sum of arbitrary positive path weights by a tropical maximum.

The quadratic central element $C = ef + fe + h^2/2$ acts with scalar $\lambda(\lambda + 2)/2$ on this Verma representation. Substituting $-\lambda - 2$ gives the same scalar. Thus a central scalar or character summary can identify a block-level relation without distinguishing the two modules or their extension structure. This is another reason not to infer module equivalence from a single spectral or compressed invariant.

Proposition 41.2 (An arbitrary cutoff is usually not a representation)**.** *Keep only $v_0, \ldots, v_{K-1}$ and set $fv_{K-1} = 0$, retaining the other displayed coefficients. Then $[e, f] - h$ vanishes on the first $K - 1$ basis vectors, but on the last one equals*

$$([e, f] - h)v_{K-1} = K(K - 1 - \lambda)v_{K-1}.$$

The cutoff satisfies this relation exactly only when $\lambda = K - 1$.

Proof. At the top vector, ef is zero, fe has coefficient $(K-1)(\lambda-K+2)$, and h has coefficient $\lambda-2(K-1)$. Subtraction and expansion give $K(K-1-\lambda)$. Away from the boundary, the original Verma computation remains unchanged. □

This defect is an explicit finite-size artifact. A large cutoff can satisfy many interior tests and still fail the algebra at its boundary. The finite quotient $L(n)$ with $K=n+1$ is different: the discarded tail is an actual submodule. Numerical truncation and categorical quotient are not synonyms. The same methodological lesson applies to the fluid/gauge laboratories and to bounded latent workspaces: inspect the boundary relations rather than infer an infinite theorem from interior agreement.

The implemented exact checker verifies $L(n)$ for $0 \le n \le 8$, several rational-weight cutoff cases, Casimir coefficients and singular-vector embeddings. Six separate Lean theorems check the coefficient identities used here. They do not constitute a Lean formalization of category $\mathcal{O}$ or exactness of the infinite resolution. The distinction is retained in the proof receipts.

41.3 Blocks, Ext data and quivers with relations

A block of category $\mathcal{O}$ can be studied through a projective generator P and its endomorphism algebra. With left-module conventions, one uses $A = \mathrm{End}(P)^{\mathrm{op}}$; equivalently, the original endomorphism algebra acts on the right of $\mathrm{Hom}(P,-)$. The opposite-ring convention matters for directed composition. In the relevant finite block setting, this gives a finite-dimensional algebraic model of the category, despite the individual Verma modules being infinite-dimensional [B2].

The quiver of a basic algebra records simple objects and first extensions, while relations and higher homological information determine how those extensions compose. A Bruhat graph, a neural architecture graph and an Ext-quiver are not automatically identical. A representation learner that stores only an adjacency matrix can discard precisely the relation and extension data needed to distinguish modules. Suitable targets include standard-filtration multiplicities, morphism dimensions, relation satisfaction and selected higher Ext groups, with explicit finite algorithms or external oracles for the examples used.

BGG reciprocity relates the multiplicity of a Verma standard in a projective standard filtration to an appropriate composition multiplicity in a Verma module. In the classical convention,

$$(P(\lambda) : M(\mu)) = [M(\mu) : L(\lambda)].$$

This is an attributed theorem of the category, not an identity obeyed by arbitrary nonnegative matrices predicted by a network [B1,B4]. It suggests a useful consistency test: predict the two sides through independently constructed views of a finite example and reject disagreement. But enforcing equality alone does not identify the correct multiplicities; an all-zero prediction would satisfy many such equalities without representing the intended objects.

Beilinson–Ginzburg–Soergel establish Koszul structures and duality patterns in the stated category-$\mathcal{O}$ settings. A Koszul grading means that the semisimple degree-zero part has a linear graded projective resolution; it is not merely a convenient integer feature on a graph. Their principal-block and parabolic–singular statements retain their hypotheses and indexing conventions [B2]. The ordinary BGG resolution of a dominant-integral finite-dimensional $L(\lambda)$ has degree-i terms built from $M(w\cdot\lambda)$ with $\ell(w)=i$. These are Verma modules, not a claim that every term is projective. The dot action is $w\cdot\lambda = w(\lambda+\rho)-\rho$.

For a neural grammar, the standards and projectives provide different types of construction object. A proposed rewrite between complexes must preserve the actual differential and its grading, not only the alternating sum of dimensions. Equality of a Grothendieck-group class is weaker than

isomorphism; a chain map is weaker than a quasi-isomorphism; a derived equivalence is stronger still. This hierarchy can become part of the model's type system and certificate vocabulary.

41.4 Exterior BGG certificates, exactness and integral obstructions

Let $E_1, \dots, E_m$ act on a finite graded space and satisfy

$$E_i^2 = 0, \qquad E_iE_j + E_jE_i = 0 \quad (i \neq j).$$

For commuting scalar coordinates x_i, define $d(x) = \sum_i x_i E_i$. Expansion gives $d(x)^2 = 0$. This is the elementary linear-complex mechanism underlying the exterior/symmetric interface. The full BGG correspondence adds grading, functors and categorical identifications; the identity alone is not the whole correspondence [B3].

Proposition 41.3 (A quantified relation defect)**.** *For bounded operators and any submultiplicative operator norm,*

$$\|d(x)^2\| \leq \sum_i |x_i|^2 \|E_i^2\| + \sum_{i<j} |x_i x_j| \|E_iE_j + E_jE_i\|.$$

Proof. Expand the square, group the two ordered cross terms for each unordered pair, and apply the triangle inequality and scalar homogeneity. No optimization or statistical assumption enters this estimate. □

The bound gives a meaningful train-time diagnostic when the matrices are learned: the declared relation residuals control a specified differential residual on a bounded coordinate domain. A numerical zero at a few sampled x values does not prove the relations everywhere. A small residual also does not imply proximity to a genuine representation without an additional stability theorem. The norm and metric must be fixed: arbitrary ill-conditioned basis changes can distort residual sizes, whereas unitary changes preserve the usual operator and Frobenius norms.

The executable example uses left exterior multiplication on a bitmask-indexed wedge basis. For any nonzero rational vector x, choose a covector a with $a(x) = 1$. Contraction ι_a satisfies

$$\iota_a \, d(x) + d(x) \, \iota_a = I.$$

The identity follows on each wedge monomial from the contraction derivation rule. It supplies an actual contracting homotopy, hence exactness of that finite exterior complex over the field. At $x = 0$, the differential vanishes and the cohomology is the whole graded space. The checker retains both cases instead of treating $d^2 = 0$ as a synonym for exactness.

Integral coefficients change the conclusion. The one-generator complex

$$0 \longrightarrow \mathbb{Z} \xrightarrow{2} \mathbb{Z} \longrightarrow 0$$

has a $\mathbb{Z}/2\mathbb{Z}$ cokernel, although it becomes contractible over $\mathbb{Q}$ using the coefficient $1/2$. Reduction modulo two kills the differential. This gives a concrete connection to the orders program: scalar extension, reduction and integral exactness preserve different information. A rational or finite-field rank check cannot replace a certificate of an integral module claim.

41.5 How to use the material in the neural research program

A defensible architecture extension stores a typed finite presentation: the algebra and field, generator matrices or actions, grading/weight labels, relations, chosen module truncation or quotient, and a metric. A graph-to-graph network can predict missing maps, propose a resolution step, retrieve a structurally analogous complex, or choose a rewrite. An independent exact or interval checker then verifies the property appropriate to that proposal. The graph-UAP discussion concerns representability of bounded encodings; category $\mathcal{O}$ supplies additional structure on the encoded objects, not a shortcut around the UAP hypotheses.

For ToricGT or TropicalGT, tropical scores can rank candidate generators, Bruhat moves or resolution components. Their scores need not themselves be module morphisms. For Mora-3, a typed graph/forest head can represent alternative derivation histories while the analogy module retrieves source complexes using condition-only keys. The relation/path checker prevents an attractive embedding match from reversing a morphism or dropping a required branch. A stronger transfer requires chain maps and commutation witnesses, not only a graph match.

A chain-map penalty has the concrete form $\|d_Y T - T d_X\|$, with sources and targets in the correct degrees. A homotopy witness for two maps requires $T - S = d_Y H + H d_X$ with the appropriate graded components. These are proposed training/verification objectives unless a particular run and certificate are supplied. They are not evidence that the current 16 real-data component runs learned category-$\mathcal{O}$ structure: those runs used likelihood and retrieval, not these algebraic losses.

Avoid vacuous algebraic success. Setting all generator matrices to zero satisfies many homogeneous relations but represents only a trivial or collapsed object. A useful objective must additionally bind the intended dimensions, weights, ranks, central character or task readout, and measure whether the constraints improve held-out prediction or search cost. Fixed exterior operators make $d^2 = 0$ true by construction; that fact alone says nothing about whether learned data features are informative.

GFlowNet construction states may include a partial complex, a proposed morphism and its pending obligations. Rewards can favor valid completed objects, novelty and compactness while a verifier enforces exact relations. The target measure must still specify whether terminal objects are presentations, histories or equivalence classes. BGG resolutions do not eliminate multiplicity from a sampling law. A learned compressor must charge the presentation, grading, certificate and retrieval memory wherever they are not public side information.

For dynamical models, a sequence of learned state updates may be tested for compatibility with a fixed algebra action or a chain complex. If the action itself changes, transport between the actions is additional data. A small commutator residual at one time does not guarantee stability of the entire switched system. Combine a domain-wide residual estimate with the earlier error-propagation or weighted-gain certificate; retain protected memory when a globally contractive workspace would erase required distinctions.

41.6 Tropical, toric, hypertoric and surface connections

Root systems and Weyl chambers provide exact combinatorial geometry in highest-weight theory. Neural activation fans and cluster tropicalizations also produce polyhedral structures, but an arbitrary rational fan need not be a Weyl fan and an arbitrary mutation need not be a reflection. A connection requires a map of lattices, compatible relations and a statement of what is preserved. Merely using the same word “chamber” supplies no theorem.

Hypertoric category $\mathcal{O}$ is a substantive external bridge rather than a visual analogy. Braden–

Licata–Proudfoot–Webster study invariant subalgebras of a Weyl algebra under torus actions and corresponding categories related to quantized hypertoric geometry. Under their regularity/integrality hypotheses, the categories have highest-weight and Koszul structures, with Gale-dual arrangements controlling a duality [B6]. The corrected version 5, including its correction to Lemma 8.25, is the referenced version. The required quantized polarized arrangement, lattice and parameter data must be constructed before applying that result to a model.

This suggests a precise research path: begin with a small integral arrangement; build its documented algebra/category; encode standards, allowed paths and relations as typed graphs; then evaluate whether a learned policy can predict or navigate those objects efficiently. Compare against exact combinatorial algorithms. A phase vector in ToricGT or a tropical support cell is not already a hypertoric quantization, so the initial stage is a construction and validation task, not an asserted categorical identity.

Gentle surface algebras offer another representation-theoretic family. The earlier surface models retain curves, grading/line fields, stops and relations. Category-$\mathcal{O}$ block algebras and gentle algebras need not share global dimension or the same derived category. The radical-square-zero oriented triangle discussed earlier has infinite global dimension; it must not be silently identified with a finite-global-dimension regular category-$\mathcal{O}$ block. A common language of quivers, complexes and derived functors is useful without identifying the families.

Soergel bimodules and their diagrammatic/compositional theory provide externally justified categorifications linked to Hecke and Coxeter structures; the Hodge-theoretic character results retain their own hypotheses [B5]. They suggest additional certified grammar relations and graded targets, but no Soergel implementation or learned character-table result is claimed in this release. A later extension should specify the coefficient field, realization, grading and diagrammatic relations, and test them independently before using their outputs as supervision.

41.7 Quantum information, orders and the GRH/Hilbert–Pólya boundary

Finite-dimensional $\mathfrak{sl}_2$ modules can be related to compact-real-form spin representations after choosing a compatible positive Hermitian form. In the basis above for $L(n)$, positive squared norms can be chosen recursively so that $\|v_i\|^2 = i(n-i+1)\|v_{i-1}\|^2$. This makes the displayed raising and lowering actions adjoint in that weighted inner product. The construction does not equip every Verma module or every object of $\mathcal{O}$ with a physical positive-definite quantum-state interpretation.

A quantum implementation additionally needs a finite encoding, admissible Hamiltonians/channels, preparation, measurement and an error budget. Nonsemisimple extensions are algebraic data, not automatically density matrices or implementable quantum channels. A categorical zero relation and a zero composition of unconditional trace-preserving channels are different notions, as established in the earlier realization chapter. A nonunitary change of module basis does not establish unitary equivalence of a physical evolution.

Integral orders require a compatible lattice and control of denominators, torsion and bad primes. The elementary 2-differential example demonstrates the failure of transferring rational exactness unchanged to an integral setting. Integral forms of enveloping algebras are not, merely by being integral forms, finite-rank orders in a fixed finite-dimensional algebra under the monograph's definition. One must identify the actual finite algebra and full lattice needed for an orders theorem.

For a Hilbert–Pólya proposal, algebraic central characters, Casimir spectra or Koszul duality do not identify the zeros of a specified arithmetic L-function. One still needs a Hilbert space, a densely defined self-adjoint operator with its domain, and the required arithmetic spectral or

determinant identity. Finite representations and exact categorical identities can organize candidate constructions, symmetry checks and trace calculations. They do not supply the missing analytic identity, a continuum convergence theorem or a proof of GRH by themselves.

Cryptographic claims remain separate as well. A complicated representation category, a large Weyl group, mutation dynamics or a learned search policy is not a security reduction. The algebraic machinery may help construct or analyze instances, detect leaked invariants, or certify a finite computation. Security requires an explicit threat model and problem-specific evidence. No cryptographic primitive is replaced by a novel unreviewed category-$\mathcal{O}$ construction in this release.

41.8 An executable agenda and the limits of the current evidence

The current algebraic implementation covers exact small $\mathfrak{sl}_2$ quotients, singular-vector coefficients, Casimir checks, cutoff defects, exterior anticommutation, $d^2 = 0$ and an explicit rational contracting homotopy. Written proofs establish the elementary general statements. The six Lean declarations formalize selected polynomial identities, not the full external category theory. These examples are constructed mathematical tests and remain separate from the genuinely measured-data experiments in the preceding chapter.

The next scientifically meaningful comparison would add a specified relation-preserving head to a frozen native architecture and compare it with matched unconstrained and trivial-representation controls on held-out algebraic structures and real tasks. It would evaluate whether the constraints improve code length, transfer, proof-search cost or robustness after charging their compute and storage. It must retain wrong-direction, boundary-truncation, bad-prime and collapsed-feature controls. No such trained BGG improvement is claimed from the present pilot.

For discovering structure in physical data, BGG methods can supply candidate algebraic descriptions, invariants and consistency checks. A learned description is persuasive when it predicts new observations, survives perturbation and transfers through verified maps. Even then, algebraic regularity or compression is not evidence of a hidden intentional message without an independently testable encoding/communication hypothesis. The purpose of this synthesis is to make those tests more exact, not to turn representation theory into an inference about intent.

Chapter 42

Reading the linked surface experiment

The explorer makes one computation inspectable across four linked views. Its verified 3×3 toric cellulation has nine black rotation cycles, eighteen white involution cycles and nine faces on thirty-six darts. The eighteen white midpoint vertices index the physical edge qubits. Euler characteristic zero, binary check ranks eight and eight, and two logical qubits are recomputed from the manifest.

Select an edge in the torus, the flat fundamental-domain chart or the token table. Its identifier is highlighted in the geometry, circuit and attention readout. The flat view repeats periodic boundary representatives without adding physical qubits. Torus curves follow the declared periodic coordinates rather than straight chords through the surface.

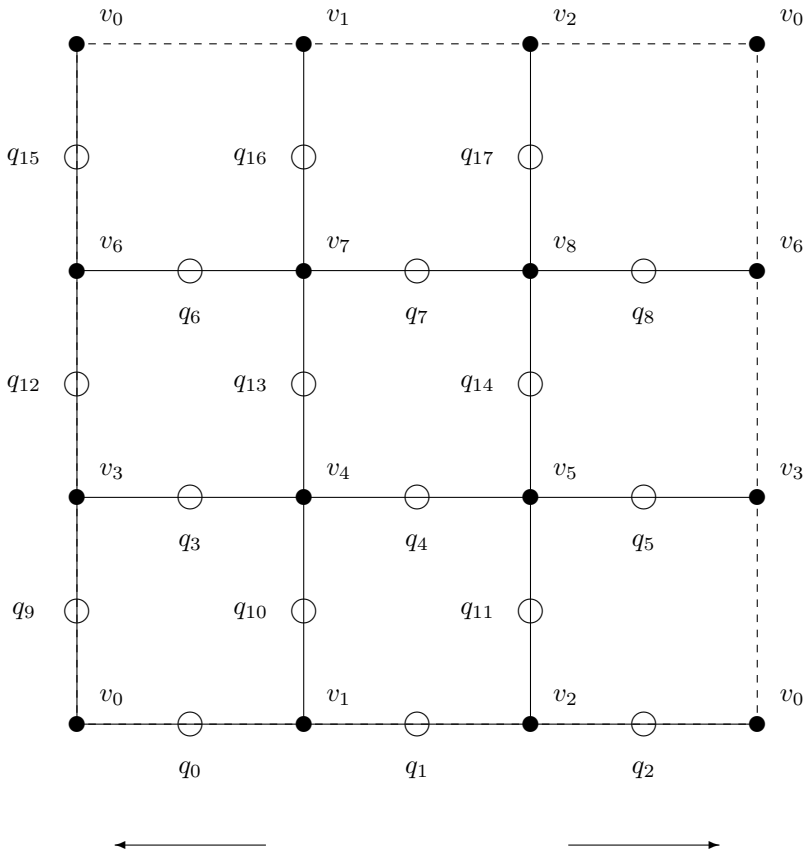
Each edge supplies a two-dimensional key and four-dimensional value derived from those coordinates. Query, gain, temperature and input signs are interactive. Softmax averages value vectors. Score-hard routing averages values at tied score maximizers. Normalized max-plus takes a separate score-plus-value maximum in each coordinate. Displayed winners reveal why the last two operations can select different controls.

The four classical outputs set four R_Y angles on the first star. Starting from the prepared eighteen-qubit state, the worker applies those rotations and three CNOTs, then calculates the target Z expectation, sixteen-outcome patch marginal, stabilizer expectations and toric-code energy. The reported marginal entropy is a measurement entropy, not an entanglement entropy.

An independent identity predicts the expectation as input-sign parity times the product of the four angle cosines. The circuit is nevertheless executed: the identity checks execution rather than replacing it. NumPy, the browser Float64 kernel and NVIDIA cuStateVec complex128 were compared on nine controller configurations. The maximum GPU/CPU amplitude discrepancy was 6.94×10^{-18}; the maximum parameter-shift versus analytic derivative discrepancy was 3.61×10^{-16}.

Rendered checks cover all controller modes, sign response, linked selection, keyboard operation, URL restoration, model-bound JSON export, stale-result rejection and mobile orientations. Input changes immediately invalidate the earlier export. Delayed worker messages cannot overwrite newer configurations. The same browser audit replays all five finite contracts through the local Lean service.

The browser and GPU runs are explicit classical simulations of the quantum model. The four-qubit half-density experiment separately checks a dense unitary evolution operator. Together these realizations make the relationship between combinatorial structure, classical control and quantum-state evolution directly reproducible. Hardware execution, learned physical identification and efficient gate compilation remain distinct operations with their own evidence requirements.



Opposite boundaries are identified.

Figure 42.1: A fundamental-domain chart of the actual 3×3 toric-code manifest. Boundary vertex labels repeat because opposite sides are identified. The eighteen white edge midpoints, not the thirty-six incidence darts, index the physical qubits. The selected first star is $\{q_0, q_2, q_9, q_{15}\}$.

Part IX

Completed computational steps and the unified research program

Chapter 43

Completed next steps: uniform arithmetic and cover-compatible operators

43.1 One continuing manuscript and one project entry point

This edition integrates the original monograph, *Certified Surface Computation*, and the arithmetic–covering–CCC continuation in the same source-built document. The chapter and theorem counters are shared, the table of contents is regenerated, and the earlier source registers remain inspectable appendices. The current project page and its existing `downloads/paper.pdf` address are the canonical entry points. Historical releases are retained for provenance, not presented as separate current installments.

The integrations below are executed next steps, not a list of intentions. Their acceptance criteria and measurements are generated from the emitted circuits, fitted coefficients and calibration samples. Their finite scope is stated separately from broader mathematical goals. Replacing a field primitive is not the same as replacing every primitive in a cryptographic-size elliptic algorithm; a finite channel experiment does not establish a continuum field theory; and more Monte Carlo draws do not repair a misspecified observational null.

43.2 Uniform reversible field primitives instead of lookup synthesis

The previous finite compiler used enumerated tables to make every operation inspectable. The new compiler generates arithmetic uniformly from the bit width and modulus. It implements carry and borrow recurrences, a schoolbook product, constant-divisor binary long division, and exponentiation for inversion. It never enumerates all possible inputs to generate a circuit. The logical alphabet remains X, CX and CCX, making the accounting compatible with the earlier reference.

For bits x, y, c, a carry step writes

$$s = x \operatorname{xor} y \operatorname{xor} c, \qquad c' = (x \wedge y) \operatorname{xor} ((x \operatorname{xor} y) \wedge c).$$

The two carry terms are disjoint, so the XOR equals their Boolean disjunction. The borrow recurrence for $x - y - b$ is

$$b' = (\neg x \wedge y) \operatorname{xor} (\neg(x \operatorname{xor} y) \wedge b).$$

Both identities can be checked on eight bit triples and then iterated inductively. Each Boolean intermediate is computed into a fresh zero wire. The result and a validity bit are copied into an arbitrary output register, after which the complete computation is reversed. This restores every saved intermediate without erasing the retained input.

Proposition 43.1 (Long-division invariant and total reversible oracle)**.** *Let p be a fixed positive n-bit modulus. Before a binary long-division step suppose $0 \le r < p$. On consuming the next bit b, set $s = 2r + b$. One conditional subtraction of p gives the unique remainder r' in $[0, p)$ with $r' \equiv s \pmod{p}$. Consequently the emitted product-and-reduction circuit computes multiplication modulo p. Its compute–copy–uncompute extension is a permutation on every input/output word with clean scratch.*

Proof. The invariant gives $0 \le s < 2p$. If $s < p$ the remainder is s; otherwise it is $s - p$. Both are in $[0, p)$ and congruent to s. Starting with remainder zero and reading the product bits from most significant to least significant proves the reduction invariant by induction. The carry construction computes the ordinary product before reduction. Finally, for retained input x, XOR by the computed output $f(x)$ is an involution on every output word. Reversing the computation restores scratch, proving the claimed reversible extension. □

Invalid field inputs contribute the all-zero output, including a false validity bit. For inversion, zero is invalid; on nonzero inputs and a declared prime modulus the binary exponent $p-2$ computes the inverse. The validity test is itself emitted and charged. Complete input/output-register tests are run on selected small moduli; wider tests include explicit boundary words and independently seeded arbitrary output registers.

An n-bit addition uses $O(n)$ elementary operations, multiplication and long division use $O(n^2)$, and binary-exponent inversion uses $O(n^3)$. Saving all intermediates gives the same polynomial orders for work space. Cleanup doubles the computation rather than changing these orders. This compiler is not claimed to have the single-ancilla efficiency of the distinct ripple-carry construction of Cuccaro and collaborators. That paper is a resource benchmark, not a name applied to a different implementation: https://arxiv.org/abs/quant-ph/0410184.

Prime	Operation	Gates	Scratch	Tested inputs
7	add	276	78	64
7	multiply	708	203	64
7	inverse	1986	569	8
13	add	347	99	256
13	multiply	1129	328	256
13	inverse	5199	1511	16
31	add	466	132	1024
31	multiply	1820	526	1024
31	inverse	11724	3393	32
251	add	727	207	112
251	multiply	4305	1256	112
251	inverse	48577	14190	256
65521	add	1439	411	112
65521	multiply	16177	4756	112

The new primitive closes the exponential-input-table gap for field arithmetic. The original total-domain elliptic circuits remain available as independent references. Integrating these primitives into a complete variable-point elliptic algorithm, proving all exceptional cases there, and selecting a

fault-tolerant architecture are distinct further tasks; the field-oracle result does not silently claim them completed.

43.3 A trained neural operator that respects covers and quantum channels

For the cycle shift S_N, define the lazy walk

$$P_N = \tfrac{1}{2}I + \tfrac{1}{4}(S_N + S_N^*), \qquad L_N = 4(I - P_N).$$

The new time-conditioned network outputs nonnegative coefficients $w_k(t)$ summing to one, with an exact identity channel at $t = 0$. It learns

$$T_{\theta,N}(t) = \sum_{k=0}^{24} w_k(t) P_N^k$$

against the exact graph heat operator $\exp(-tL_N)$. The same coefficients are used at every size. Training uses sizes three and six; validation selects the checkpoint on distinct times; testing uses previously unseen sizes nine, fifteen and twenty-four and a separate time grid. Three independently initialized runs are retained.

Proposition 43.2 (Structure is preserved independently of fitting accuracy)**.** *Every $T_{\theta,N}(t)$ is a nonnegative doubly stochastic matrix. If a covering pullback satisfies $P_N U = U P_M$, then $T_{\theta,N}(t)U = UT_{\theta,M}(t)$. The associated quantum channel*

$$\mathcal{P}_N(\rho) = \tfrac{1}{2}\rho + \tfrac{1}{4}S_N \rho S_N^* + \tfrac{1}{4}S_N^* \rho S_N$$

and every convex mixture $\sum_k w_k \mathcal{P}_N^k$ are completely positive and trace-preserving. Their diagonal action agrees with $T_{\theta,N}$.

Proof. Powers and convex combinations preserve nonnegativity and both unit sums. Induction proves $P_N^k U = U P_M^k$, and linearity proves the polynomial identity. Each shift conjugation is a unitary channel. Convex combinations and compositions preserve complete positivity and trace preservation, proving the quantum statement. Taking diagonals commutes with these permutation conjugations, giving the classical walk and its powers. □

This construction connects the covering and quantum-coarse-graining interfaces with an actual learned operator. Its structural correctness does not depend on a small training loss. Its approximation quality is evaluated independently: for the circulant matrices the operator-norm error is the maximum difference of their Fourier eigenvalues. Browser inference evaluates the exported network and finite walk, not a pre-rendered loss curve.

The exact heat operator has a stronger analytic same-degree baseline. Its infinite walk expansion has Poisson weights with mean $4t$. Truncate after degree K and assign the omitted probability to the identity term. The operator-norm error is at most

$$2\Pr\{\mathrm{Poisson}(4t) > K\},$$

because every walk power has norm at most one. The neural model is compared with this baseline; a learned model is not presumed preferable when the physics already supplies the coefficients.

Seed	Held-out cases	Maximum learned error	Maximum analytic error
17	21	0.033905	5.21e-07
29	21	0.030755	5.21e-07
47	21	0.015812	5.21e-07

The results establish a trained, cover-compatible, probability-preserving finite operator and its CPTP realization. They do not establish an advantage over the analytic kernel. The experiment does not replace the learned Ising channel: it tests size transport and an exact conserved-structure contract that the earlier four-site fit did not test.

43.4 Higher-resolution CCC calibration with explicit simulation uncertainty

The original WMAP products, 72 templates, three methods and learned QNN parameters are retained. An additional exploratory precision study uses the joint phase action, preserving the complex cross-products between the three map products. It refreshes the independent reference normalization with 512 draws, uses 4,095 calibration realizations and tests 1,024 additional null realizations in bounded batches. Its minimum rank is 1/4096, rather than 1/512. This is a new declared precision study, not a backdated replacement for the earlier result.

If K of B independent null simulations exceed the fixed observed score, the Monte Carlo rank is $(1 + K)/(B + 1)$. The associated binomial interval describes uncertainty in the conditional exceedance probability, with the observed statistic, normalization and surrogate model held fixed. It does not measure the probability that CCC is true. Exact binomial intervals are also reported for the independent null-validation fraction.

Precision-study quantity	Value
Calibration draws	4095
Observed exceedances	154
Finite rank	0.037842
Independent null rejections	57/1024
Conditional tail interval (95%)	[0.031990, 0.043895]
Null-validation interval (95%)	[0.042429, 0.071521]

Model error is not simulation error. For an event A and proposed null law Q, a justified bound $\|P - Q\|_{\mathrm{TV}} \leq \delta$ would give $P(A) \leq Q(A) + \delta$. No numerical δ for the real foreground/instrument pipeline is established here. Increasing B shrinks simulation uncertainty under Q but supplies no such missing physical bound. Common-reference sensitivity and frequency-residual results remain alongside the precision study, including disagreements.

43.5 Completion ledger and remaining research obligations

The uniform field generator, total-domain small-register tests, trained cross-size operator, structural proofs, larger joint calibration, and same-page replay are implemented and executed. The integrated build checks continuous manuscript references and rejects missing source inputs. Browser tests exercise the new controls alongside the original six laboratories, including restoration of shared-page state and exported numerical records.

Broader tasks remain explicitly identified: full polynomial-size elliptic scalar arithmetic with a physical resource model; general trained-QNN infinite-width limits beyond the specified independent/dependency regimes; an observational likelihood incorporating actual instrument and foreground uncertainties; and analytic continuum or physical claims requiring those premises. Each now has more explicit computational or mathematical interfaces against which further completion can be tested. None is deferred merely because of the history of the problem, and none is inferred from a favorable finite example.

Chapter 44

A worked bridge: orders, periodic complexes and spectra

The preceding chapters provide several legitimate connections, each with a specified domain of validity. We now place them beside one another in a single calculation. Its purpose is to identify exactly which structures survive a change of viewpoint. The example is elementary enough to calculate without appealing to an existence theorem, but rich enough to distinguish a generic semisimple algebra, an integral order, a singular special fiber, an infinite projective resolution, and a positive finite-dimensional Hodge operator. It also directly uses the nodal order occurring in the inherited research program [137–139].

44.1 The order and its two complementary elements

Let R be a commutative integral domain, let $u \in R$ be nonzero, and set

$$\Lambda = R[X]/(X^2 - uX), \qquad x = \overline{X}, \qquad y = u - x.$$

Then $xy = yx = 0$. The polynomial is monic, so every element of Λ has a unique expression $a + bx$, with $a, b \in R$. Relative to the basis $(1, x)$, multiplication by x and y is represented by

$$M_x = \begin{pmatrix} 0 & 0 \\ 1 & u \end{pmatrix}, \qquad M_y = \begin{pmatrix} u & 0 \\ -1 & 0 \end{pmatrix}.$$

These matrices exhibit the zero-composition condition as an identity over the coefficient ring, before any choice of scalar field or inner product.

Proposition 44.1 (An exact periodic complex)**.** *There is an exact complex of free Λ-modules*

$$\cdots \xrightarrow{x} \Lambda \xrightarrow{y} \Lambda \xrightarrow{x} \Lambda \xrightarrow{y} \Lambda \xrightarrow{x} \cdots .$$

Here an arrow labelled by an element denotes multiplication by that element. Moreover, with $M = \Lambda/x\Lambda$, the augmented sequence

$$\cdots \xrightarrow{x} \Lambda \xrightarrow{y} \Lambda \xrightarrow{x} \Lambda \longrightarrow M \longrightarrow 0$$

is a free resolution of M.

Proof. For $a + bx \in \Lambda$,

$$x(a + bx) = (a + ub)x, \qquad y(a + bx) = a(u - x) = ay.$$

Uniqueness of the basis expansion gives

$$\ker M_x = R(x - u) = Ry, \qquad \operatorname{im} M_y = Ry.$$

It also gives $\ker M_y = Rx$: the coefficient of x in ay is $-a$, so $ay = 0$ implies $a = 0$. Finally $\operatorname{im} M_x = Rx$, since $x = x \cdot 1$. Thus the kernel of each arrow is the image of the preceding one. The quotient map onto M has kernel $x\Lambda = Rx$, which verifies exactness at the augmented term. Every unaugmented term is the free module Λ of rank one. □

The ring need not be a domain: indeed $xy = 0$. This is not a defect in the argument. The coefficient ring R is a domain, whereas the order Λ is a ring with zero divisors. Keeping those two statements separate prevents an erroneous cancellation of x or y. In particular, this is not a regular-sequence Koszul resolution. It is a periodic resolution arising from the explicit annihilator identities.

44.2 Tensor products detect information that ranks omit

The module M is isomorphic to R as an R-module, with x acting as zero. The action of y on M is consequently multiplication by u. Tensoring the resolution with M over Λ produces the chain complex

$$\cdots \xrightarrow{0} R \xrightarrow{u} R \xrightarrow{0} R \xrightarrow{u} R \xrightarrow{0} R.$$

The rightmost copy is in degree zero.

Proposition 44.2 (An explicit derived obstruction)**.** *For $M = \Lambda/(x)$ as above,*

$$\operatorname{Tor}^{\Lambda}_{2j+1}(M, M) \cong R/uR \quad (j \geq 0), \qquad \operatorname{Tor}^{\Lambda}_{2j}(M, M) = 0 \quad (j \geq 1).$$

In degree zero $M \otimes_\Lambda M \cong R$.

Proof. The displayed free resolution computes the derived tensor product. Its odd differentials are zero, and its positive even differentials are multiplication by u. Since R is a domain and $u \neq 0$, the latter maps are injective. The homology in positive even degree is therefore zero. In odd degree the kernel is all of R and the incoming image is uR, giving R/uR. In degree zero the incoming differential is zero. This establishes all the assertions directly from the definition of homology. □

If u is a unit, the positive Tor groups vanish, and the generic splitting is already available over R. If u is a nonunit, the odd groups record a defect supported on its vanishing locus. Their R-rank is zero, so taking generic ranks erases this information completely. This is a concrete reason to retain the exact category of lattices and derived functors when comparing orders. A rational representation and its integral model can agree in rank while disagreeing substantially in extension theory.

This example also guards against a misleading computational inference. Efficient manipulation of the two small matrices computes this particular resolution. It says nothing about a polynomial-time procedure for arbitrary module classification, arbitrary syzygies, or a cryptographic search problem. The presentation size, coefficient bit lengths, and demanded output are part of any complexity statement.

44.3 Generic splitting and singular specialization

Over the fraction field K of R, the two elements

$$e_x = x/u, \qquad e_y = y/u$$

are orthogonal idempotents with $e_x + e_y = 1$. Thus $K \otimes_R \Lambda \cong K \times K$. The denominators matter: if u is not invertible in R, these idempotents need not belong to Λ itself. The generic decomposition therefore need not be an integral decomposition. This is the same elementary obstruction that appears when an overorder introduces idempotents unavailable in the smaller order.

For $R = k[u]$ or $k[[u]]$, reduction modulo u gives

$$\Lambda/u\Lambda \cong k[X]/(X^2).$$

The special fiber has a nonzero nilpotent, whereas the generic fiber is a product of fields. The periodic maps reduce to multiplication by X and $-X$, whose image and kernel both equal kX. Thus exactness survives this particular reduction. This conclusion follows from an explicit calculation; it must not be generalized to a statement that every exact complex remains exact after arbitrary nonflat base change. Proposition 44.2 already demonstrates the contrary phenomenon for another tensor operation.

44.4 The finite Hodge operator of the same complex

We next work over $\mathbb{R}$ and evaluate the parameter at a real number t. This is a finite matrix family. Equip every copy of $\mathbb{R}^2$ with the standard Euclidean inner product, and take the consecutive differentials

$$A_t = \begin{pmatrix} 0 & 0 \\ 1 & t \end{pmatrix}, \qquad B_t = \begin{pmatrix} t & 0 \\ -1 & 0 \end{pmatrix}.$$

Their composite vanishes for every t, and each matrix has rank one. The middle Hodge Laplacian is

$$\Delta_t = A_t A_t^T + B_t^T B_t.$$

Proposition 44.3 (Exactness and an explicit metric gap)**.** *For every $t \in \mathbb{R}$,*

$$\Delta_t = (1 + t^2) I_2.$$

Its kernel is zero and its least eigenvalue is $1 + t^2 \geq 1$. For every $\eta > 0$, replacing both differentials by $\eta A_t, \eta B_t$ preserves their images, kernels and exactness but replaces the Laplacian by $\eta^2(1 + t^2)I_2$.

Proof. Direct multiplication gives

$$A_t A_t^T = \begin{pmatrix} 0 & 0 \\ 0 & 1 + t^2 \end{pmatrix}, \qquad B_t^T B_t = \begin{pmatrix} 1 + t^2 & 0 \\ 0 & 0 \end{pmatrix}.$$

Adding proves the identity. A scalar multiple of the identity has the stated eigenvalues. Multiplication of a linear map by a nonzero scalar changes neither its image nor its kernel. Its contribution to $AA^T + B^T B$ is multiplied by the square of that scalar, which proves the final assertion. □

Several lessons are visible simultaneously. A singular algebraic fiber does not force this chosen finite Hodge gap to close. Conversely, exactness and the same algebraic rank stratum do not provide a positive lower bound when the scale of the differential is allowed to vary. The inner product and normalization of the operator are essential data. The latter statement is the finite counterpart of the much larger distinction between a topological index and a uniform physical spectral gap. It supplies a counterexample to an inference; it does not model the interacting Yang–Mills Hamiltonian.

44.5 How the cryptographic comparison enters

The derived obstruction R/uR and the cyclic hidden subgroup in the cryptography chapter are both quotients of explicitly presented modules. That common language is useful: one can write maps, calculate kernels and compare exact sequences. Their computational roles are nevertheless different. A presentation of R/uR includes its relation u. In an elliptic curve discrete logarithm instance $Q = dP$, the unknown d is precisely the parameter of the hidden kernel

$$H_d = \{(-dt, t) : t \in \mathbb{Z}/q\mathbb{Z}\}.$$

Writing H_d symbolically does not compute d. Fourier sampling obtains information about its annihilator using a quantum evaluation procedure for the public map $(a, b) \mapsto aP + bQ$. A useful algorithmic bridge must explain how representations of the relevant maps are obtained without first solving the hidden problem.

Similarly, writing down a lattice basis after knowing a secret scalar is not a reduction from an unknown scalar to lattice reduction. The basis of the endomorphism relation lattice in the cryptography chapter is available when the action of the specified endomorphism on the cyclic subgroup is known. That availability is an input condition. The abstract existence of an endomorphism in an order does not automatically provide its efficient evaluation, its scalar action on a selected subgroup, or a short basis in a specified metric.

Chapter 45

A quantitative synthesis and a research program

45.1 The information that each theory actually supplies

The most useful common objects are modules with extra structure and complexes of maps between them. The relevant extra structure changes from chapter to chapter. An order-lattice is an integral module inside a rational representation. A cryptographic lattice is supplied with a norm, representation and probability distribution. An elliptic complex is supplied with differential operators and an exact principal-symbol sequence away from the zero section. A quantum theory is supplied with positive correlation functionals, spacetime covariance, locality and a reconstruction procedure.

These additions are not optional embellishments to an otherwise complete algebraic answer. They determine which questions are meaningful. The matrix equation $BA = 0$ determines a variety of complexes, but no quantum resource cost. A K-theory class determines the stable topological input to an index pairing, but no absolute time scale. A free abelian group determines neither the length of a shortest vector nor the distribution of a sampled error.

Interface	Data retained	Additional data needed for the next question
Order to rational algebra	Scalar extension, representation, rank	Integral conductor, local structure and lattice class
Order to geometric lattice	An embedding into a real vector space	Positive form, basis, coefficient sizes and distortion bounds
Complex to cohomology	Kernel modulo image	Chosen inner products and quantitative singular values
Elliptic symbol to index	Symbol K-class and characteristic pairing	Domains, boundary conditions and estimates on positive spectrum
Group action to quantum algorithm	An abstract operation or oracle	Reversible implementation, cleanup, precision and total resource accounting
Lattice action to LWE problem	Module arithmetic	Error law, parameter regime, sample distribution and a reduction
Euclidean correlations to a QFT	Positive continuum functionals and axioms	Control of phases, quantitative large-distance decay and spectral support

The table organizes a rigorous synthesis because each row is an interface with testable requirements. It prevents a conclusion from one category being silently reinterpreted as a conclusion in another.

45.2 A general comparison estimate

There is a simple quantitative principle behind several of these interfaces. Let V be a finite-dimensional real inner-product space, let Δ be a positive semidefinite self-adjoint operator, and suppose its kernel is K. A gap statement is a lower bound on the orthogonal complement of K:

$$\langle v, \Delta v\rangle \geq c\|v\|^2 \qquad (v \perp K).$$

Algebraic identifications transport vectors. To transport this inequality, they must also control norms and quadratic forms.

Proposition 45.1 (Transport of a coercivity bound)**.** *Let V, W be finite-dimensional real inner-product spaces, $T : W \to V$ a linear isomorphism, and q a nonnegative quadratic form on V with kernel K. Suppose*

$$q(v) \geq c\,\operatorname{dist}(v, K)^2 \quad (v \in V), \qquad \|Tw\| \geq m\|w\| \quad (w \in W)$$

with $c, m > 0$. Put $q_T(w) = q(Tw)$ and $K_T = T^{-1}K$. Then

$$q_T(w) \geq cm^2\,\operatorname{dist}(w, K_T)^2.$$

Proof. For every $k \in K$, write $k = Tz$ with $z \in K_T$. Then $\|Tw - k\| = \|T(w - z)\| \geq m\|w - z\|$. Taking the infimum over k, equivalently over z, gives $\operatorname{dist}(Tw, K) \geq m\operatorname{dist}(w, K_T)$. Apply the assumed bound for q and square the latter inequality. □

This estimate has deliberately modest scope. It does not construct T, prove that m stays away from zero in a limit, or identify a finite matrix with a physical Hamiltonian. It shows exactly what must be added to an algebraic identification to transport a lower spectral bound. The same smallest singular value controls distortion of short vectors in a lattice. The same absence of a uniform bound can invalidate a proposed continuum-limit argument or a proposed cryptographic reduction.

Corollary 45.2 (Why pointwise equivalence is insufficient)**.** *For a family of isomorphisms T_j, positive constants m_j for each fixed j do not imply a uniform positive lower bound for the transported quadratic forms. On $V = W = \mathbb{R}$, with $j \geq 1$ and $q(v) = \|v\|^2$, the choice $T_j = j^{-1}I$ gives $q_{T_j}(w) = j^{-2}\|w\|^2$.*

Proof. The displayed identity follows from homogeneity of the norm. Its least Rayleigh quotient is j^{-2}, which tends to zero. Each T_j remains an isomorphism. □

45.3 A diagram with explicit logical boundaries

There is a defensible route from the two books to the other topics. Integral orders organize endomorphism rings and modules. These produce complexes and deformation problems, whose finite algebra can be studied through rank strata. Smooth geometric counterparts have symbol complexes and index classes. Gauge theory supplies a particularly important elliptic deformation complex. Finite abelian groups and modules also support the character theory underlying quantum

hidden-subgroup algorithms. The links are structural and, in examples such as the preceding chapter, computationally explicit.

What is absent is a theorem making the path transitive at the level of the most difficult conclusions. An endomorphism order does not turn an ASD index into an ECC discrete logarithm algorithm. A positive finite Hodge Laplacian does not establish a uniform infinite-volume Yang–Mills gap. A functional equation and real coefficients do not identify an arithmetic determinant with a self-adjoint operator. These omissions are the mathematical content of the remaining research questions, rather than details to be supplied by an unlabelled analogy.

45.4 Three concrete research directions

45.4.1 Arithmetic modules with explicit metric control

A productive first direction is to choose a family of orders and record, together, its integral presentation, archimedean embeddings, positive form, local conductor, and a basis with certified coefficient sizes. The codifferent calculation in the cryptography chapter provides the correct dual object under a trace pairing. The shear example shows why a module isomorphism alone cannot certify a short-vector comparison. A useful theorem would provide a uniform, explicit distortion bound for a specified family of transformations, with the bound expressed in the same parameter as the input bit length.

For an imaginary quadratic order, one can begin with explicit ideal lattices and the complex absolute-value metric. One must state which ideals are invertible, whether the order is maximal, how the discriminant varies, and which basis or reduced representative is produced. These restrictions are essential for a valid theorem about dual ideals or norm comparison. Extending such a theorem to noncommutative orders requires specifying left versus right modules and the chosen involution or positive representation.

The computational companion provides small exact examples, not an asymptotic complexity result. A subsequent study should measure both coefficient growth and geometric distortion. Reporting only a favorable rank or module decomposition would leave the main quantitative issue unanswered.

45.4.2 Resource-certified elliptic-curve quantum arithmetic

The finite reproducibility program is executed in Chapter 15, within this same monograph. It fixes affine/infinity encoding, total invalid-word behavior, logical gates, clean work, coherent signed decomposition and independent resource accounting. The measured decomposition costs more for the original bounded compiler; this result is retained rather than replaced by an unsupported efficiency claim.

Chapter 43 completes the next field-arithmetic step with uniform carry, borrow, product, reduction and inversion circuits, together with executed small-domain and wider sampled checks. This replaces exponential input-table synthesis for those primitives, without claiming that a complete large-curve elliptic algorithm or fault-tolerant compilation has already been built. Chapter 31 develops the covering, integral-order and learned-channel constructions; Chapter 33 and the completed-steps chapter supply the observed-sky tests and their calibration limits. Their statements, measurements and interactive examples belong to the same continuing document and project page.

45.4.3 From elliptic estimates to a uniform physical gap

A third direction begins with the Yang–Mills existence framework reviewed in this manuscript. The independent assessment separates existence from the mass-gap problem [45]. A further theorem needs a quantitative bound strong enough to survive ultraviolet and infinite-volume limiting procedures, together with a correct identification of the reconstructed observables and vacuum sector.

The spectral theorem argument in the gauge chapter gives a precise target: uniform exponential decay on the orthogonal complement of the vacuum is equivalent to a positive lower spectral threshold, with the stated domain and cyclicity conditions. The UV counterexample there proves that matching leading short-distance behavior cannot alone establish this target. A new mass-gap argument must therefore add infrared information. The ASD index formula may organize classical deformation spaces, but it does not provide the required estimate for the interacting quantum Hamiltonian.

Formalization could assist by making the limiting and reconstruction interfaces explicit. The hard step would remain a proved analytic estimate, not a declaration naming a gap constant. A sound Lean statement must quantify over the relevant cutoffs and volumes, give the correct normalization, and make positivity of the common lower bound an actual conclusion or a clearly identified hypothesis. The formalization appendix explains this distinction using smaller declarations that can be compiled in the present project.

45.5 What this monograph establishes

This work supplies a coherent collection of explicit constructions, elementary proofs, counterexamples to invalid transfers, a source audit, selected machine-checked algebra, and reproducible finite experiments. It does not assert a proof of GRH, a Yang–Mills mass gap, a generally faster ECC algorithm, or a new security reduction for lattice cryptography. Those would require additional theorems not contained in the source material or proved here.

The strongest common principle obtained from the comparison is quantitative: an algebraic or categorical connection becomes an algorithmic or spectral connection only after its representations, metrics, probability laws, normalizations and limiting operations are controlled. The examples make each part of that principle testable. The research directions above specify what a successful next theorem would have to say and what evidence would make it independently assessable.

Appendix A

Reproducible finite experiments and interactive geometry

A.1 What the computations verify

The computational companion has two independent implementations of the principal finite examples: Python for the paper's exact checks and figures, and JavaScript for the interactive page. Both are deliberately small. Their purpose is to make definitions inspectable, expose exceptional cases, and check consistency between formulas and displays. Neither is an implementation of production cryptography or a simulator of continuum Yang–Mills theory.

The Python script `code/verify_examples.py` uses only the standard library.

It writes `verification/examples.json`, including a SHA-256 hash of the script, exact finite counts, and numerical residuals. The figure script additionally uses NumPy and Matplotlib. The browser experiments are buildless and use no remote JavaScript dependency. A separate Node script `site/verify-math.cjs` checks the shared mathematical functions used by the interface. The environment and invocation commands are recorded in the project README.

It is important to distinguish three kinds of evidence. A symbolic proof establishes an identity for all objects under its assumptions. Exhaustive finite enumeration establishes a property of the enumerated finite structure. A floating-point residual checks a sampled numerical implementation up to a tolerance. The companion labels these categories explicitly. For example, checking $BA = 0$ to 10^{-12} at sample points on a projected cone is not a proof of exactness, and a table of 19 elliptic-curve points is not a proof of the group law for every elliptic curve.

A.2 The 19-element elliptic-curve example

The curve

$$E: \quad y^2 = x^3 + 2x + 2 \quad \text{over } \mathbb{F}_{17}$$

has 18 affine points and one point O at infinity. The enumeration checks every pair $(x, y) \in \mathbb{F}_{17}^2$ directly. It also checks that $P = (5, 1)$ has order 19 and generates the entire set. The affine formulas treat the identity, inverse pairs and tangent cases separately. This is essential: blindly inverting a zero denominator would invalidate the implementation.

k	kP	k	kP
0	O	10	$(7, 11)$
1	$(5, 1)$	11	$(13, 10)$
2	$(6, 3)$	12	$(0, 11)$
3	$(10, 6)$	13	$(16, 4)$
4	$(3, 1)$	14	$(9, 1)$
5	$(9, 16)$	15	$(3, 16)$
6	$(16, 13)$	16	$(10, 11)$
7	$(0, 6)$	17	$(6, 14)$
8	$(13, 7)$	18	$(5, 16)$
9	$(7, 6)$	19	O

The exact test enumerates all $19^3 = 6,859$ triples and checks associativity with modular integer arithmetic. It checks closure and commutativity on all pairs. Finally, for each of the 19 choices of d, it enumerates the kernel of $(a, b) \mapsto aP + b(dP)$ and compares it to $\{(-dt, t) : t \in \mathbb{Z}/19\mathbb{Z}\}$. This last comparison connects the concrete curve display to the hidden-subgroup theorem in the text.

The browser lets the reader select a scalar and see its corresponding point. The small modulus makes every operation inspectable. It also makes this example unsuitable as a security parameter: exhaustive enumeration is intentional. The page does not accept real keys or provide an attack service.

A.3 The Fourier annihilator experiment

For the group $(\mathbb{Z}/7\mathbb{Z})^2$ and $d = 3$, the exact Fourier support is

$$(0, 0), (1, 3), (2, 6), (3, 2), (4, 5), (5, 1), (6, 4).$$

Each outcome has probability $1/7$. Every other grid point has probability zero. The exact proof is the geometric-sum identity in the cryptography chapter. The Python computation evaluates the complex exponential sum numerically and compares its squared modulus with that exact probability table. The discrepancy is recorded rather than rounded into a claim of exact floating-point equality.

The interactive version varies the hidden slope d. The seven supported points change accordingly, while the normalization remains one. This is a finite mathematical simulation of the distribution. It does not simulate the physical implementation of a quantum Fourier transform, and the page's knowledge of d is part of the demonstration setup. A real algorithm must obtain its information from the oracle evaluations described in the proof.

A.4 A genuine four-coordinate variety

The variety of complexes of dimension vector $(1, 2, 1)$ has coordinates

$$A = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \qquad B = \begin{pmatrix} y_1 & y_2 \end{pmatrix}, \qquad BA = x_1y_1 + x_2y_2 = 0.$$

The browser displays the real locus in $\mathbb{R}^4$. It does not identify this four-real-dimensional ambient space with the complex variety's real eight-dimensional ambient space. The real cone has dimension three away from its vertex. Its unit-sphere link is two-dimensional.

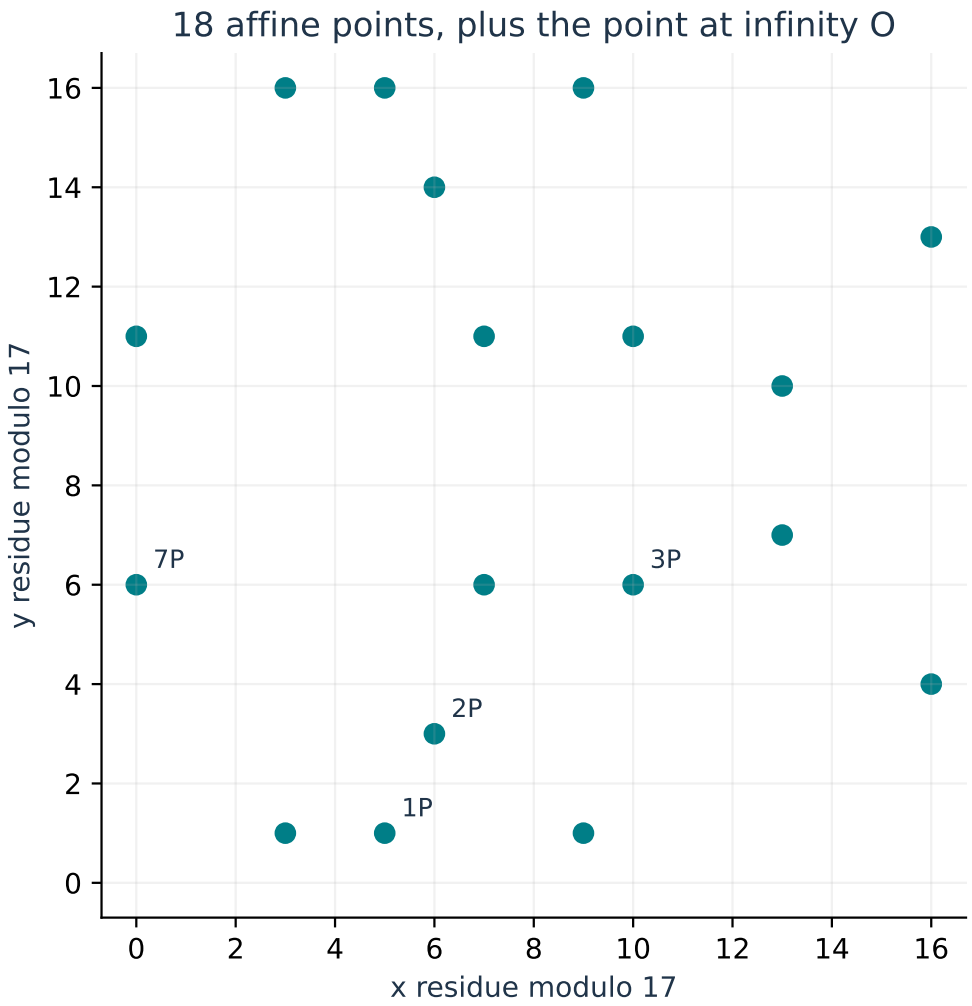


Figure A.1: The finite-field curve is a discrete set of residues. The straight layout of the residue grid is a display convention; it is not the real plane cubic, and Euclidean distances between the plotted dots have no role in its group law. The point O is handled separately.

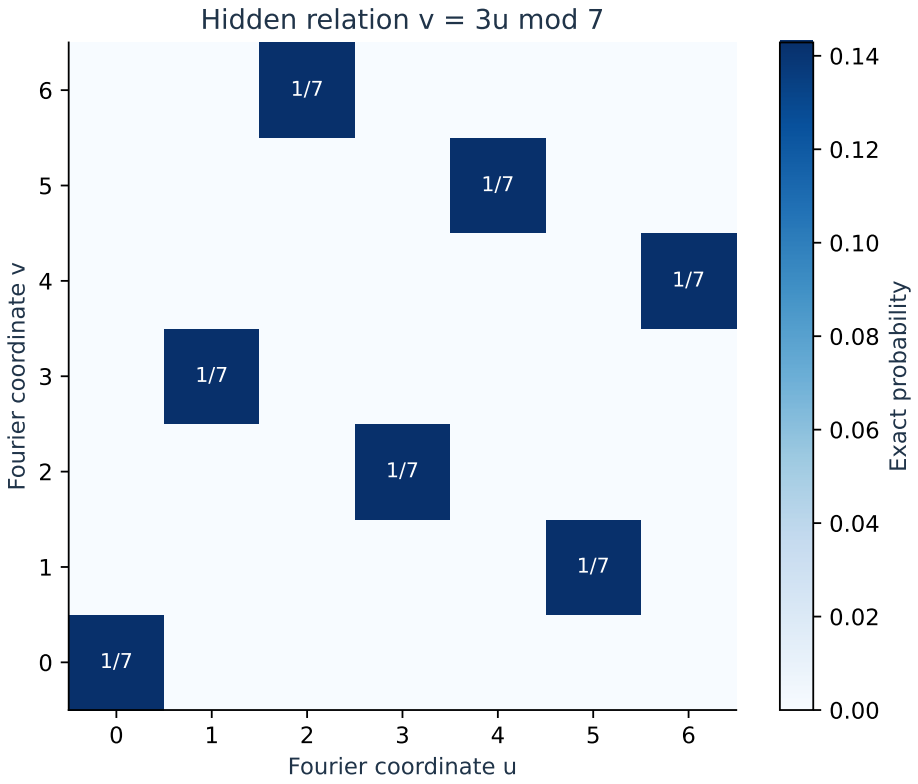


Figure A.2: Fourier mass lies on the annihilator $v = 3u \pmod 7$. The color bar measures probability, not amplitude. Six of the seven possible outcomes have nonzero u and therefore recover d by the stated modular inverse.

Set

$$u = 2^{-1/2}(\cos\alpha, \sin\alpha), \qquad v = 2^{-1/2}(\cos\beta, \sin\beta), \qquad x = (u+v)/\sqrt{2}, \quad y = (u-v)/\sqrt{2}.$$

Then

$$x \cdot y = \tfrac{1}{2}(\|u\|^2 - \|v\|^2) = 0, \qquad \|x\|^2 + \|y\|^2 = \|u\|^2 + \|v\|^2 = 1.$$

The coordinate transformation is orthogonal and invertible. Its inverse sends every point of the unit link to two vectors of norm $1/\sqrt{2}$. Hence the link is $S^1 \times S^1$. This derivation is the reason the parameterization is used; a generic torus-shaped drawing would not establish that it represents the intended algebraic set.

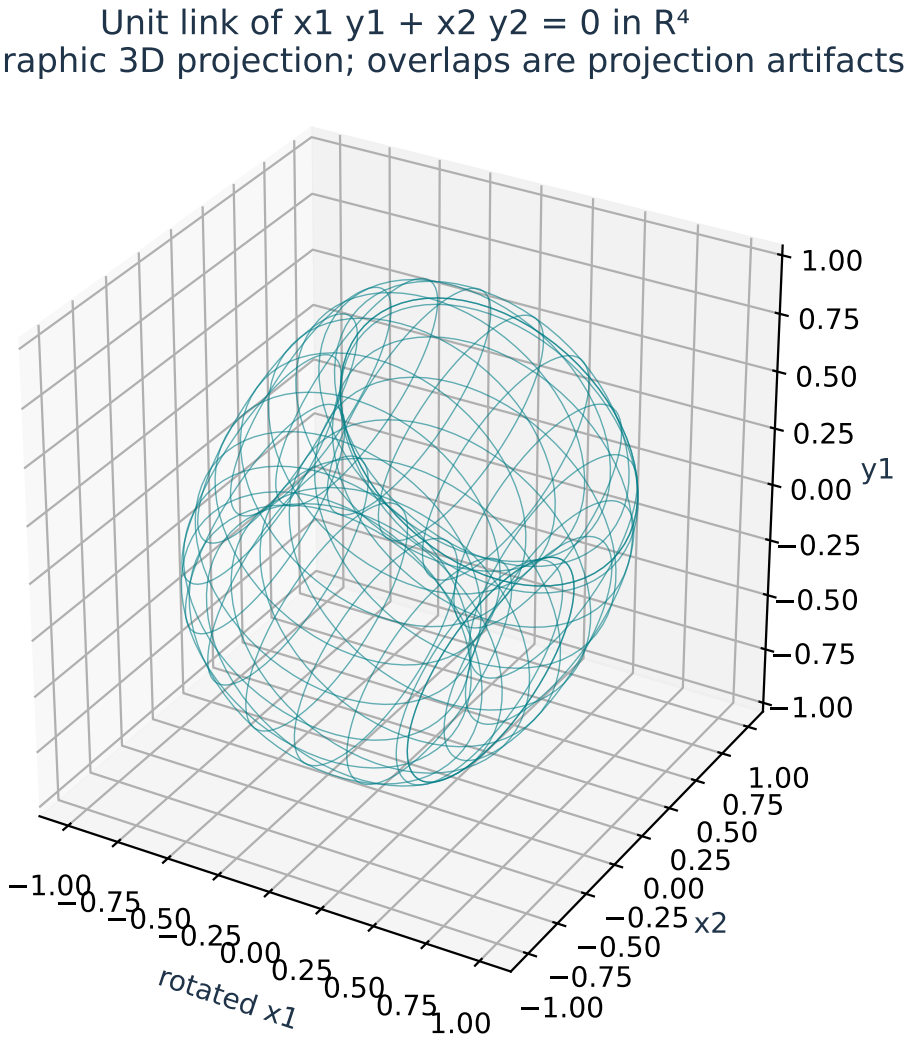


Figure A.3: A specified rotation in four-dimensional coordinate space followed by orthographic projection into three dimensions. Apparent intersections can be caused by projection. The mathematical link lies in $S^3 \subset \mathbb{R}^4$; it is not an embedded surface in ordinary three-dimensional space under every projection.

The page exposes rotation and projection controls and reports the defining equation's residual. Rank strata are derived from whether A and B vanish, rather than inferred from the appearance of

the projected image. The origin is handled separately from the link. The tangent calculation is also separate: a direction satisfying the linearized equation may retain a quadratic obstruction, as proved and partially formalized in the companion.

A.5 Index, kernel and gap are separate readouts

For

$$A_\varepsilon = (\varepsilon, 0)^T, \qquad B = (0, 1),$$

the middle Laplacian is $\mathrm{diag}(\varepsilon^2, 1)$. For $\varepsilon \neq 0$ the complex is exact in the middle and its least eigenvalue is positive. As $\varepsilon \to 0$ the lowest eigenvalue approaches zero. At $\varepsilon = 0$ the kernel dimension changes. The Euler characteristic of the underlying three-term complex remains $1 - 2 + 1 = 0$. At the endpoint, the smallest *positive* eigenvalue is 1, whereas the smallest eigenvalue is 0; these are different readouts and the page distinguishes them.

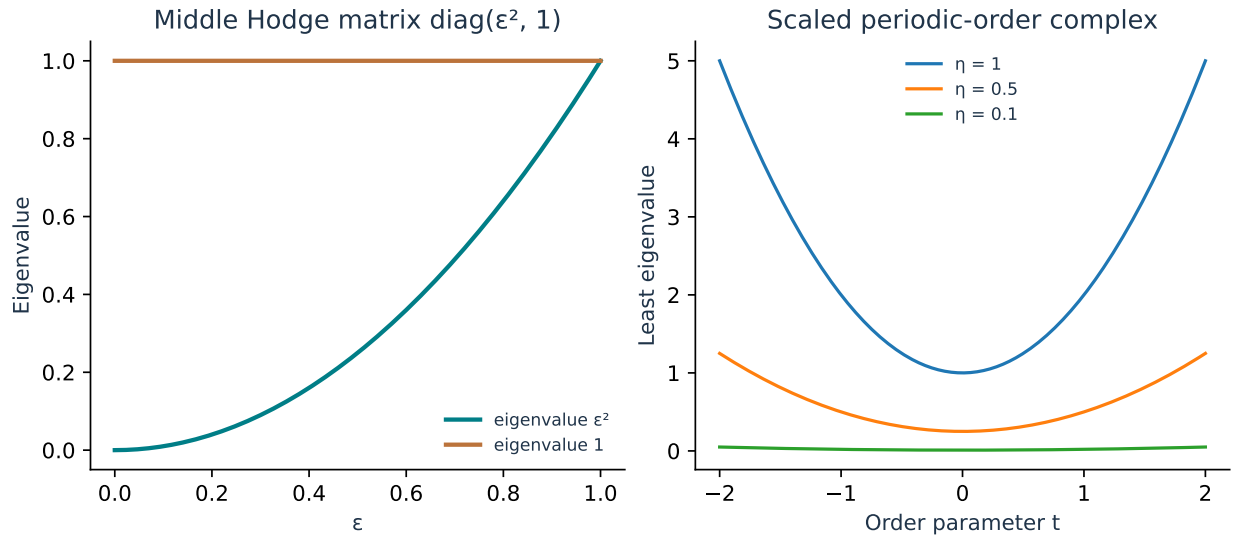


Figure A.4: Left: a fixed Euler characteristic with changing kernel and lower eigenvalue. Right: the exact periodic-order example has eigenvalue $\eta^2(1 + t^2)$ after rescaling both differentials. Algebraic exactness does not determine a metric scale. These are finite matrices, not physical Yang–Mills Hamiltonians.

The twisted flat-torus spectrum provides an infinite-dimensional counterpart whose eigenvalues are still explicit. With $L_1 = L_2 = 1$ and automorphy parameters $(s, 0)$ for $s \in [-1/2, 1/2]$, its lowest scalar eigenvalue is $4\pi^2 s^2$. At the trivial character a zero mode appears. This plot belongs to a free rank-one flat connection, and it must not be read as an interacting nonabelian quantum mass-gap computation.

A.6 The torus and plaquette displays

The elliptic-torus laboratory starts with $\mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$, where $\mathrm{Im}\,\tau > 0$. A lattice Gram matrix records the flat metric of the fundamental parallelogram. The familiar doughnut in three-dimensional space is a topological visualization of the quotient, not an isometric embedding of that flat metric. Changing

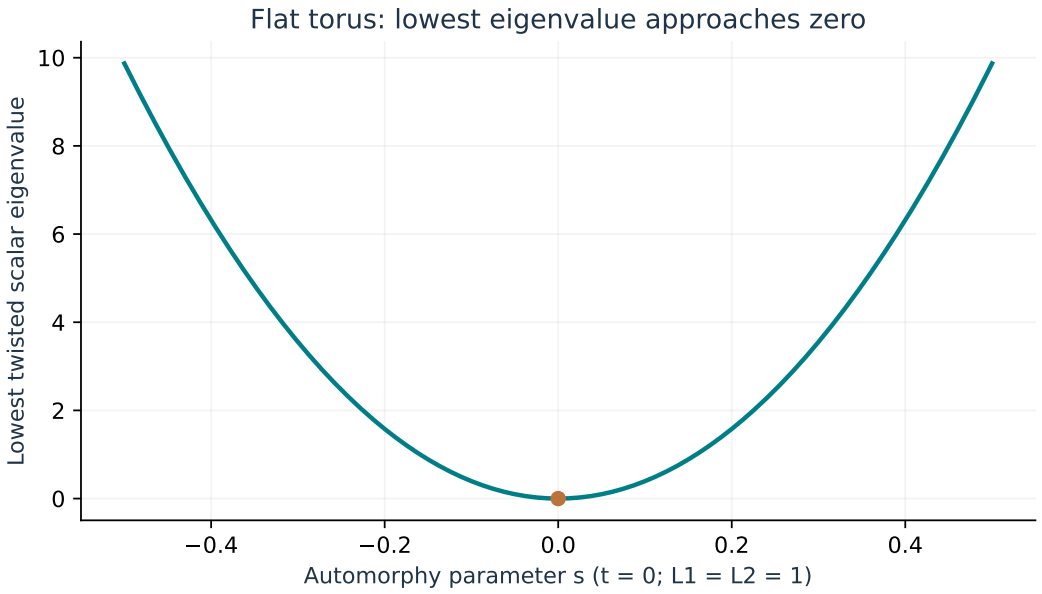


Figure A.5: An exact slice of the flat-torus spectral family. The horizontal coordinate is an automorphy parameter modulo one; it is not time, energy, or a lattice-spacing cutoff.

τ changes the complex structure and the displayed lattice data; an attractive three-dimensional surface cannot substitute for those coordinates.

The optional gauge laboratory uses a single $U(1)$ plaquette. If edge variables are U_i and vertex transformations are g_i, then

$$U_i \longmapsto g_i U_i g_{i+1}^{-1}$$

preserves the oriented product around the plaquette, since consecutive factors cancel. For a nonabelian group the product is conjugated at the base vertex and its trace is invariant. The browser displays the abelian identity and a numerical residual. It does not generate a Yang–Mills ensemble, take a continuum limit, or estimate a physical mass.

A.7 Reproduction and interpretation

The README gives commands for regenerating the JSON evidence and five publication figures. Browser controls use native keyboard-accessible inputs; the camera also supports keyboard rotation. Motion is user-controlled and the reduced-motion preference is respected. Equations, parameter labels and numeric summaries remain available alongside the canvas. The static application uses relative links, so the same files can be served from a local checkout or a project-hosting subdirectory.

Reproduction should begin with the small exact tests. If they fail, the display should not be trusted. If they pass, that establishes consistency for the tested examples, not every theorem in the paper. The next appendix describes the different and stronger guarantee supplied by the selected Lean declarations, together with the exact limits of that guarantee.

A.8 Exact algebraic companion checks

The gentle-invariant verification script checks all $3^{10} = 59{,}049$ ten-coordinate vectors with entries in $\{0, 1, 2\}$, finds the 257 satisfying the four matching equations, and decomposes every one using the constructive proof. It also checks 100 larger constructed examples, the two corrected exponent relations, and 135 exact block determinants for the Kronecker pencil. Direct pair arithmetic in $\mathbb{F}_{17}[u]/(u^2 - 3)$ counts 323 points on the companion elliptic curve, agreeing with the independently derived zeta recurrence. These results are recorded in `extension/verification/invariants.json`.

The interactive algebra companion has separate exact checks: 6,561 small Plucker matrices; all fifteen tableaux of shape $(3, 2)$ on three letters; 2,478 admissible gentle rank cases, tested by independent elimination in characteristics two and seven; and 6,561 matching-generator monomials with 4,529 distinct exponent vectors. These finite tests verify implementations of the proved formulas; they are not substitutes for the unbounded proofs in the algebraic chapters.

The quantum companion and its GPU results are described in the quantum-surface chapter. The browser worker was also exercised through its actual controls: the 42-gate encoder produced support on 256 basis states with norm and all stabilizers equal to one to the displayed precision; an X error on edge zero flipped exactly the two adjacent face checks; and a live gradient update reduced the displayed loss with all sixteen predictions computed. Browser layout checks included a 375-pixel content viewport without page-level horizontal overflow.

Appendix B

The Lean 4 companion: verified statements and explicit limits

B.1 Current seven-module verification

The current companion contains 121 selected theorem declarations in seven modules.

Module	Declarations
`BridgeTheorems`	25
`QuantumInterfaces`	7
`DiagramGrammar`	16
`FluidInterfaces`	34
`ReductionContracts`	16
`TransportContracts`	17
`WeightedGainContracts`	6

The final verification records a fresh compiler check, exact printed types, source hashes and axiom dependencies for every declaration, with no warnings or unexpected axioms. Its authoritative record is `lean/final-verification/verification.json` and the seven module logs. The formal type map identifies 74 declarations with precise partial components of 27 manuscript environments; no whole composite environment is labelled completely formalized. This verification covers the selected local interfaces, not the full upstream continuum proof or cryptographic security.

The inherited local sources have a separate fresh audit of 368 declarations in 37 files, comprising 361 analytic declarations and seven arithmetic declarations. Together with the 121 selected extension declarations, this gives 489 audited declarations across 44 source files. The inherited manifest is `extension/lean/inherited-verification/all/verification.json`; it binds the actual sources, compiler logs and current evidence-policy hash. These counts describe the inspected local packages and do not establish the inherited research program's unresolved arithmetic goals.

The following section preserves the original 25-declaration module's exact verification record, which remains a component of the current package.

B.2 Original module verification result

The file `lean/BridgeTheorems.lean` contains 25 named theorem declarations. It was compiled successfully with Lean 4.24.0 on Linux, using mathlib revision

f897ebcf72cd16f89ab4577d0c826cd14afaafc7.

The compiler exited with status zero. Every declaration has a corresponding `#print axioms` command in the audit block. The resulting dependencies are subsets of the standard set

{propext, Classical.choice, Quot.sound}.

No declaration in this companion uses `sorry`, `sorryAx`, or a new project axiom. The actual compiler output is preserved in `lean/verification.log`; the toolchain, source hash, declaration names and axiom lists are preserved in `lean/verification.json`. Those files describe a completed local check, not a proposed future build.

The source hash of the verified Lean file is

a85b897824b333fd8e7962bfb7959875e1c7e5d7870e4694b02e345855b895ba.

If the file is edited, that hash and the build evidence must be regenerated. The pinned toolchain and dependency commit make the verification target unambiguous. They do not by themselves prove that a later modified file has been checked.

B.3 Why this is selected formalization

The present development formalizes an elementary mathematical spine through the monograph. It does not formalize every theorem cited or proved in the paper. In particular it does not contain a formal proof of the full Buchsbaum–Eisenbud acyclicity criterion, the Atiyah–Singer index theorem, the correctness and complexity of Shor's physical circuit model, a worst-case LWE security reduction, the inherited arithmetic determinant conjectures, or the Yang–Mills construction. It proves neither GRH nor a Yang–Mills mass gap. Stating this scope is necessary to interpret what the compiler has actually certified.

There are three different proof layers in the deliverable. First, the paper gives complete elementary proofs of many finite and algebraic results. Second, it states deep external theorems with citations and their relevant hypotheses. Third, the current companion checks 121 selected formal declarations across the seven modules listed above. The original four modules contribute 82; the reduction, transport and weighted-gain modules add 16, 17 and 6, respectively. Appendix C.1 identifies the earlier fluid-era checks and distinguishes them from the current aggregate verification. These layers overlap, but they are not identical. A citation to a formal upstream project is also different from compiling that project during this work. The Yang–Mills dependency manifest records the frozen external source and assessment, and explicitly states that the full external development was not rebuilt here.

B.4 Linear complexes and rank

The first group of declarations works with arbitrary linear maps $A : V_0 \to V_1$ and $B : V_1 \to V_2$ over a field. The zero-composition hypothesis is an equality of linear maps, not a numerical tolerance:

$$B \circ A = 0.$$

The declaration `complex_range_le_kernel` proves $\operatorname{im} A \subseteq \ker B$. Under finite dimensionality of the middle space, `complex_rank_bound` proves the adjacent-rank inequality, and `middle_exact_iff_rank_sum` proves that equality is equivalent to middle exactness. These statements are useful formal interfaces because they distinguish the complex condition from the stronger exactness condition.

The Euler-rank identity is stated over signed integers. This avoids a common formalization trap: subtraction on natural numbers is truncated, whereas the ordinary algebraic cancellation of alternating dimensions assumes signed arithmetic or hypotheses ensuring every subtraction is valid. The geometric interpretation of the ranks as dimensions of actual cohomology groups is supplied by the preceding linear-algebra results and the paper.

B.5 The cone and its tangent calculation

The definition `cone` is the polynomial $x_1y_1 + x_2y_2$ over a commutative ring. Its perturbation theorem is an exact polynomial identity with a quadratic remainder. Given a point on the cone and a direction satisfying the linearized equation, the second theorem identifies the residual as $t^2(u_1v_1 + u_2v_2)$. It does not incorrectly assert that every tangent vector gives a straight line lying on the variety.

The sum-and-difference identity supplies the algebra behind the real torus-link parameterization. It is stated without division, so it is valid over every commutative ring. The geometric step that divides by $\sqrt{2}$ and identifies circles takes place over $\mathbb{R}$ in the paper. The Lean theorem checks the polynomial part; it does not claim a formal homeomorphism of the link with a torus. Similarly, the two-generator Koszul declaration proves zero composition only. Regular-sequence hypotheses are needed for an exactness theorem and are not silently inferred from commutativity.

B.6 Hodge positivity and the characteristic polynomial

The generic matrix theorem proves that

$$AA^* + B^*B$$

is positive semidefinite for finite real matrices. This statement does not require $BA = 0$. That is mathematically appropriate: positivity of the sum of two squares is weaker than its interpretation as the Laplacian of a complex. The paper uses the complex condition when identifying its kernel with cohomology.

The coordinate calculation uses $A = (x_1, x_2)^T$ and $B = (y_1, y_2)$ and proves the quadratic-form identity

$$z^T\Delta_1 z = (x_1z_1 + x_2z_2)^2 + (y_1z_1 + y_2z_2)^2.$$

The energy is nonnegative, and vanishes exactly when both displayed linear forms vanish. The characteristic determinant is checked before imposing the complex equation:

$$\det(tI - \Delta_1) = (t - \|x\|^2)(t - \|y\|^2) - (x \cdot y)^2.$$

On the cone $x \cdot y = 0$, it factors into the two squared-length factors. Retaining the defect term in the general identity helps prevent accidental use of the factorization off the variety.

The final Hodge declarations check the diagonal family and the existence of arbitrarily small energy along a fixed unit direction. This is a formal finite-dimensional obstruction to inferring a uniform coercivity constant from exactness alone. It is not a formal theorem about the spectrum of an interacting quantum-field Hamiltonian.

B.7 Integral shears, cyclic kernels and noisy arithmetic

The shear declarations compute the determinant, inverse, image of a basis vector, and its squared norm. Together they make precise the distinction between an invertible integral map and a map with bounded Euclidean distortion. They do not formalize a complete lattice-reduction algorithm or its asymptotic runtime.

The cyclic-kernel declarations work in `ZMod q`. They prove the algebraic relation underlying $aP + b(dP) = (a + bd)P$ and the collision equation. The formal file does not construct a quantum state, a unitary Fourier transform, a reversible elliptic-curve circuit, or a success-probability measure. Those require a substantially larger formal specification. The Fourier proof in the paper and the finite numerical experiment have their own clearly stated status.

The noise-cancellation declaration is a matrix identity over a commutative ring, using compatible finite index types. It proves that the public linear terms cancel and isolates the remaining error. The binary-noise bound is a real inequality used in an elementary correctness argument. Neither proves computational indistinguishability, chosen-ciphertext security, or the security of ML-KEM. Correctness algebra and security reductions answer different questions.

B.8 Reading an axiom audit

The appearance of `Classical.choice` in a declaration does not mean that an arbitrary new cryptographic or analytic hypothesis has been added. It identifies use of Lean's standard classical foundation. Likewise `propext` and `Quot.sound` are standard logical principles in the mathlib environment. What would materially change this audit is a new project axiom asserting a difficult target, or an unproved placeholder that introduces `sorryAx`.

Absence of such placeholders is necessary, but is not a complete review of a formal result. One must still read the theorem's type. A theorem of the form "if a uniform gap exists, then a uniform gap exists" can compile without any placeholder and yet supply no new mathematics. An empty index type can make a supposedly universal statement vacuous. An unconnected construction witness can prove an existential object while failing to connect it to the intended lattice approximants. These are specification issues, not failures of the kernel. The external Yang–Mills review discussed in the main text illustrates why keeping the construction witness in the acceptance target matters.

The present file makes each assumption explicit with `set_option autoImplicit false`. Its declarations are short enough to inspect directly. This improves auditability, but it does not erase the need to compare their types with the intended mathematical prose.

B.9 Rebuilding and extending the companion

The tracked `lean-toolchain` and `lakefile.toml` specify the compiler and exact mathlib revision. A conventional build obtains that toolchain, resolves the pinned package, obtains the corresponding compiled cache if desired, and runs `lake build`. The README also records the actual environment used for this verification. Downloaded toolchains and mathlib caches are dependencies, not vendored proof evidence.

Any extension should add a theorem-to-text mapping and preserve the axiom audit. For example, formalizing the nodal periodic resolution would require a quotient algebra, a basis equivalence and module maps, followed by kernel and image computations. Formalizing the torus-link topology would require continuous maps and an inverse homeomorphism. Formalizing the Fourier sampling theorem would require finite-dimensional complex Hilbert spaces, a normalized state, a unitary operator and

a probability interpretation. Each is a meaningful project. Replacing those requirements by a single axiom bearing the desired theorem's name would not satisfy them.

B.10 Seven additional quantum and binary-chain declarations

The additional file `QuantumInterfaces.lean` was compiled against the same pinned toolchain. Its seven declarations have printed theorem types and foundational dependency reports in `quantum-verification.log`; the independent `quantum-verification.json` records its hash and inventory. The original 25-declaration file and its verification manifest remain unchanged. These two files supply the first 32 selected declarations; the second-pass module below adds sixteen.

Three theorems verify $X^2 = Z^2 = I$ and $XZ = -ZX$ for the actual integer matrices $X = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and $Z = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$. Further declarations define the sign character χ of $\mathbb{F}_2$ and prove

$$\chi(a + b) = \chi(a)\chi(b).$$

For a binary vector space C, a linear functional $\ell : C \to \mathbb{F}_2$, and an amplitude function $\psi : C \to R$ with R any commutative ring, the implemented operators are

$$(X_a\psi)(v) = \psi(v + a), \qquad (Z_\ell\psi)(v) = \chi(\ell(v))\psi(v).$$

The theorem `shift_phase_commutation` derives $X_a Z_\ell = \chi(\ell(a))Z_\ell X_a$ as equality of functions. The general coefficient ring includes complex amplitudes, but the formal declaration does not assert an additional analytical Hilbert-space theorem.

For actual maps $d_2 : C_2 \to C_1$ and $d_1 : C_1 \to C_0$ over $\mathbb{F}_2$ with $d_1 d_2 = 0$, `boundary_phase_shift_commute` sets $a = d_2 f$ and uses the functional ℓd_1. Their pairing vanishes by the given composition equation, so the corresponding shift and phase operators commute. The result thus verifies a concrete algebraic interface used in CSS constructions; it does not assume the desired commutation as its own premise. The seventh theorem derives

$$E - (V - 1) - (F - 1) = 2g$$

from the separately supplied integer equation $V - E + F = 2 - 2g$. Surface classification, the ranks of cellular boundary maps, and error-correction performance are outside this arithmetic substitution. No complete quantum circuit or QNN training procedure is claimed formalized.

B.11 Sixteen additional diagram and transport declarations

The second-pass file `DiagramGrammar.lean` was compiled with exit status zero under the same pinned Lean 4.24.0 and mathlib environment. Its actual output is `diagram-verification.log`; its declaration and axiom inventory is `diagram-verification.json`. The source hash is

`871d34de90d99a263094f7abc4aadb51f5012d613c671b04fff43f5a0c35c1c6`.

All 48 declarations across the three modules were also freshly rechecked during the second pass, with source hashes preserved in the audit ledger.

The typed syntax has objects, identity maps, typed generators, and sequential composition. A context has one typed hole. The declarations `context_sound`, `rewriteStep_sound`, and `derivation_sound` prove that semantically equal rules remain equal in such contexts and through finite derivations. Structural identity and associativity rules have their own soundness proof. An explicit family of linear-equivalence generators adds inverse cancellation, and

`equiv_derivation_sound` verifies finite derivations using those actual rules. These statements verify the sequential component of Theorem 17.7; they do not encode the full graphical tensor, dagger, trace, measurement, or channel syntax.

Separate finite-matrix declarations verify tensor interchange, trace factorization, rectangular trace cyclicity, and the dimension value of a closed identity loop. They are algebraic component theorems, not an implicit coherence theorem for arbitrary drawings. The complex-transport declaration uses actual linear equivalences to preserve a zero composite. The exactness-transport declaration preserves image–kernel equality, and `homotopy_implies_exact` derives exactness from an explicit contracting homotopy identity. These hypotheses are meaningful data, not names for the desired conclusion.

Finally `matrix_hodge_conjugacy` proves a finite complex-matrix identity with the necessary endpoint conditions $U_0^\dagger U_0 = I$ and $U_2^\dagger U_2 = I$. Its algebraic identity allows an arbitrary middle matrix U_1; unitarity of U_1 is the further condition used for the spectral interpretation in Proposition 17.9. The formal statement therefore supports the displayed finite conjugacy component. It does not assert an infinite-dimensional gap or a continuum Yang–Mills result.

B.12 Complete verified companion sources

The following listing reproduces the verified file so that the mathematical types, proof scripts and audit commands can be read together. The separate source file is the authoritative build input.

```
import Mathlib.LinearAlgebra.FiniteDimensional.Lemmas
import Mathlib.LinearAlgebra.Matrix.PosDef
import Mathlib.LinearAlgebra.Matrix.Notation
import Mathlib.LinearAlgebra.Matrix.Determinant.Basic
import Mathlib.Data.ZMod.Basic
import Mathlib.Tactic.Ring
import Mathlib.Tactic.Linarith
import Mathlib.Tactic.FinCases

/-!
# Selected exact interfaces for the research monograph

These declarations formalize elementary statements used in the exposition. They
do not formalize the Buchsbaum-Eisenbud acyclicity theorem, Atiyah-Singer theorem,
Shor's quantum circuit, an LWE security reduction, or quantum Yang-Mills theory.

The scalar cone models complexes of dimension vector (1,2,1). The finite matrix
Laplacian is the middle Hodge Laplacian. The ZMod computation is the algebra of a
cyclic hidden subgroup, and the cancellation identity is only the correctness
algebra of a linear encryption template. No security claim follows from it.

Each named theorem is followed by `#print axioms` in the final audit block.
-/

set_option autoImplicit false

open Matrix
open scoped Matrix ComplexOrder

namespace ResearchBridge

section LinearComplex

variable {K V₀ V₁ V₂ : Type*} [Field K]
variable [AddCommGroup V₀] [AddCommGroup V₁] [AddCommGroup V₂]
variable [Module K V₀] [Module K V₁] [Module K V₂]

/-- A zero composite puts the incoming image inside the outgoing kernel. -/
theorem complex_range_le_kernel (A : V₀ →ₗ[K] V₁) (B : V₁ →ₗ[K] V₂)
```

```
    (h : B.comp A = 0) : LinearMap.range A ≤ LinearMap.ker B :=
  LinearMap.range_le_ker_iff.mpr h

/-- The adjacent ranks of a finite-dimensional complex cannot exceed the
dimension of the middle vector space. -/
theorem complex_rank_bound [FiniteDimensional K V₁]
    (A : V₀ →ₗ[K] V₁) (B : V₁ →ₗ[K] V₂) (h : B.comp A = 0) :
    Module.finrank K (LinearMap.range A) + Module.finrank K (LinearMap.range B)
      ≤ Module.finrank K V₁ := by
  have hr := Submodule.finrank_mono (complex_range_le_kernel A B h)
  have hn := B.finrank_range_add_finrank_ker
  omega

/-- With the complex condition already present, middle exactness is equivalent
to equality in the rank bound. -/
theorem middle_exact_iff_rank_sum [FiniteDimensional K V₁]
    (A : V₀ →ₗ[K] V₁) (B : V₁ →ₗ[K] V₂) (h : B.comp A = 0) :
    LinearMap.range A = LinearMap.ker B ↔
      Module.finrank K (LinearMap.range A) + Module.finrank K (LinearMap.range B)
        = Module.finrank K V₁ := by
  have hn := B.finrank_range_add_finrank_ker
  constructor
  · intro he
    rw [he]
    omega
  · intro he
    apply Submodule.eq_of_le_of_finrank_eq (complex_range_le_kernel A B h)
    omega

end LinearComplex

/-- Once the cohomology dimensions are expressed as a-r, b-r-s, c-s,
their alternating sum is independent of the ranks. Signed integers avoid
silently using truncated subtraction. -/
theorem euler_rank_identity (a b c r s : ℤ) :
    (a - r) - (b - r - s) + (c - s) = a - b + c := by
  ring

section Cone

variable {R : Type*} [CommRing R]

/-- The (1,2,1) complex equation in four scalar coordinates. -/
def cone (x₁ x₂ y₁ y₂ : R) : R := x₁ * y₁ + x₂ * y₂

/-- The exact polynomial expansion, including its quadratic remainder. -/
theorem cone_perturbation (x₁ x₂ y₁ y₂ u₁ u₂ v₁ v₂ t : R) :
    cone (x₁ + t * u₁) (x₂ + t * u₂) (y₁ + t * v₁) (y₂ + t * v₂) =
      cone x₁ x₂ y₁ y₂ +
        t * (y₁ * u₁ + y₂ * u₂ + v₁ * x₁ + v₂ * x₂) +
        t ^ 2 * cone u₁ u₂ v₁ v₂ := by
  simp only [cone]
  ring

/-- A tangent direction removes the constant and linear terms, but generally
leaves a quadratic obstruction. -/
theorem cone_tangent_second_order (x₁ x₂ y₁ y₂ u₁ u₂ v₁ v₂ t : R)
    (h₀ : cone x₁ x₂ y₁ y₂ = 0)
    (h₁ : y₁ * u₁ + y₂ * u₂ + v₁ * x₁ + v₂ * x₂ = 0) :
    cone (x₁ + t * u₁) (x₂ + t * u₂) (y₁ + t * v₁) (y₂ + t * v₂) =
      t ^ 2 * cone u₁ u₂ v₁ v₂ := by
  rw [cone_perturbation, h₀, h₁]
  ring

/-- Unnormalized sum/difference coordinates diagonalize the split quadric. -/
theorem cone_split_identity (u₁ u₂ v₁ v₂ : R) :
    cone (u₁ + v₁) (u₂ + v₂) (u₁ - v₁) (u₂ - v₂) =
      (u₁ ^ 2 + u₂ ^ 2) - (v₁ ^ 2 + v₂ ^ 2) := by
```

```
  simp only [cone]
  ring

/-- This is the equal-radius equation underlying the torus link. The
topological assertion that the link is a torus is not formalized here. -/
theorem cone_equal_norm_iff (u₁ u₂ v₁ v₂ : R) :
    cone (u₁ + v₁) (u₂ + v₂) (u₁ - v₁) (u₂ - v₂) = 0 ↔
      u₁ ^ 2 + u₂ ^ 2 = v₁ ^ 2 + v₂ ^ 2 := by
  rw [cone_split_identity, sub_eq_zero]

/-- The two-generator Koszul maps compose to zero over any commutative ring.
Exactness requires additional regular-sequence assumptions. -/
theorem koszul_zero_composite (x y c : R) : x * (-y * c) + y * (x * c) = 0 := by
  ring

end Cone

section Hodge

/-- Middle Hodge matrix for A=(x₁,x₂)^T and B=(y₁,y₂). -/
def hodgeMatrix (x₁ x₂ y₁ y₂ : ℝ) : Matrix (Fin 2) (Fin 2) ℝ :=
  !![x₁ ^ 2 + y₁ ^ 2, x₁ * x₂ + y₁ * y₂;
     x₁ * x₂ + y₁ * y₂, x₂ ^ 2 + y₂ ^ 2]

def hodgeEnergy (x₁ x₂ y₁ y₂ z₁ z₂ : ℝ) : ℝ :=
  (x₁ * z₁ + x₂ * z₂) ^ 2 + (y₁ * z₁ + y₂ * z₂) ^ 2

/-- Every finite Hodge matrix is positive semidefinite. The complex condition
BA=0 is not needed for this particular conclusion. -/
theorem hodge_positive_semidefinite {m n p : Type*}
    [Fintype m] [Fintype n] [Fintype p]
    (A : Matrix n m ℝ) (B : Matrix p n ℝ) :
    (A * Aᴴ + Bᴴ * B).PosSemidef :=
  (Matrix.posSemidef_self_mul_conjTranspose A).add
    (Matrix.posSemidef_conjTranspose_mul_self B)

/-- The middle quadratic form is the sum of the two squared differential
components, with coordinates made explicit. -/
theorem hodge_quadratic_identity (x₁ x₂ y₁ y₂ z₁ z₂ : ℝ) :
    ![z₁, z₂] ᵥ* (hodgeMatrix x₁ x₂ y₁ y₂ *ᵥ ![z₁, z₂]) =
      hodgeEnergy x₁ x₂ y₁ y₂ z₁ z₂ := by
  simp [hodgeMatrix, hodgeEnergy, Matrix.mulVec, dotProduct, Fin.sum_univ_two]
  ring

theorem hodge_energy_nonnegative (x₁ x₂ y₁ y₂ z₁ z₂ : ℝ) :
    0 ≤ hodgeEnergy x₁ x₂ y₁ y₂ z₁ z₂ := by
  exact add_nonneg (sq_nonneg _) (sq_nonneg _)

/-- Zero energy means both differential components vanish. -/
theorem hodge_energy_zero_iff (x₁ x₂ y₁ y₂ z₁ z₂ : ℝ) :
    hodgeEnergy x₁ x₂ y₁ y₂ z₁ z₂ = 0 ↔
      x₁ * z₁ + x₂ * z₂ = 0 ∧ y₁ * z₁ + y₂ * z₂ = 0 := by
  constructor
  · intro h
    unfold hodgeEnergy at h
    have hx := sq_nonneg (x₁ * z₁ + x₂ * z₂)
    have hy := sq_nonneg (y₁ * z₁ + y₂ * z₂)
    constructor <;> nlinarith
  · rintro ⟨hx, hy⟩
    simp [hodgeEnergy, hx, hy]

/-- The characteristic determinant includes the squared complex defect. -/
theorem hodge_characteristic_identity (x₁ x₂ y₁ y₂ t : ℝ) :
    Matrix.det (t • (1 : Matrix (Fin 2) (Fin 2) ℝ) - hodgeMatrix x₁ x₂ y₁ y₂) =
      (t - (x₁ ^ 2 + x₂ ^ 2)) * (t - (y₁ ^ 2 + y₂ ^ 2)) -
        (cone x₁ x₂ y₁ y₂) ^ 2 := by
  simp [Matrix.det_fin_two, hodgeMatrix, cone]
  ring
```

```
/-- On the variety of complexes, the middle characteristic determinant
factors into the two squared-length eigenvalue factors. -/
theorem hodge_characteristic_factor (x₁ x₂ y₁ y₂ t : ℝ)
    (h : cone x₁ x₂ y₁ y₂ = 0) :
    Matrix.det (t • (1 : Matrix (Fin 2) (Fin 2) ℝ) - hodgeMatrix x₁ x₂ y₁ y₂) =
      (t - (x₁ ^ 2 + x₂ ^ 2)) * (t - (y₁ ^ 2 + y₂ ^ 2)) := by
  rw [hodge_characteristic_identity, h]
  ring

theorem hodge_family_diagonal (t : ℝ) :
    hodgeMatrix t 0 0 1 = !![t ^ 2, 0; 0, 1] := by
  simp [hodgeMatrix]

/-- Positive t keeps the two scalar maps nonzero, but the energy in a fixed
unit direction can be arbitrarily small. This is a finite-dimensional
counterexample to deriving a uniform positive energy bound from rank alone. -/
theorem hodge_family_arbitrarily_small_energy (ε : ℝ) (hε : 0 < ε) :
    ∃ t : ℝ, 0 < t ∧ cone t 0 0 1 = 0 ∧
      hodgeEnergy t 0 0 1 1 0 < ε := by
  let t : ℝ := min ε 1 / 2
  have hp : 0 < min ε 1 := lt_min hε zero_lt_one
  have h₁ : min ε 1 ≤ ε := min_le_left _ _
  have h₂ : min ε 1 ≤ 1 := min_le_right _ _
  have hsq : (min ε 1) ^ 2 ≤ min ε 1 := by
    nlinarith [mul_nonneg (le_of_lt hp) (sub_nonneg.mpr h₂)]
  refine ⟨t, by dsimp [t]; linarith, by simp [cone], ?_⟩
  simp only [hodgeEnergy, mul_one, zero_mul, add_zero, zero_add]
  dsimp [t]
  nlinarith

end Hodge

section Shear

/-- An integral basis change with determinant one. -/
def shear (n : ℤ) : Matrix (Fin 2) (Fin 2) ℤ := !![1, n; 0, 1]

theorem shear_determinant_one (n : ℤ) : Matrix.det (shear n) = 1 := by
  simp [shear, Matrix.det_fin_two]

theorem shear_inverse (n : ℤ) : shear n * shear (-n) = 1 := by
  ext i j
  fin_cases i <;> fin_cases j <;>
    simp [shear, Matrix.mul_apply, Fin.sum_univ_two]

theorem shear_e2_image (n : ℤ) : shear n *ᵥ ![0, 1] = ![n, 1] := by
  ext i
  fin_cases i <;> simp [shear, Matrix.mulVec, dotProduct, Fin.sum_univ_two]

/-- The squared Euclidean norm of the transformed unit basis vector grows
quadratically even though the integral determinant remains one. -/
theorem shear_e2_norm_squared (n : ℤ) :
    ((shear n *ᵥ ![0, 1]) 0 : ℝ) ^ 2 +
      ((shear n *ᵥ ![0, 1]) 1 : ℝ) ^ 2 = (n : ℝ) ^ 2 + 1 := by
  rw [shear_e2_image]
  simp

end Shear

section CyclicKernel

variable {q : ℕ}

/-- Coordinates for the additive cyclic hidden-subgroup equation. -/
def hiddenLinear (x a b : ZMod q) : ZMod q := a + x * b

theorem hidden_kernel_relation (x a b : ZMod q) :
```

```
    hiddenLinear x a b = 0 ↔ a = -(x * b) := by
  simp only [hiddenLinear, add_eq_zero_iff_eq_neg]

/-- Equal function values differ by an element of the cyclic kernel. -/
theorem hidden_collision_relation (x a b c d : ZMod q) :
    hiddenLinear x a b = hiddenLinear x c d ↔
      (a - c) + x * (b - d) = 0 := by
  constructor
  · intro h
    calc
      (a - c) + x * (b - d) = hiddenLinear x a b - hiddenLinear x c d := by
        simp only [hiddenLinear]
        ring
      _ = 0 := sub_eq_zero.mpr h
  · intro h
    apply sub_eq_zero.mp
    calc
      hiddenLinear x a b - hiddenLinear x c d = (a - c) + x * (b - d) := by
        simp only [hiddenLinear]
        ring
      _ = 0 := h

end CyclicKernel

section Noise

/-- Algebraic cancellation in a linear encryption template. This identity
holds over any commutative ring, including modular rings; correctness of
decoding additionally requires a quantitative noise bound. -/
theorem noise_cancellation {R m n : Type*} [CommRing R]
    [Fintype m] [Fintype n]
    (A : Matrix m n R) (s : n → R) (e r : m → R) (μ : R) :
    (((A *ᵥ s + e) ⬝ᵥ r) + μ) - (s ⬝ᵥ (Aᵀ *ᵥ r)) =
      e ⬝ᵥ r + μ := by
  rw [add_dotProduct, Matrix.mulVec_transpose, dotProduct_comm s,
    ← Matrix.dotProduct_mulVec, dotProduct_comm r]
  ring

/-- A concrete two-coordinate error bound for binary selector coefficients.
It is a correctness estimate, not an LWE hardness assertion. -/
theorem binary_noise_bound (e₁ e₂ r₁ r₂ E : ℝ)
    (hE : 0 ≤ E) (h₁ : |e₁| ≤ E) (h₂ : |e₂| ≤ E)
    (hr₁ : r₁ = 0 ∨ r₁ = 1) (hr₂ : r₂ = 0 ∨ r₂ = 1) :
    |e₁ * r₁ + e₂ * r₂| ≤ 2 * E := by
  rcases hr₁ with rfl | rfl <;> rcases hr₂ with rfl | rfl
  · simpa using (show 0 ≤ 2 * E by linarith)
  · simpa using (show |e₂| ≤ 2 * E by linarith)
  · simpa using (show |e₁| ≤ 2 * E by linarith)
  · simpa using (show |e₁ + e₂| ≤ 2 * E from
      le_trans (abs_add_le e₁ e₂) (by linarith))

end Noise

end ResearchBridge

-- Kernel-auditable declaration inventory. Standard foundational axioms from
-- Lean/mathlib may occur; there must be no ‘sorryAx‘ or project-specific axiom.
#print axioms ResearchBridge.complex_range_le_kernel
#print axioms ResearchBridge.complex_rank_bound
#print axioms ResearchBridge.middle_exact_iff_rank_sum
#print axioms ResearchBridge.euler_rank_identity
#print axioms ResearchBridge.cone_perturbation
#print axioms ResearchBridge.cone_tangent_second_order
#print axioms ResearchBridge.cone_split_identity
#print axioms ResearchBridge.cone_equal_norm_iff
#print axioms ResearchBridge.koszul_zero_composite
#print axioms ResearchBridge.hodge_positive_semidefinite
#print axioms ResearchBridge.hodge_quadratic_identity
```

```
#print axioms ResearchBridge.hodge_energy_nonnegative
#print axioms ResearchBridge.hodge_energy_zero_iff
#print axioms ResearchBridge.hodge_characteristic_identity
#print axioms ResearchBridge.hodge_characteristic_factor
#print axioms ResearchBridge.hodge_family_diagonal
#print axioms ResearchBridge.hodge_family_arbitrarily_small_energy
#print axioms ResearchBridge.shear_determinant_one
#print axioms ResearchBridge.shear_inverse
#print axioms ResearchBridge.shear_e2_image
#print axioms ResearchBridge.shear_e2_norm_squared
#print axioms ResearchBridge.hidden_kernel_relation
#print axioms ResearchBridge.hidden_collision_relation
#print axioms ResearchBridge.noise_cancellation
#print axioms ResearchBridge.binary_noise_bound
```

The additional quantum and binary-chain file follows. Its source and audit are independent of the preceding file.

```
import Mathlib.LinearAlgebra.Matrix.Notation
import Mathlib.LinearAlgebra.FiniteDimensional.Lemmas
import Mathlib.Data.ZMod.Basic
import Mathlib.Tactic.FinCases
import Mathlib.Tactic.Ring
import Mathlib.Tactic.Linarith

/-!
# Selected finite quantum and homological interfaces

The Pauli calculations use actual integer matrices. The binary-character
calculations act on functions with values in an arbitrary commutative ring
(and therefore specialize to complex amplitudes). Boundary commutation is
derived from the zero-composition equation for actual linear maps over ZMod 2.
The Euler statement is only an integer substitution, not a formal proof of
surface classification, cellular homology, or a quantum error-correction theorem.
No full QNN training algorithm, quantum circuit, or physical reconstruction is
formalized here. Every named theorem has a printed foundational dependency audit.
-/

set_option autoImplicit false

namespace ResearchQuantum

def pauliX : Matrix (Fin 2) (Fin 2) ℤ := !![0, 1; 1, 0]
def pauliZ : Matrix (Fin 2) (Fin 2) ℤ := !![1, 0; 0, -1]

theorem pauliX_square : pauliX * pauliX = 1 := by
  decide

theorem pauliZ_square : pauliZ * pauliZ = 1 := by
  decide

theorem pauli_anticommutation : pauliX * pauliZ = -(pauliZ * pauliX) := by
  decide

section BinaryCharacters

variable {R C : Type*} [CommRing R]
variable [AddCommGroup C] [Module (ZMod 2) C]

/-- The sign character of the additive two-element group. -/
def binarySign (a : ZMod 2) : R := if a = 0 then 1 else -1

theorem binarySign_add (a b : ZMod 2) :
    binarySign (R := R) (a + b) = binarySign (R := R) a * binarySign (R := R) b := by
  have htwo : (2 : ZMod 2) = 0 := by decide
  fin_cases a <;> fin_cases b <;> norm_num [binarySign, htwo]
  intro h
  exact (h htwo).elim
```

```
/-- Translation of the binary configuration indexing an amplitude function. -/
def bitShift (a : C) (psi : C → R) : C → R := fun v => psi (v + a)

/-- The phase operator associated with a binary linear functional. -/
def parityPhase (ell : C →₁[ZMod 2] ZMod 2) (psi : C → R) : C → R :=
  fun v => binarySign (ell v) * psi v

theorem shift_phase_commutation (a : C) (ell : C →₁[ZMod 2] ZMod 2)
    (psi : C → R) :
    bitShift a (parityPhase ell psi) =
      fun v => binarySign (ell a) * parityPhase ell (bitShift a psi) v := by
  funext v
  simp only [bitShift, parityPhase, map_add, binarySign_add]
  ring

end BinaryCharacters

section ChainComplex

variable {R C0 C1 C2 : Type*} [CommRing R]
variable [AddCommGroup C0] [AddCommGroup C1] [AddCommGroup C2]
variable [Module (ZMod 2) C0] [Module (ZMod 2) C1] [Module (ZMod 2) C2]

/-- Binary boundary and coboundary operators commute as phase/shift actions
because the actual successive linear boundary maps compose to zero. -/
theorem boundary_phase_shift_commute
    (d1 : C1 →₁[ZMod 2] C0) (d2 : C2 →₁[ZMod 2] C1)
    (hcomplex : d1.comp d2 = 0) (ell : C0 →₁[ZMod 2] ZMod 2)
    (f : C2) (psi : C1 → R) :
    bitShift (d2 f) (parityPhase (ell.comp d1) psi) =
      parityPhase (ell.comp d1) (bitShift (d2 f) psi) := by
  have hd : d1 (d2 f) = 0 := by
    simpa using congrArg (fun L : C2 →₁[ZMod 2] C0 => L f) hcomplex
  rw [shift_phase_commutation]
  funext v
  simp [LinearMap.comp_apply, hd, binarySign]

end ChainComplex

/-- This is the arithmetic substitution after ranks V-1 and F-1 and Euler
characteristic 2-2g have been separately justified in surface topology. -/
theorem euler_surface_rank_substitution (V E F g : ℤ)
    (hEuler : V - E + F = 2 - 2 * g) :
    E - (V - 1) - (F - 1) = 2 * g := by
  linarith

end ResearchQuantum

#check ResearchQuantum.pauliX_square
#check ResearchQuantum.pauliZ_square
#check ResearchQuantum.pauli_anticommutation
#check ResearchQuantum.binarySign_add
#check ResearchQuantum.shift_phase_commutation
#check ResearchQuantum.boundary_phase_shift_commute
#check ResearchQuantum.euler_surface_rank_substitution

#print axioms ResearchQuantum.pauliX_square
#print axioms ResearchQuantum.pauliZ_square
#print axioms ResearchQuantum.pauli_anticommutation
#print axioms ResearchQuantum.binarySign_add
#print axioms ResearchQuantum.shift_phase_commutation
#print axioms ResearchQuantum.boundary_phase_shift_commute
#print axioms ResearchQuantum.euler_surface_rank_substitution
```

The second-pass typed diagram, matrix contraction, and complex-transport module follows. Its exact statement types and foundational dependency commands are part of the authoritative source.

```
import Mathlib.LinearAlgebra.FiniteDimensional.Lemmas
import Mathlib.Data.Complex.Basic
import Mathlib.LinearAlgebra.Matrix.Kronecker
import Mathlib.LinearAlgebra.Matrix.Trace

/-!
# A typed sequential fragment of diagram rewriting

Objects are labels and each label is assigned a module. The syntax below has
only identity, typed generators, and sequential composition. A context has one
typed hole; replacing that hole by an equal linear map preserves denotation.
Sound basic rewrite rules therefore remain sound under every finite sequence of
contextual rewrites. No termination, confluence, semantic completeness, tensor
word syntax, graphical isotopy, traced symmetric-monoidal coherence, dagger,
measurement, classical control, or CP/TP channel theorem is asserted here.

Separate finite-matrix lemmas record tensor interchange and closed-loop
contractions. They are exact algebraic components, not an identification of the
sequential syntax with a full tensor graph grammar. Algebraic changes of bases
preserve complexes and exactness; a separate finite matrix identity gives
Hodge conjugacy under its stated unitary endpoint assumptions. No uniform
analytic bound or infinite-dimensional spectral theorem is asserted.
-/

set_option autoImplicit false

universe u v w z

namespace DiagramGrammar

inductive Diagram {Obj : Type u} (Gen : Obj → Obj → Type v) :
    Obj → Obj → Type (max u v) where
  | ident (a : Obj) : Diagram Gen a a
  | atom {a b : Obj} (g : Gen a b) : Diagram Gen a b
  | seq {a b c : Obj} (f : Diagram Gen a b) (g : Diagram Gen b c) : Diagram Gen a c

inductive Context {Obj : Type u} (Gen : Obj → Obj → Type v) (a b : Obj) :
    Obj → Obj → Type (max u v) where
  | hole : Context Gen a b a b
  | before {x y z : Obj} (f : Diagram Gen x y) (C : Context Gen a b y z) :
      Context Gen a b x z
  | after {x y z : Obj} (C : Context Gen a b x y) (f : Diagram Gen y z) :
      Context Gen a b x z

def fill {Obj : Type u} {Gen : Obj → Obj → Type v} {a b x y : Obj}
    (C : Context Gen a b x y) (d : Diagram Gen a b) : Diagram Gen x y :=
  match C with
  | .hole => d
  | .before f C => .seq f (fill C d)
  | .after C f => .seq (fill C d) f

section Semantics

variable {Obj : Type u} {Gen : Obj → Obj → Type v}
variable {K : Type z} [Semiring K]
variable {V : Obj → Type w}
variable [∀ a, AddCommMonoid (V a)] [∀ a, Module K (V a)]

def denote (rho : ∀ (a b : Obj), Gen a b → V a →ₗ[K] V b) :
    ∀ {a b}, Diagram Gen a b → V a →ₗ[K] V b
  | _, _, .ident _ => LinearMap.id
  | _, _, .atom g => rho _ _ g
  | _, _, .seq f g => (denote rho g).comp (denote rho f)

/-- Equality of the interpreted hole is preserved by every typed sequential context. -/
theorem context_sound (rho : ∀ (a b : Obj), Gen a b → V a →ₗ[K] V b)
    {a b x y : Obj} {d e : Diagram Gen a b}
```

```
    (h : denote rho d = denote rho e) (C : Context Gen a b x y) :
    denote rho (fill C d) = denote rho (fill C e) := by
  induction C with
  | hole => exact h
  | before f C ih => simp only [fill, denote, ih]
  | after C f ih => simp only [fill, denote, ih]

inductive StructuralRule : ∀ {a b : Obj}, Diagram Gen a b → Diagram Gen a b → Prop
  | left_id {a b} (d : Diagram Gen a b) :
      StructuralRule (.seq (.ident a) d) d
  | right_id {a b} (d : Diagram Gen a b) :
      StructuralRule (.seq d (.ident b)) d
  | assoc {a b c d} (f : Diagram Gen a b) (g : Diagram Gen b c) (h : Diagram Gen c d) :
      StructuralRule (.seq (.seq f g) h) (.seq f (.seq g h))

/-- The actual identity and associativity rewrite rules are semantically sound. -/
theorem structural_sound (rho : ∀ (a b : Obj), Gen a b → V a →₁[K] V b)
    {a b : Obj} {d e : Diagram Gen a b} (h : StructuralRule d e) :
    denote rho d = denote rho e := by
  cases h <;> ext v <;> rfl

inductive RewriteStep
    (Rules : ∀ {a b : Obj}, Diagram Gen a b → Diagram Gen a b → Prop) :
    ∀ {x y : Obj}, Diagram Gen x y → Diagram Gen x y → Prop
  | in_context {a b x y} {d e : Diagram Gen a b}
      (h : Rules d e) (C : Context Gen a b x y) :
      RewriteStep Rules (fill C d) (fill C e)

def Derives
    (Rules : ∀ {a b : Obj}, Diagram Gen a b → Diagram Gen a b → Prop)
    {x y : Obj} (d e : Diagram Gen x y) : Prop :=
  Relation.ReflTransGen (RewriteStep Rules) d e

/-- A single basic sound rule can be applied within any sequential context. -/
theorem rewriteStep_sound
    (rho : ∀ (a b : Obj), Gen a b → V a →₁[K] V b)
    (Rules : ∀ {a b : Obj}, Diagram Gen a b → Diagram Gen a b → Prop)
    (valid : ∀ {a b} {d e : Diagram Gen a b}, Rules d e → denote rho d = denote rho e)
    {x y : Obj} {d e : Diagram Gen x y} (h : RewriteStep Rules d e) :
    denote rho d = denote rho e := by
  cases h with
  | in_context h C => exact context_sound rho (valid h) C

/-- Every finite derivation using sound contextual rules preserves the exact linear map. -/
theorem derivation_sound
    (rho : ∀ (a b : Obj), Gen a b → V a →₁[K] V b)
    (Rules : ∀ {a b : Obj}, Diagram Gen a b → Diagram Gen a b → Prop)
    (valid : ∀ {a b} {d e : Diagram Gen a b}, Rules d e → denote rho d = denote rho e)
    {x y : Obj} {d e : Diagram Gen x y} (h : Derives Rules d e) :
    denote rho d = denote rho e := by
  induction h with
  | refl => rfl
  | tail first last ih => exact ih.trans (rewriteStep_sound rho Rules valid last)

/-- An exact consequence for the fixed identity/associativity rewrite system. -/
theorem structural_derivation_sound
    (rho : ∀ (a b : Obj), Gen a b → V a →₁[K] V b)
    {x y : Obj} {d e : Diagram Gen x y} (h : Derives StructuralRule d e) :
    denote rho d = denote rho e :=
  derivation_sound rho StructuralRule (structural_sound rho) h

/-- Linear equivalences supply an actual inverse-cancellation rule. This does
not assert that arbitrary generators, measurements, or channels are invertible. -/
theorem linearEquiv_cancellation {a b : Obj} (e : V a ≃₁[K] V b) :
    denote (fun (i j : Obj) (g : V i ≃₁[K] V j) => g.toLinearMap)
      (.seq (.atom e) (.atom e.symm)) =
    denote (fun (i j : Obj) (g : V i ≃₁[K] V j) => g.toLinearMap) (.ident a) := by
  ext x
```

```
  exact e.symm_apply_apply x

/-- A concrete basic rule family for generators interpreted as linear
 equivalences: structural rules and inverse-pair cancellation. -/
inductive EquivRule : ∀ {a b : Obj},
    Diagram (fun i j => V i ≃ₗ[K] V j) a b →
    Diagram (fun i j => V i ≃ₗ[K] V j) a b → Prop
  | structural {a b} {d e : Diagram (fun i j => V i ≃ₗ[K] V j) a b}
      (h : StructuralRule d e) : EquivRule d e
  | cancel {a b} (e : V a ≃ₗ[K] V b) :
      EquivRule (.seq (.atom e) (.atom e.symm)) (.ident a)

/-- Each rule of the explicit equivalence-generator system has an actual proof. -/
theorem equivRule_sound {a b : Obj}
    {d e : Diagram (fun i j => V i ≃ₗ[K] V j) a b} (h : EquivRule d e) :
    denote (fun (i j : Obj) (g : V i ≃ₗ[K] V j) => g.toLinearMap) d =
    denote (fun (i j : Obj) (g : V i ≃ₗ[K] V j) => g.toLinearMap) e := by
  cases h with
  | structural hs => exact structural_sound _ hs
  | cancel e => exact linearEquiv_cancellation e

/-- Any finite sequence of contextual structural/inverse rewrites preserves
 the exact linear map, without assuming every pair of equivalent maps can be
 connected by this limited rule set. -/
theorem equiv_derivation_sound {a b : Obj}
    {d e : Diagram (fun i j => V i ≃ₗ[K] V j) a b} (h : Derives EquivRule d e) :
    denote (fun (i j : Obj) (g : V i ≃ₗ[K] V j) => g.toLinearMap) d =
    denote (fun (i j : Obj) (g : V i ≃ₗ[K] V j) => g.toLinearMap) e :=
  derivation_sound _ EquivRule equivRule_sound h

end Semantics

section MatrixComponents

open scoped Kronecker

variable {R : Type u} [CommSemiring R]

/-- Tensor interchange is a finite matrix identity with explicit product indices. -/
theorem tensor_interchange {a b c d e f : Type v} [Fintype b] [Fintype e]
    (A : Matrix a b R) (B : Matrix b c R) (C : Matrix d e R) (D : Matrix e f R) :
    (A ⊗ₖ C) * (B ⊗ₖ D) = (A * B) ⊗ₖ (C * D) :=
  (Matrix.mul_kronecker_mul A B C D).symm

/-- Two disjoint closed matrix contractions multiply their scalar values. -/
theorem tensor_trace_factorization {a b : Type v} [Fintype a] [Fintype b]
    (A : Matrix a a R) (B : Matrix b b R) :
    Matrix.trace (A ⊗ₖ B) = Matrix.trace A * Matrix.trace B :=
  Matrix.trace_kronecker A B

/-- Cyclicity also permits the two sides to have different finite dimensions. -/
theorem rectangular_trace_cyclicity {a b : Type v} [Fintype a] [Fintype b]
    (A : Matrix a b R) (B : Matrix b a R) :
    Matrix.trace (A * B) = Matrix.trace (B * A) :=
  Matrix.trace_mul_comm A B

/-- A closed identity wire contributes the dimension, not the scalar one. -/
theorem identity_loop_value {a : Type v} [Fintype a] [DecidableEq a] :
    Matrix.trace (1 : Matrix a a R) = (Fintype.card a : R) :=
  Matrix.trace_one

end MatrixComponents

section ChangeOfBasis

variable {K V0 V1 V2 W0 W1 W2 : Type*} [Semiring K]
variable [AddCommMonoid V0] [AddCommMonoid V1] [AddCommMonoid V2]
variable [AddCommMonoid W0] [AddCommMonoid W1] [AddCommMonoid W2]
```

```
variable [Module K V0] [Module K V1] [Module K V2]
variable [Module K W0] [Module K W1] [Module K W2]

/-- Independent changes of bases preserve the zero-composition equation.
The equivalences are algebraic; no norm or spectral bound is asserted. -/
theorem complex_transport
    (A : V0 →ₗ[K] V1) (B : V1 →ₗ[K] V2) (h : B.comp A = 0)
    (e0 : V0 ≃ₗ[K] W0) (e1 : V1 ≃ₗ[K] W1) (e2 : V2 ≃ₗ[K] W2) :
    (e2.toLinearMap.comp (B.comp e1.symm.toLinearMap)).comp
      (e1.toLinearMap.comp (A.comp e0.symm.toLinearMap)) = 0 := by
  ext x
  have hz := LinearMap.congr_fun h (e0.symm x)
  simpa using congrArg e2 hz

/-- A middle contracting-homotopy identity proves exactness, provided the
successive maps also form a complex. This abstracts the exactness step in the
Koszul and finite Hodge arguments without formalizing exterior algebra. -/
theorem homotopy_implies_exact
    (A : V0 →ₗ[K] V1) (B : V1 →ₗ[K] V2) (hcomplex : B.comp A = 0)
    (s : V1 →ₗ[K] V0) (t : V2 →ₗ[K] V1)
    (hhomotopy : A.comp s + t.comp B = LinearMap.id) :
    LinearMap.range A = LinearMap.ker B := by
  apply le_antisymm (LinearMap.range_le_ker_iff.mpr hcomplex)
  intro v hv
  refine ⟨s v, ?_⟩
  have h := LinearMap.congr_fun hhomotopy v
  have hv0 : B v = 0 := hv
  simpa [hv0] using h

/-- Exactness itself is preserved under the independent algebraic changes of
bases; this includes an explicit witness for every vector in the new kernel. -/
theorem exact_transport
    (A : V0 →ₗ[K] V1) (B : V1 →ₗ[K] V2)
    (hexact : LinearMap.range A = LinearMap.ker B)
    (e0 : V0 ≃ₗ[K] W0) (e1 : V1 ≃ₗ[K] W1) (e2 : V2 ≃ₗ[K] W2) :
    LinearMap.range (e1.toLinearMap.comp (A.comp e0.symm.toLinearMap)) =
      LinearMap.ker (e2.toLinearMap.comp (B.comp e1.symm.toLinearMap)) := by
  have hcomplex : B.comp A = 0 := LinearMap.range_le_ker_iff.mp hexact.le
  apply le_antisymm
    (LinearMap.range_le_ker_iff.mpr (complex_transport A B hcomplex e0 e1 e2))
  intro w hw
  change e2 (B (e1.symm w)) = 0 at hw
  have hb : B (e1.symm w) = 0 := e2.injective (by simpa using hw)
  have hr : e1.symm w ∈ LinearMap.range A := by
    rw [hexact]
    exact hb
  rcases hr with ⟨v, hv⟩
  refine ⟨e0 v, ?_⟩
  simpa using congrArg e1 hv

end ChangeOfBasis

section HodgeConjugacy

open Matrix

/-- A finite matrix Hodge identity under unitary changes at the endpoint
spaces. The middle matrix may be arbitrary for this identity; choosing it
unitary yields the usual unitary-conjugacy formula. No spectral theorem or
infinite-dimensional operator domain is formalized by this statement. -/
theorem matrix_hodge_conjugacy
    {n0 n1 n2 : Type*} [Fintype n0] [Fintype n1] [Fintype n2]
    [DecidableEq n0] [DecidableEq n2]
    (A : Matrix n1 n0 ℂ) (B : Matrix n2 n1 ℂ)
    (U0 : Matrix n0 n0 ℂ) (U1 : Matrix n1 n1 ℂ) (U2 : Matrix n2 n2 ℂ)
    (h0 : U0ᴴ * U0 = 1) (h2 : U2ᴴ * U2 = 1) :
    (U1 * A * U0ᴴ) * (U1 * A * U0ᴴ)ᴴ +
      (U2 * B * U1ᴴ)ᴴ * (U2 * B * U1ᴴ) =
```

```
      U1 * (A * A^H + B^H * B) * U1^H := by
  simp only [Matrix.conjTranspose_mul, Matrix.conjTranspose_conjTranspose]
  calc
    _ = U1 * A * (U0^H * U0) * A^H * U1^H +
        U1 * B^H * (U2^H * U2) * B * U1^H := by simp only [Matrix.mul_assoc]
    _ = U1 * A * A^H * U1^H + U1 * B^H * B * U1^H := by
      rw [h0, h2]
      simp only [Matrix.mul_one]
    _ = _ := by simp only [Matrix.mul_add, Matrix.add_mul, Matrix.mul_assoc]

end HodgeConjugacy

end DiagramGrammar

-- Exact declaration types, followed by foundational dependency inventories.
#check DiagramGrammar.context_sound
#check DiagramGrammar.structural_sound
#check DiagramGrammar.rewriteStep_sound
#check DiagramGrammar.derivation_sound
#check DiagramGrammar.structural_derivation_sound
#check DiagramGrammar.linearEquiv_cancellation
#check DiagramGrammar.equivRule_sound
#check DiagramGrammar.equiv_derivation_sound
#check DiagramGrammar.tensor_interchange
#check DiagramGrammar.tensor_trace_factorization
#check DiagramGrammar.rectangular_trace_cyclicity
#check DiagramGrammar.identity_loop_value
#check DiagramGrammar.complex_transport
#check DiagramGrammar.homotopy_implies_exact
#check DiagramGrammar.exact_transport
#check DiagramGrammar.matrix_hodge_conjugacy

#print axioms DiagramGrammar.context_sound
#print axioms DiagramGrammar.structural_sound
#print axioms DiagramGrammar.rewriteStep_sound
#print axioms DiagramGrammar.derivation_sound
#print axioms DiagramGrammar.structural_derivation_sound
#print axioms DiagramGrammar.linearEquiv_cancellation
#print axioms DiagramGrammar.equivRule_sound
#print axioms DiagramGrammar.equiv_derivation_sound
#print axioms DiagramGrammar.tensor_interchange
#print axioms DiagramGrammar.tensor_trace_factorization
#print axioms DiagramGrammar.rectangular_trace_cyclicity
#print axioms DiagramGrammar.identity_loop_value
#print axioms DiagramGrammar.complex_transport
#print axioms DiagramGrammar.homotopy_implies_exact
#print axioms DiagramGrammar.exact_transport
#print axioms DiagramGrammar.matrix_hodge_conjugacy
```

Appendix C

Fluid verification, reproducibility and proof dependencies

C.1 Current verification boundary

The final local companion has 121 selected declarations across seven modules, including 16 reduction contracts, 17 transport declarations and six weighted-gain contracts. All seven were freshly checked with Lean 4.24.0 and the pinned mathlib revision, with zero warnings and only the disclosed standard foundational axioms. Exact types, source hashes and logs are retained under `lean/final-verification`. The earlier 82-declaration fluid-era records below remain identifiable historical checks of the first four modules. Neither record completes the separate upstream Navier–Stokes full build, which remains blocked by missing rate-limited dependency artifacts.

C.2 What was checked and what each check means

The fluid extension has three distinct verification layers. The manuscript gives explicit proofs of the displayed finite statements and analytical interfaces. The local Lean package checks selected statements with precisely typed hypotheses. The executable examples compare independently implemented finite calculations. None of these layers is interchangeable with the complete upstream continuum development.

The additional file `extension/lean/FluidInterfaces.lean` contains 34 theorem declarations. It was elaborated successfully by Lean 4.24.0 against the same pinned mathlib revision as the earlier companion,

`f897ebcf72cd16f89ab4577d0c826cd14afaafc7`.

The check returned exit code zero and no warnings. Every declaration has a printed type and an axiom report. The only reported foundations are `propext`, `Classical.choice` and `Quot.sound`; no admitted theorem or project-specific axiom is introduced. The verified source SHA-256 is

`9ac752e3450fd07f9fb23cb4d0283ecee1b70d7a3188ece1cdcff5d8b66934b2`.

At that earlier checkpoint, all four modules were freshly rechecked together, with unchanged source hashes and zero warnings. The individual logs and declaration-by-declaration axiom comparisons are retained in `extension/review/navier-stokes/local-kernel-verification.json`. That earlier package contained 82 selected declarations: 25 original interfaces, seven quantum interfaces, sixteen diagram interfaces and these 34 fluid interfaces. This is a count of exact declarations, not a claim that

82 complete manuscript theorems, every numerical implementation, or the continuum singularity construction have been formalized here.

Family	Count	Checked content and its boundary
Finite energy	11	Finite real matrix work identities, pressure cancellation under the stated constraint, nonnegative factorized viscosity, and an actual time derivative of coordinate kinetic energy. No continuum regularity conclusion.
Orthogonal reduction	5	Kernel residual, orthogonality, Pythagorean energy split, energy nonincrease and best approximation for a symmetric idempotent matrix. No closure or trajectory error theorem.
Chain continuity	2	Boundary flux has zero subsequent boundary, and adding such a flux preserves the divergence. The complex equation is explicit.
Schur elimination	1	Equivalence of a block system and its eliminated system when the required block is a linear equivalence. No assumption that every block is invertible.
Discrete memory	3	All-step hidden-state expansion, resolved memory identity and the obstruction to an instantaneous closure. No decay assumption is inserted into the recurrence.
Forbidden composition	1	A strict lower block, then a block diagonal map, then another strict lower block has zero composite. This is one exact operator pattern from the Volterra discussion.
Quadratic order	11	The faithful order representation, addition and multiplication, norm determinant, matrix polynomial, determinant of every power, and absence of a commuting real linear complex structure. The topological covering interpretation is proved separately in the manuscript.

The component map in `extension/review/navier-stokes/formalization-map.json` and the main claim register link these declarations to the statements they actually support. A matrix energy identity is a component of the finite ODE existence argument, not a Lean proof of that entire existence theorem. The exact discrete memory recurrence is not a Lean proof of the continuous integral formula. Likewise the determinant of J^k is not, by itself, a formalized covering-space degree theorem.

C.3 The separate upstream Navier–Stokes formalization

The released project uses a different toolchain: Lean 4.34.0-rc2 and its own pinned mathlib dependency. Its repository is pinned at

`f9e8bc5b38b6e212696e8a30e3e91517af887bbd`.

The audited final solution target has a local import closure of 609 modules. A comment- and string-aware source scan of this closure finds no admitted placeholders or project axiom declarations, and no challenge module is imported by that closure. This is useful source evidence but is not a kernel check.

A fresh upstream build was attempted with the exact verified compiler archive. Dependency-cache retrieval encountered HTTP 429 responses. The subsequent direct target check stopped because imported build artifacts were missing, before theorem-body elaboration. The full upstream kernel check and independent Comparator check therefore remain uncompleted in this work. The supplied `release-build-verification.json` records these outcomes, while `release-verify.py` supplies a pinned reproduction route. A failed dependency download neither verifies nor refutes the theorem. The local 34-declaration result is kept entirely separate from this upstream result.

The complete 166-page paper was read across the two mathematical reviewers, with every page attributed in `release-source-manifest.json`. That reading supports the exposition and the precise statement of the forced whole-space and periodic results. It does not substitute for the missing fresh formal build, nor does it turn a proof of forced breakdown into an unforced theorem.

C.4 Independent finite computational checks

The new program `extension/code/verify_fluid_interfaces.py` constructs exact integer cell complexes for six periodic square grids, with $L = 3, \ldots, 8$. It checks the zero boundary composite and the expected finite ranks. It independently evaluates symbolic Taylor–Green equations and dense complex-matrix circuit operations. Its machine-readable output is `extension/verification/fluids-python.json`.

The browser core is separately checked by `extension/site/verify-fluids.cjs`, which consumes independent Python reference data. The recorded comparison covers sixteen circuit inputs, 64 parameter-feature derivative cases and 8,865 complete toral-fibre points. The maximum amplitude discrepancy is about 2.22×10^{-16}; the maximum character-pullback discrepancy is about 1.01×10^{-13}; and the heat-channel probability discrepancy is about 3.47×10^{-18}. These are floating-point discrepancies, not exact-real proof bounds. The finite-difference momentum residual decreases from 3.03×10^{-6} to 7.57×10^{-7} to 1.89×10^{-7} under successive halvings of the step, consistent with the method's second-order error on this manufactured solution.

The channel checks distinguish the live diagonal-state simulation from the separate coherent-state matrix test. Positivity, trace preservation and complete positivity follow analytically from the Kraus form in Chapter 27; finite Choi and density-matrix calculations test the implemented instances. The circuit is simulated classically and is untrained. The circuit register and the periodic position register are distinct systems. No NVIDIA acceleration is required or claimed for these small new examples; the earlier GPU surface-code and ToricGT experiments retain their separate recorded execution evidence.

C.5 Reproduction commands

The weighted-gain continuation adds 15 test methods, including 256 seeded exact comparisons, 84 invalid-scalar cases, 16 evidence mutations and 21 isolated CLI cases. Four retained witness/replay pairs cover a contracting bound, a rejected rate, a holding noncontracting rate and the exact rational reconstruction of a historical floating-point false positive. The separate quantum feedback example checks 180 channel parameter cases, 2,700 Choi principal minors, 2,160 coherent matrix/Bloch comparisons, 2,160 fixed-control distance equalities and 256 exact controller block bounds. It additionally retains ideal and perturbed trajectories, eight path horizons, invalid-input checks and independent accepted/rejected gain replays. These are finite implementation checks; the global comparison and physical channel arguments are the written proofs in Chapter 30.

The exact records and rational artifacts are stored in these directories:

```
extension/review/final-pass/weighted-gain/
extension/review/final-pass/quantum-gain-control/
```

From the repository root, the new checks can be run without a GPU or quantum device:

```
python3 -B extension/code/test_weighted_gain.py
python3 extension/code/check_weighted_gain_witness.py \
  extension/review/final-pass/weighted-gain/witness-accepted.json
python3 extension/code/verify_quantum_gain_control.py \
  --output-dir /tmp/grh-quantum-gain-evidence
```

The independent witness checker works under optimized Python. The broader quantum example imports an existing verifier that deliberately refuses optimized execution; run that driver normally, as shown. Neither program claims a security proof, a trained fluid model or a native quantum run.

From the repository root, the finite computations can be repeated with:

```
python3 extension/code/verify_fluid_interfaces.py
node extension/site/verify-fluids.cjs
```

The first script requires the numerical and symbolic dependencies stated in its source. The second uses the generated reference data. The local formal check can be repeated using the installed pinned toolchain:

```
python3 extension/review/navier-stokes/formalization-verify.py \
  --toolchain /absolute/path/to/lean-4.24.0-linux \
  --mathlib /absolute/path/to/mathlib4-v4.24.0
```

Alternatively, from a correctly resolved Lake package, run

```
lake env lean FluidInterfaces.lean
```

The manuscript entry points share the same chapter sources. With the listed TeX packages installed:

```
python3 extension/manuscript/build_pdf.py
python3 extension/manuscript/build_pdf.py \
  --main fluid-extension --min-pages 1
```

The full monograph keeps the 100-page lower gate. The standalone supplement is not padded to a target length. Both builds reject undefined references, missing characters and overfull boxes, and retain their actual compiler logs and source hashes. Visual review is recorded separately from those mechanical gates.

C.6 The complete selected fluid formalization

The following is the source actually checked, including its type and axiom audit block. Its finite types are part of the theorem statements; replacing them with function spaces would create a different theorem and new analytic obligations.

```
import Mathlib.LinearAlgebra.Matrix.PosDef
import Mathlib.LinearAlgebra.Matrix.ToLin
import Mathlib.LinearAlgebra.Matrix.Notation
import Mathlib.LinearAlgebra.Matrix.Determinant.Basic
import Mathlib.Algebra.Module.LinearMap.End
import Mathlib.Analysis.Calculus.Deriv.Add
import Mathlib.Analysis.Calculus.Deriv.Mul
import Mathlib.Tactic.Ring
import Mathlib.Tactic.Linarith

/-!
# Finite fluid and information-reduction interfaces
```

```
The matrix results below concern actual finite real matrices, not a continuum
existence or regularity theorem. In particular, an energy identity for every
finite discretization supplies no mesh-uniform derivative or compactness bound.
Skew transport, positive factorized closures, orthogonal projection, and chain
complexes provide independently checkable structural constraints for a fluid
model on a surface cellulation. A learned closure may depend on the current
state: the corresponding matrix identities apply pointwise at that state.

The block-elimination and memory results describe exact finite linear reduction.
They do not identify its memory kernel with the nonlinear Navier--Stokes one,
and do not assert that an instantaneous closure is exact.

This file adds selected mathematical interfaces to the existing formalization.
It does not reprove or import the OpenAI Navier--Stokes release.
-/

set_option autoImplicit false

open Matrix
open scoped BigOperators Matrix

namespace FluidInterfaces

section Energy

variable {n m p : Type*} [Fintype n] [Fintype m] [Fintype p]

/-- Finite Euclidean kinetic energy with unit mass matrix. -/
noncomputable def kineticEnergy (u : n → ℝ) : ℝ := (u ⬝ᵥ u) / 2

theorem dot_self_nonneg (u : n → ℝ) : 0 ≤ u ⬝ᵥ u := by
  exact Finset.sum_nonneg (fun i _ => mul_self_nonneg (u i))

theorem kineticEnergy_nonneg (u : n → ℝ) : 0 ≤ kineticEnergy u :=
  div_nonneg (dot_self_nonneg u) (by norm_num)

/-- The transpose identity is valid also for rectangular finite matrices. -/
theorem transpose_work (L : Matrix m n ℝ) (u : n → ℝ) (v : m → ℝ) :
    u ⬝ᵥ (Lᵀ *ᵥ v) = (L *ᵥ u) ⬝ᵥ v := by
  rw [dotProduct_mulVec, vecMul_transpose]

/-- Actual matrix skew-symmetry implies exactly zero quadratic work. -/
theorem skew_transport_work (J : Matrix n n ℝ) (hJ : Jᵀ = -J)
    (u : n → ℝ) : u ⬝ᵥ (J *ᵥ u) = 0 := by
  have h := transpose_work J u u
  rw [hJ, neg_mulVec, dotProduct_neg, dotProduct_comm (J *ᵥ u) u] at h
  linarith

/-- A factorized dissipative closure has exactly squared-flux work. -/
theorem factorized_viscosity_work (L : Matrix m n ℝ) (u : n → ℝ) :
    u ⬝ᵥ ((Lᵀ * L) *ᵥ u) = (L *ᵥ u) ⬝ᵥ (L *ᵥ u) := by
  rw [← mulVec_mulVec, transpose_work]

theorem factorized_viscosity_nonneg (L : Matrix m n ℝ) (u : n → ℝ) :
    0 ≤ u ⬝ᵥ ((Lᵀ * L) *ᵥ u) := by
  rw [factorized_viscosity_work]
  exact dot_self_nonneg _

/-- Discrete gradient pressure does no work on a divergence-free velocity. -/
theorem pressure_work_zero (B : Matrix p n ℝ) (u : n → ℝ) (pressure : p → ℝ)
    (hu : B *ᵥ u = 0) : u ⬝ᵥ (Bᵀ *ᵥ pressure) = 0 := by
  rw [transpose_work, hu, zero_dotProduct]

/-- Exact finite energy-work balance including pressure, physical viscosity,
and a separately factorized learned closure. No continuum limit is asserted. -/
theorem finite_fluid_work_identity
    (J : Matrix n n ℝ) (hJ : Jᵀ = -J)
```

```
    (D : Matrix m n ℝ) (L : Matrix p n ℝ) (ν : ℝ)
    (u force : n → ℝ) :
    u ⬝ᵥ (J *ᵥ u - ν • ((Dᵀ * D) *ᵥ u) - (Lᵀ * L) *ᵥ u + force) =
      u ⬝ᵥ force - ν * ((D *ᵥ u) ⬝ᵥ (D *ᵥ u)) -
        (L *ᵥ u) ⬝ᵥ (L *ᵥ u) := by
  rw [dotProduct_add, dotProduct_sub, dotProduct_sub,
    skew_transport_work J hJ, dotProduct_smul,
    factorized_viscosity_work, factorized_viscosity_work]
  simp only [smul_eq_mul]
  ring

theorem unforced_work_nonpositive
    (J : Matrix n n ℝ) (hJ : Jᵀ = -J)
    (D : Matrix m n ℝ) (L : Matrix p n ℝ) (ν : ℝ) (hν : 0 ≤ ν)
    (u : n → ℝ) :
    u ⬝ᵥ (J *ᵥ u - ν • ((Dᵀ * D) *ᵥ u) - (Lᵀ * L) *ᵥ u) ≤ 0 := by
  have h := finite_fluid_work_identity J hJ D L ν u 0
  simp only [add_zero, dotProduct_zero] at h
  rw [h]
  have hv := mul_nonneg hν (dot_self_nonneg (D *ᵥ u))
  have hl := dot_self_nonneg (L *ᵥ u)
  linarith

/-- This is an actual time derivative, proved from the coordinate derivatives. -/
theorem kinetic_energy_derivative (u : ℝ → n → ℝ) (v : n → ℝ) (t : ℝ)
    (hu : ∀ i, HasDerivAt (fun s => u s i) (v i) t) :
    HasDerivAt (fun s => kineticEnergy (u s)) ((u t) ⬝ᵥ v) t := by
  have hi : ∀ i, HasDerivAt (fun s => u s i * u s i / 2) (u t i * v i) t := by
    intro i
    convert ((hu i).mul (hu i)).div_const 2 using 1; ring
  have hs := HasDerivAt.fun_sum (u := Finset.univ) (fun i _ => hi i)
  simpa only [kineticEnergy, dotProduct, Finset.sum_div] using hs

/-- The exact semidiscrete ODE yields a genuine differential energy identity.
It remains a finite model: no uniform continuum estimate is implicit. -/
theorem finite_fluid_energy_derivative
    (J : Matrix n n ℝ) (hJ : Jᵀ = -J)
    (D : Matrix m n ℝ) (L : Matrix p n ℝ) (ν : ℝ)
    (u : ℝ → n → ℝ) (force : n → ℝ) (t : ℝ)
    (hu : ∀ i, HasDerivAt (fun s => u s i)
      ((J *ᵥ u t - ν • ((Dᵀ * D) *ᵥ u t) - (Lᵀ * L) *ᵥ u t + force) i) t) :
    HasDerivAt (fun s => kineticEnergy (u s))
      (u t ⬝ᵥ force - ν * ((D *ᵥ u t) ⬝ᵥ (D *ᵥ u t)) -
        (L *ᵥ u t) ⬝ᵥ (L *ᵥ u t)) t := by
  have h := kinetic_energy_derivative u _ t hu
  rwa [finite_fluid_work_identity J hJ D L ν (u t) force] at h

end Energy

section OrthogonalReduction

variable {n : Type*} [Fintype n]

/-- Idempotence sends the discarded part into the projector's kernel. -/
theorem projection_residual_kernel (P : Matrix n n ℝ) (hP : P * P = P)
    (u : n → ℝ) : P *ᵥ (u - P *ᵥ u) = 0 := by
  rw [mulVec_sub, mulVec_mulVec, hP, sub_self]

/-- Symmetry and idempotence imply orthogonality of retained and discarded data. -/
theorem projection_residual_orthogonal (P : Matrix n n ℝ)
    (hs : Pᵀ = P) (hi : P * P = P) (u v : n → ℝ) :
    (u - P *ᵥ u) ⬝ᵥ (P *ᵥ v) = 0 := by
  rw [← hs, transpose_work, hs, projection_residual_kernel P hi, zero_dotProduct]

/-- Exact information loss of an orthogonal linear reduction in kinetic energy. -/
theorem projection_energy_split (P : Matrix n n ℝ)
    (hs : Pᵀ = P) (hi : P * P = P) (u : n → ℝ) :
    kineticEnergy u = kineticEnergy (P *ᵥ u) + kineticEnergy (u - P *ᵥ u) := by
```

```
  have ho := projection_residual_orthogonal P hs hi u u
  simp only [sub_dotProduct] at ho
  have hc := dotProduct_comm u (P *ᵥ u)
  simp only [kineticEnergy, sub_dotProduct, dotProduct_sub]
  linarith

theorem projection_energy_nonincrease (P : Matrix n n ℝ)
    (hs : Pᵀ = P) (hi : P * P = P) (u : n → ℝ) :
    kineticEnergy (P *ᵥ u) ≤ kineticEnergy u := by
  have hsplit := projection_energy_split P hs hi u
  have hnonneg := kineticEnergy_nonneg (u - P *ᵥ u)
  linarith

/-- The same split proves optimality against every vector in the retained range. -/
theorem projection_best_approximation (P : Matrix n n ℝ)
    (hs : Pᵀ = P) (hi : P * P = P) (u v : n → ℝ) (hv : P *ᵥ v = v) :
    kineticEnergy (u - P *ᵥ u) ≤ kineticEnergy (u - v) := by
  have hsplit := projection_energy_split P hs hi (u - v)
  have hp : P *ᵥ (u - v) = P *ᵥ u - v := by rw [mulVec_sub, hv]
  have hr : u - v - (P *ᵥ u - v) = u - P *ᵥ u := by abel
  rw [hp, hr] at hsplit
  rw [hsplit]
  exact le_add_of_nonneg_left (kineticEnergy_nonneg _)

end OrthogonalReduction

section Continuity

variable {n m p : Type*} [Fintype n] [Fintype p]

/-- A flux represented by the next boundary automatically preserves divergence. -/
theorem boundary_flux_divergence_zero (B : Matrix m n ℝ) (C : Matrix n p ℝ)
    (hcomplex : B * C = 0) (q : p → ℝ) : B *ᵥ (C *ᵥ q) = 0 := by
  rw [mulVec_mulVec, hcomplex, zero_mulVec]

/-- The exact algebraic continuity constraint survives every such flux update. -/
theorem boundary_flux_update_preserves_divergence
    (B : Matrix m n ℝ) (C : Matrix n p ℝ) (hcomplex : B * C = 0)
    (u : n → ℝ) (q : p → ℝ) :
    B *ᵥ (u + C *ᵥ q) = B *ᵥ u := by
  rw [mulVec_add, boundary_flux_divergence_zero B C hcomplex, add_zero]

end Continuity

section BlockElimination

variable {K X Y : Type*} [Field K]
variable [AddCommGroup X] [Module K X] [AddCommGroup Y] [Module K Y]

/-- The exact reduced block; invertibility is an explicit analytic obstruction. -/
def schurComplement (A : X →ₗ[K] X) (B : Y →ₗ[K] X)
    (C : X →ₗ[K] Y) (D : Y ≃ₗ[K] Y) : X →ₗ[K] X :=
  A - B.comp (D.symm.toLinearMap.comp C)

/-- Eliminating the hidden variable is equivalent to the original block system,
not merely a one-way necessary condition. -/
theorem schur_elimination_iff
    (A : X →ₗ[K] X) (B : Y →ₗ[K] X)
    (C : X →ₗ[K] Y) (D : Y ≃ₗ[K] Y)
    (x : X) (y : Y) (f : X) (g : Y) :
    (A x + B y = f ∧ C x + D y = g) ↔
      (schurComplement A B C D x = f - B (D.symm g) ∧
        y = D.symm (g - C x)) := by
  have hy : C x + D y = g ↔ y = D.symm (g - C x) := by
    constructor
    · intro h
      have hd : D y = g - C x := by rw [← h]; abel
      simpa only [D.symm_apply_apply] using congrArg D.symm hd
```

```
    · intro h
      rw [h, D.apply_symm_apply]
      abel
  rw [hy]
  constructor
  · rintro ⟨hf, hy'⟩
    refine ⟨?_, hy'⟩
    rw [hy', map_sub, map_sub] at hf
    simp only [schurComplement, LinearMap.sub_apply, LinearMap.comp_apply,
      LinearEquiv.coe_coe]
    rw [← hf]
    abel
  · rintro ⟨hf, hy'⟩
    refine ⟨?_, hy'⟩
    simp only [schurComplement, LinearMap.sub_apply, LinearMap.comp_apply,
      LinearEquiv.coe_coe] at hf
    rw [hy', map_sub, map_sub]
    rw [sub_eq_iff_eq_add] at hf
    rw [hf]
    abel

end BlockElimination

section ExactMemory

variable {K X Y : Type*} [Field K]
variable [AddCommGroup X] [Module K X] [AddCommGroup Y] [Module K Y]

/-- Eliminating a hidden linear recurrence produces its entire past dependence. -/
theorem hidden_recurrence_expansion (C : X →ₗ[K] Y) (D : Module.End K Y)
    (p : ℕ → X) (q : ℕ → Y)
    (hstep : ∀ n : ℕ, q (n + 1) = D (q n) + C (p n)) (n : ℕ) :
    q n = (D ^ n) (q 0) +
      ∑ k ∈ Finset.range n, (D ^ (n - 1 - k)) (C (p k)) := by
  induction n with
  | zero => simp
  | succ n ih =>
    rw [hstep, ih, map_add, map_sum, Finset.sum_range_succ]
    simp only [Nat.add_sub_cancel, Nat.sub_self, pow_zero, Module.End.one_apply]
    have hsum :
        (∑ k ∈ Finset.range n, D ((D ^ (n - 1 - k)) (C (p k)))) =
          ∑ k ∈ Finset.range n, (D ^ (n - k)) (C (p k)) := by
      apply Finset.sum_congr rfl
      intro k hk
      have hklt : k < n := Finset.mem_range.mp hk
      have he : n - k = (n - 1 - k) + 1 := by omega
      rw [he, pow_succ', Module.End.mul_apply]
    rw [hsum, pow_succ', Module.End.mul_apply]
    abel

/-- The exact reduced dynamics contains a forcing from the hidden initial state
and the finite convolution with kernel B D^j C. -/
theorem resolved_memory_identity (A : Module.End K X) (B : Y →ₗ[K] X)
    (C : X →ₗ[K] Y) (D : Module.End K Y) (p : ℕ → X) (q : ℕ → Y)
    (hp : ∀ n : ℕ, p (n + 1) = A (p n) + B (q n))
    (hq : ∀ n : ℕ, q (n + 1) = D (q n) + C (p n)) (n : ℕ) :
    p (n + 1) = A (p n) + B ((D ^ n) (q 0)) +
      ∑ k ∈ Finset.range n, B ((D ^ (n - 1 - k)) (C (p k))) := by
  rw [hp, hidden_recurrence_expansion C D p q hq n, map_add, map_sum]
  abel

/-- There need be no exact autonomous deterministic closure on retained data:
two states with the same retained coordinates can have distinct next states. -/
theorem hidden_state_prevents_autonomous_closure
    (A : Module.End K X) (B : Y →ₗ[K] X) (y : Y) (hy : B y ≠ 0) :
    ¬ ∃ F : X → X, ∀ (x : X) (z : Y), F x = A x + B z := by
  rintro ⟨F, hF⟩
  have h₀ := hF 0 0
```

```
  have hy' := hF 0 y
  simp only [map_zero, add_zero, zero_add] at h₀ hy'
  exact hy (hy'.symm.trans h₀)

end ExactMemory

section NilpotentBlocks

variable {K X Y : Type*} [Field K]
variable [AddCommGroup X] [Module K X] [AddCommGroup Y] [Module K Y]

/-- A strict lower rectangular block on the direct sum of retained spaces. -/
def lowerBlock (L : X →ₗ[K] Y) : Module.End K (X × Y) where
  toFun xy := (0, L xy.1)
  map_add' a b := by simp
  map_smul' c a := by simp

/-- Any diagonal block between two strict lower blocks leaves zero composite.
The left and right lower blocks may differ, covering a parameter derivative. -/
theorem lower_diagonal_lower_zero (L M : X →ₗ[K] Y)
    (DX : Module.End K X) (DY : Module.End K Y) :
    (lowerBlock L).comp ((DX.prodMap DY).comp (lowerBlock M)) = 0 := by
  apply LinearMap.ext
  intro xy
  change (0, L (DX 0)) = (0, 0)
  simp

end NilpotentBlocks

section QuadraticOrder

/-- The integer matrix for multiplication by sqrt(2), in a conjugate basis. -/
def quadraticGenerator : Matrix (Fin 2) (Fin 2) ℤ := !![-1, 1; 1, 1]

/-- The degree-fourteen matrix appearing in the reviewed proof's toral stage. -/
def renormalizationMatrix : Matrix (Fin 2) (Fin 2) ℤ := !![3, 1; 1, 5]

/-- An exact faithful representation of the underlying quadratic order. -/
def quadraticOrderMatrix (a b : ℤ) : Matrix (Fin 2) (Fin 2) ℤ :=
  !![a - b, b; b, a + b]

theorem renormalization_is_quadratic_order_element :
    renormalizationMatrix = quadraticOrderMatrix 4 1 := by
  decide

theorem quadratic_order_one :
    quadraticOrderMatrix 1 0 = (1 : Matrix (Fin 2) (Fin 2) ℤ) := by
  decide

theorem quadratic_order_addition (a b c d : ℤ) :
    quadraticOrderMatrix a b + quadraticOrderMatrix c d =
      quadraticOrderMatrix (a + c) (b + d) := by
  ext i j
  fin_cases i <;> fin_cases j <;> simp [quadraticOrderMatrix] <;> ring

theorem quadratic_generator_square : quadraticGenerator * quadraticGenerator =
    (2 : ℤ) • (1 : Matrix (Fin 2) (Fin 2) ℤ) := by
  decide

theorem renormalization_quadratic_identity :
    renormalizationMatrix * renormalizationMatrix - 8 • renormalizationMatrix +
      (14 : ℤ) • (1 : Matrix (Fin 2) (Fin 2) ℤ) = 0 := by
  decide

theorem renormalization_determinant : Matrix.det renormalizationMatrix = 14 := by
  norm_num [renormalizationMatrix, Matrix.det_fin_two]

theorem renormalization_power_determinant (k : ℕ) :
```

```
    Matrix.det (renormalizationMatrix ^ k) = 14 ^ k := by
  rw [Matrix.det_pow, renormalization_determinant]

theorem quadratic_order_multiplication (a b c d : ℤ) :
    quadraticOrderMatrix a b * quadraticOrderMatrix c d =
      quadraticOrderMatrix (a * c + 2 * b * d) (a * d + b * c) := by
  ext i j
  fin_cases i <;> fin_cases j <;>
    simp [quadraticOrderMatrix, Matrix.mul_apply, Fin.sum_univ_two] <;> ring

theorem quadratic_order_determinant (a b : ℤ) :
    Matrix.det (quadraticOrderMatrix a b) = a ^ 2 - 2 * b ^ 2 := by
  simp [quadraticOrderMatrix, Matrix.det_fin_two]
  ring

theorem quadratic_order_representation_injective :
    Function.Injective (fun ab : ℤ × ℤ => quadraticOrderMatrix ab.1 ab.2) := by
  intro ab cd h
  have h₀ := congrArg (fun M : Matrix (Fin 2) (Fin 2) ℤ => M 0 0) h
  have h₁ := congrArg (fun M : Matrix (Fin 2) (Fin 2) ℤ => M 0 1) h
  change ab.1 - ab.2 = cd.1 - cd.2 at h₀
  change ab.2 = cd.2 at h₁
  apply Prod.ext
  · linarith
  · exact h₁

/-- This real toral endomorphism cannot be complex-linear for any real linear
complex structure: no commuting real matrix squares to minus the identity. -/
theorem no_commuting_complex_structure (I : Matrix (Fin 2) (Fin 2) ℝ)
    (hc : I * !![3, 1; 1, 5] = !![3, 1; 1, 5] * I) :
    I * I ≠ -(1 : Matrix (Fin 2) (Fin 2) ℝ) := by
  intro hsq
  have hc₀ := congrArg (fun M : Matrix (Fin 2) (Fin 2) ℝ => M 0 0) hc
  have hs₀ := congrArg (fun M : Matrix (Fin 2) (Fin 2) ℝ => M 0 0) hsq
  simp [Matrix.mul_apply, Fin.sum_univ_two] at hc₀ hs₀
  have hb : I 0 1 = I 1 0 := by linarith
  rw [← hb] at hs₀
  nlinarith [sq_nonneg (I 0 0), sq_nonneg (I 0 1)]

end QuadraticOrder

end FluidInterfaces

/-! Exact theorem types and trusted-kernel foundational dependency audit. -/
#check FluidInterfaces.renormalization_is_quadratic_order_element
#check FluidInterfaces.quadratic_order_one
#check FluidInterfaces.quadratic_order_addition
#check FluidInterfaces.dot_self_nonneg
#check FluidInterfaces.kineticEnergy_nonneg
#check FluidInterfaces.transpose_work
#check FluidInterfaces.skew_transport_work
#check FluidInterfaces.factorized_viscosity_work
#check FluidInterfaces.factorized_viscosity_nonneg
#check FluidInterfaces.pressure_work_zero
#check FluidInterfaces.finite_fluid_work_identity
#check FluidInterfaces.unforced_work_nonpositive
#check FluidInterfaces.kinetic_energy_derivative
#check FluidInterfaces.finite_fluid_energy_derivative
#check FluidInterfaces.projection_residual_kernel
#check FluidInterfaces.projection_residual_orthogonal
#check FluidInterfaces.projection_energy_split
#check FluidInterfaces.projection_energy_nonincrease
#check FluidInterfaces.projection_best_approximation
#check FluidInterfaces.boundary_flux_divergence_zero
#check FluidInterfaces.boundary_flux_update_preserves_divergence
#check FluidInterfaces.schur_elimination_iff
#check FluidInterfaces.hidden_recurrence_expansion
#check FluidInterfaces.resolved_memory_identity
```

```
#check FluidInterfaces.hidden_state_prevents_autonomous_closure
#check FluidInterfaces.lower_diagonal_lower_zero
#check FluidInterfaces.quadratic_generator_square
#check FluidInterfaces.renormalization_quadratic_identity
#check FluidInterfaces.renormalization_determinant
#check FluidInterfaces.renormalization_power_determinant
#check FluidInterfaces.quadratic_order_multiplication
#check FluidInterfaces.quadratic_order_determinant
#check FluidInterfaces.quadratic_order_representation_injective
#check FluidInterfaces.no_commuting_complex_structure

#print axioms FluidInterfaces.renormalization_is_quadratic_order_element
#print axioms FluidInterfaces.quadratic_order_one
#print axioms FluidInterfaces.quadratic_order_addition
#print axioms FluidInterfaces.dot_self_nonneg
#print axioms FluidInterfaces.kineticEnergy_nonneg
#print axioms FluidInterfaces.transpose_work
#print axioms FluidInterfaces.skew_transport_work
#print axioms FluidInterfaces.factorized_viscosity_work
#print axioms FluidInterfaces.factorized_viscosity_nonneg
#print axioms FluidInterfaces.pressure_work_zero
#print axioms FluidInterfaces.finite_fluid_work_identity
#print axioms FluidInterfaces.unforced_work_nonpositive
#print axioms FluidInterfaces.kinetic_energy_derivative
#print axioms FluidInterfaces.finite_fluid_energy_derivative
#print axioms FluidInterfaces.projection_residual_kernel
#print axioms FluidInterfaces.projection_residual_orthogonal
#print axioms FluidInterfaces.projection_energy_split
#print axioms FluidInterfaces.projection_energy_nonincrease
#print axioms FluidInterfaces.projection_best_approximation
#print axioms FluidInterfaces.boundary_flux_divergence_zero
#print axioms FluidInterfaces.boundary_flux_update_preserves_divergence
#print axioms FluidInterfaces.schur_elimination_iff
#print axioms FluidInterfaces.hidden_recurrence_expansion
#print axioms FluidInterfaces.resolved_memory_identity
#print axioms FluidInterfaces.hidden_state_prevents_autonomous_closure
#print axioms FluidInterfaces.lower_diagonal_lower_zero
#print axioms FluidInterfaces.quadratic_generator_square
#print axioms FluidInterfaces.renormalization_quadratic_identity
#print axioms FluidInterfaces.renormalization_determinant
#print axioms FluidInterfaces.renormalization_power_determinant
#print axioms FluidInterfaces.quadratic_order_multiplication
#print axioms FluidInterfaces.quadratic_order_determinant
#print axioms FluidInterfaces.quadratic_order_representation_injective
#print axioms FluidInterfaces.no_commuting_complex_structure
```

Appendix D

Source concordance, review coverage, and mathematical provenance

D.1 What the review establishes

The mathematical argument in this monograph rests on several different kinds of evidence. A bibliographic inventory identifies sources. Text extraction makes a document searchable. Reading a passage identifies its claims and hypotheses. Reconstructing a proof checks its inference. A successful Lean compilation checks a precisely stated formal declaration relative to its dependencies. These activities support different conclusions, and this chapter records their actual extent. The resulting account is substantial but is not a claim to have proved, formalized, or independently verified every statement in every linked document.

There are two relevant authorship layers. The supplied classical books on lattices over orders are by Klaus W. Roggenkamp, with Verena Huber-Dyson as coauthor of Volume I; Volume II is authored by Roggenkamp [133, 135]. The modern research repositories and linked surface-order manuscripts are associated with Amelie Schreiber [137–139]. Neither author's complete research oeuvre has been reviewed. Our use of selected additional papers by Roggenkamp concerns concrete questions about orders, irreducible lattices, and Bass-type phenomena [131, 132, 134]. It does not license a statement about everything either author has written.

D.2 The fixed source snapshot

The primary repository was examined at commit `dbd954eede479a5129831e4e7c30537389ae973d` [137]. Four important source files received complete text review: `book/main.tex`, the generated chapter `book/chapters/reviewed-mathematics.tex`, and `papers/core.tex` and `papers/details.tex`. Their measured coverage is as follows. Words are whitespace-delimited source tokens, so the counts include mathematical notation and LaTeX commands; they are coverage measurements rather than estimates of polished prose length.

Source file	Lines	Source words
Book entry point	67	466
Reviewed mathematics chapter	1453	23652
Core paper	1072	6123
Details paper	1621	9893
Total	4213	40134

Complete reading of these files does not mean that every transitive `input` was reviewed in the same audit. The book entry point includes many other chapters. The two papers also include separate short-row-pair and Möbius-splitting files. Each source retains its own coverage status in the companion review. Likewise, reading a discussion of an external theorem does not constitute an independent proof audit of that theorem's original paper. The distinction matters especially for analytic number theory, operator algebras, derived equivalences, and arithmetic descent.

The linked mathematics submodule was inventoried at commit `4eaee87548250b2682affda7d05d947c05554bae`. Its tree contains 174,524 files. This is a file-tree inventory, not a claim that those files were all read. The repository link inventory contains 409 distinct external URLs. A URL may identify a paper, a documentation page, a repository, or a repeated discussion of a result; it is not automatically one independent theorem or one reviewed document. These 409 destinations have not all been reviewed. The two linked author manuscripts, with 223 and 265 PDF pages at their pinned snapshots, were extracted and selectively reviewed [138, 139]. Their complete proofs were not audited page by page. These explicit limits prevent a large source manifest from being mistaken for a completed universal verification.

D.3 The classical books: extraction and passage coverage

All 309 PDF pages of Volume I were extracted by optical character recognition, and all resulting page text was read. Selected mathematical statements were also checked against page images. Formula reconstruction remains selective: OCR can lose subscripts, reverse inclusion signs, merge displayed matrices, or turn a quantified condition into an apparently unconditional statement. Consequently "all page text read" is stronger than a table-of-contents survey but weaker than a verified transcription or a complete proof audit [135].

All 392 PDF pages of Volume II were likewise processed by OCR, and all resulting page text has now been read across the review team. The initial review covered 158 pages, including the complete extracted prose of Chapters VII and VIII, selected passages in the remaining chapters, and the final errata. Two complementary supplemental reads covered the remaining 101 and 133 pages without gaps. The extraction contains approximately 99,149 whitespace-delimited words. It is a working aid, not an authoritative mathematical edition. The page coverage manifests distinguish extraction, complete OCR-content reading, and selected visual verification. Some densely typeset classifications and matrix displays remain badly garbled in OCR; their presence and role were reviewed, but their formulas and proofs were not thereby certified [133].

Volume II also uses two page coordinate systems. Its printed page 114 is PDF page 119; printed page 139 is PDF page 144; and the printed errata page 387 is PDF page 392. Keeping these coordinates explicit makes the most important comparisons reproducible. Full copyrighted scans and bulk OCR are not reproduced as new research content. The public companion instead contains original notes, coverage metadata, bibliographic identifiers, and the new derivations needed by the present argument.

D.4 Three decisive anchors in Volume II

The first anchor is the discussion of topological connections on printed page 114. It motivates a bridge from modules over rings of continuous functions to vector bundles, but a modern statement must specify which K-theory is meant. For a compact Hausdorff space, the finite projective module–vector bundle correspondence gives the familiar interpretation of $K_0(C(X,\mathbb{C}))$ as complex topological $K^0(X)$ [133, 162]. This does not identify algebraic and topological K-theory in every

degree. Already at a point, $K_1^{\mathrm{alg}}(\mathbb{C}) = \mathbb{C}^\times$, whereas $K_1^{\mathrm{top}}(\mathbb{C}) = 0$. The present manuscript therefore uses the historical connection with its category, topology, and degree specified, rather than allowing the notation K_1 alone to carry an unjustified comparison.

The second anchor is the weighted Gram-rank argument on printed page 139. For a real matrix X and a positive diagonal matrix W,

$$\ker(X^{\mathsf{T}}WX) = \ker X, \qquad \operatorname{rank}(X^{\mathsf{T}}WX) = \operatorname{rank} X.$$

Indeed, $v^{\mathsf{T}}X^{\mathsf{T}}WXv$ is a sum of positive weights times squares. Its vanishing forces every component of Xv to vanish. This elementary argument, checked against the scan, supplies a genuine connection to the Hodge quadratic forms developed in the complexes chapter [133]. It also identifies the necessary hypothesis: positivity of the coefficient form. With an indefinite weight, cancellation can invalidate the kernel identity. Over a finite field, a nonzero vector can be isotropic, so a real positivity argument cannot simply be transferred. The algebraic identity and the ordered-field inference must remain separate.

The third anchor is the 1970 errata on printed page 387. It includes a restriction involving totally definite quaternion components for a statement in Volume I and a correction to the argument on Volume I, chapter-page IV.13. The latter concerns an unjustified order-linear splitting of torsion and torsion-free parts. The replacement argument uses the appropriate local and faithfully flat module-theoretic methods [133, 135]. This correction is material to any attempt to reuse the books for equivalences of module categories. Our matrix-ring Morita calculation has an independent direct proof and does not invoke the faulty splitting. A historical source should be read together with its published corrections, especially when the correction changes the mechanism of a proof.

D.5 What survives scrutiny in the upstream research program

The fully read source files contain substantial mathematical repairs and careful negative results. They distinguish infinite orders from finite quotients, path degree from the formal order parameter, and algebraic, formal, Laurent, and operator-norm completions. Their nodal examples have explicit alternating resolutions. Their semi-invariant constructions distinguish determinant weights from ordinary invariants, and their representation-theoretic examples state where Schur constructions cease to preserve additivity or a chosen scalar action [137]. These are useful structural results even when they do not settle the eventual analytic problem.

The same source explicitly identifies limitations of several proposed arithmetic transfers. A complex dessin does not by itself recover rational descent data. Inserting a known two-dimensional arithmetic matrix into a larger projective module changes an ordinary determinant by multiplicity. A later branch coefficient functor can remove a particular multiplicity, but it changes the determinant target and gives up specified symmetries. Tame descent can preserve an endomorphism order while failing to preserve the desired integral column module. None of these distinctions can be replaced by a generic assertion that the categories are "equivalent" [137].

The analytic papers are equally explicit about the missing global step. Self-adjoint finite spectral data yield a critical-line conclusion only after an appropriately normalized determinant converges locally uniformly to the intended arithmetic function. Strong operator convergence alone does not provide trace-norm determinant convergence. Finite Euler products cannot supply a locally uniform zero-bearing limit on a region where the approximants remain zero-free and the limit is nonzero. Positivity of an unrelated matrix kernel does not prove positivity of an arithmetic moment form. The papers' conditional closure statements are useful because they locate these obligations precisely; they do not discharge them [137].

There is one editorial qualification worth retaining in future revisions. An earlier discussion says that no multiplicity-removing functor has been supplied, whereas a later chapter constructs a branch functor under changed requirements. The earlier sentence is accurate when restricted to the naive tensor insertion being discussed. Read globally, it is too broad. The mathematical distinction is between failure of that insertion and success of a narrower construction that relinquishes symmetry. The review records this qualification rather than silently replacing either claim.

D.6 From sources to the new theorems in this monograph

The complexes chapters distinguish three kinds of dependence. Elementary results such as the dimension of a two-step rank stratum, its tangent equation, finite Hodge decomposition, and the orthogonal quadric calculation receive proofs in the text. The Buchsbaum–Eisenbud exactness criterion and the geometry of varieties of complexes are attributed to their separate primary sources; the latter attribution includes De Concini and Strickland [27, 40, 41]. The Atiyah–Singer index theorem and the analytical existence theory behind elliptic Hodge decomposition remain external advanced results, stated with their hypotheses rather than presented as newly proved here [12, 13].

This discipline matters at the gauge-theoretic interface. The covariant derivative squares to a curvature action. A deformation sequence is a complex under specified background conditions, such as flatness or the anti-self-dual equation in the appropriate projected sequence. Its index concerns a classical elliptic problem and cannot, without another theorem, be read as a positive spectral gap for an interacting quantum Hamiltonian [11, 16]. The Yang–Mills chapter separately assesses the recent construction and its stated acceptance criteria. The source concordance does not enlarge those claims.

D.7 A finite-field correction in the source matrix reductions

A further qualification arises in Volume II, printed page 351 (chapter-page X 65; PDF page 356), in the reduction of extension matrices over a cubic finite-field extension [133]. The scan asserts that, outside characteristic 3, one may choose a primitive element whose cube belongs to the base field. Visual inspection confirms that this is an overstatement in the source, rather than an OCR error. The weaker relation used subsequently has the following direct justification.

Proposition D.1 (Trace-zero cubic presentation)**.** *Let S/k be an extension of finite fields of degree 3, with* $\operatorname{char} k \neq 3$*. There is an element α generating S over k, and elements $a, b \in k$ with $b \neq 0$, such that*

$$\alpha^3 = a\alpha + b, \qquad \alpha^{-1} = b^{-1}\alpha^2 - ab^{-1}.$$

There need not be a generator whose cube lies in k.

Proof. Choose $\beta \in S \setminus k$. The tower formula shows that $[k(\beta) : k]$ divides 3 and is greater than 1, hence equals 3. Thus β generates S. The field trace is k-linear, and the trace of $c \in k$ is $3c$: multiplication by c on the three-dimensional vector space S has matrix cI_3. Division by 3 is permitted, so put

$$\alpha = \beta - \frac{\operatorname{Tr}_{S/k}(\beta)}{3}.$$

Then $\operatorname{Tr}_{S/k}(\alpha) = 0$. Since the subtracted quantity lies in k, the element α remains outside k and generates S.

Its monic minimal polynomial has degree 3. In the basis $1, \alpha, \alpha^2$, multiplication by α has the corresponding companion matrix. Its trace is the negative of the quadratic coefficient, so that

coefficient vanishes. We may therefore write the polynomial as $T^3 - aT - b$. Its constant term cannot vanish: otherwise T would divide this irreducible cubic. Consequently $b \neq 0$, and the polynomial equation gives $\alpha(\alpha^2 - a) = b$. Multiplication by b^{-1} proves the inverse formula.

For the last assertion, take $S = \mathbb{F}_8$ and $k = \mathbb{F}_2$. Every polynomial $T^3 - c$, with $c \in \mathbb{F}_2$, has a root in $\mathbb{F}_2$: choose 0 when $c = 0$ and 1 when $c = 1$. It is therefore reducible. It cannot be the degree-three minimal polynomial of a generator of $\mathbb{F}_8/\mathbb{F}_2$. □

This repair establishes exactly the depressed-cubic relation and quadratic inverse expression needed at that point. It leaves the source's separate characteristic-3 argument separate. It also illustrates the scope of the review: correcting this local assertion does not certify every subsequent matrix transformation or formalize the classification theorem in Lean.

D.8 Formal proofs, computations, and remaining coverage

The original Lean companion file contains 25 selected declarations compiled against Lean 4.24.0 and the pinned mathlib revision recorded in its manifest. Their types concern actual finite-dimensional maps, concrete matrix identities, real positive semidefiniteness, a determinant factorization, modular kernel relations, and noise cancellation. Each declaration has a printed dependency report; the successful build contains no project-specific axiom or `sorry`. Standard foundational dependencies are recorded rather than described as an absence of all axioms. This is a precise verification result for the declarations present in the file.

An additional seven-declaration file verifies finite Pauli identities, binary phase and shift commutation, the chain-boundary specialization, and a conditional integer Euler-rank substitution. Its independent successful build and axiom audit supplied the first 32 selected declarations. The second audit adds sixteen declarations in DiagramGrammar, bringing the three-module inventory to 48; its scope is recorded in the formalization appendix.

It is not a formalization of every theorem in either classical book, the entire source repository, Shor's quantum circuit, the Atiyah–Singer index theorem, or a Yang–Mills continuum construction. Upstream reports of earlier formal components likewise retain their literal theorem scopes. A numerical example, a source-code formula, and a proved quantified theorem are different deliverables even when they concern the same matrix. The companion manifests make those distinctions visible.

The synthesis chapter's periodic order example was also checked directly: the annihilators give the stated exact alternating resolution; tensoring with the quotient gives the claimed odd-degree torsion groups; and the real finite matrix specialization has middle Hodge operator $(1+t^2)I$. Scaling the differentials then preserves middle exactness while sending the least Hodge eigenvalue to zero. These calculations show why the review's main limitation is mathematical as well as bibliographic: an algebraic identification does not supply the uniform estimates required by a spectral limit or the cost bounds required by an algorithmic claim. Completing the unreviewed corpus would improve coverage; proving those additional estimates would require genuinely additional mathematics.

Appendix E

Primary sources and source-pinned implementation register

Reading scope: the native theorem/code passages discussed in the audit were inspected at the commits below. External entries are attributed results, not claims that their complete proofs were reverified in Lean. The present review checked primary metadata and the relevant accessible statements; selected sections of Yun, ConvexTok, the syntax–semantics paper and the BGG papers were examined. Unread surrounding arguments are not represented as a complete full-text review. Numerical experiments have a separate source/hash/run record. No third-party full text is redistributed.

P1 Tropical Quivers of Architectures research program. Commit `2b6ab79a685e2c6233fb8892e59589951feeddad`; especially graph tokenization, graph-to-graph approximation and bounded grammar claims. https://github.com/amelie-iska/Tropical_Quivers_of_Archs/tree/2b6ab79a685e2c6233fb8892e59589951feeddad

P2 ToricGT source and manuscript. Commit `a5199bc989a6fb170e21220e5e87bf95bc2fc0b4`; tokenizer, graph model, GFlowNet primitive and fixed-cell/hybrid approximation statements. https://github.com/amelie-iska/ToricGT/tree/a5199bc989a6fb170e21220e5e87bf95bc2fc0b4

P3 TropicalGT implementation. Commit `f8ffc7150afdd83bace8078c86fde7c8eeb0e08e`; manuscript and tropical attention implementation. https://github.com/amelie-iska/TropicalGT/tree/f8ffc7150afdd83bace8078c86fde7c8eeb0e08e

P4 Mora-3 implementation. Commit `75ac2d696f42da0a673fc43a20d84686a2ceed40`; TokenGT, tropical attention, continuous GFN, forest selection and learned analogy interfaces. Historical working-preprint status is not a current training-completion claim. https://github.com/amelie-iska/Mora-3/tree/75ac2d696f42da0a673fc43a20d84686a2ceed40

P5 ToricGT-Program. Commit `f91ba1811b5a4f70d53c7120acf8ca2a0049f85b`; dynamical, memory, trajectory and algebraic-training manuscripts. https://github.com/amelie-iska/ToricGT-Program/tree/f91ba1811b5a4f70d53c7120acf8ca2a0049f85b

R1 Yun et al. *Are Transformers universal approximators of sequence-to-sequence functions?* ICLR 2020; version 2. Theorems 2–3 use their stated L^p metric. https://arxiv.org/abs/1912.10077v2

R2 Kim et al. *Pure Transformers are Powerful Graph Learners.* NeurIPS 2022; version 2. TokenGT construction and graph expressivity. https://arxiv.org/abs/2207.02505v2

R3 Leshno, Lin, Pinkus and Schocken. *Multilayer feedforward networks with a nonpolynomial activation function can approximate any function.* Neural Networks 6 (1993), 861–867. https://doi.org/10.1016/S0893-6080(05)80131-5

R4 Pérez, Barceló and Marinkovic. *Attention is Turing-Complete.* JMLR 22(75), 2021. https://jmlr.org/papers/v22/20-302.html

R5 Merrill and Sabharwal. *The Expressive Power of Transformers with Chain of Thought.* ICLR 2024; version 5. https://arxiv.org/abs/2310.07923v5

R6 Tempus, Whittington, Schmidt, Komm and Pimentel. *Tokenisation via Convex Relaxations.* 2026 preprint, version 1. Integer formulation, LP relaxation and objective-specific certification.

https://arxiv.org/abs/2605.22821v1

R7 Malkin et al. *Trajectory balance: Improved credit assignment in GFlowNets.* NeurIPS 2022. https://arxiv.org/abs/2201.13259

R8 Lahlou et al. *A theory of continuous generative flow networks.* ICML 2023. Generalized measure-based state spaces. https://proceedings.mlr.press/v202/lahlou23a.html

R9 Watrous. *The Theory of Quantum Information.* Cambridge University Press, 2018. Channels, distinguishability and operational resource definitions. https://cs.uwaterloo.ca/~watrous/TQI/

R10 Fomin and Zelevinsky. *Cluster algebras I: Foundations.* 2001 preprint; JAMS 2002. https://arxiv.org/abs/math/0104151

R11 Buan et al. *Tilting theory and cluster combinatorics.* Advances in Mathematics 204 (2006). https://arxiv.org/abs/math/0402054

R12 Keller and Yang. *Derived equivalences from mutations of quivers with potential.* Advances in Mathematics 226 (2011), 2118–2168. https://arxiv.org/abs/0906.0761

R13 Assem, Brüstle, Charbonneau-Jodoin and Plamondon. *Gentle algebras arising from surface triangulations.* https://arxiv.org/abs/0903.3347

R14 Opper, Plamondon and Schroll. *A geometric model for the derived category of gentle algebras.* https://arxiv.org/abs/1801.09659

R15 Lekili and Polishchuk. *Derived equivalences of gentle algebras via Fukaya categories.* Version 5, 2019. Homologically smooth graded gentle algebras, boundary stops and line fields. https://arxiv.org/abs/1801.06370v5

R16 Chomsky, interview by Polychroniou, 3 May 2023. *Noam Chomsky Speaks on What ChatGPT Is Really Good For.* https://chomsky.info/20230503-2/

R17 Marcolli, Berwick and Chomsky. *Syntax-semantics interface: an algebraic model.* Version 1, 2023; especially Section 7, which relates attention to the algebraic framework. https://arxiv.org/abs/2311.06189v1

R18 Marcolli and Port. *Graph Grammars, Insertion Lie Algebras, and Quantum Field Theory.* Mathematics in Computer Science 9 (2015), 391–408. https://arxiv.org/abs/1502.07796

R19 Kallini et al. *Mission: Impossible Language Models.* ACL 2024. https://aclanthology.org/2024.acl-long.787/

R20 Hunter. *Kallini et al. (2024) Do Not Compare Impossible Languages with Constituency-based Ones.* Computational Linguistics 51(2), 2025. https://aclanthology.org/2025.cl-2.7/

R21 Gulordava et al. *Colorless Green Recurrent Networks Dream Hierarchically.* NAACL 2018. https://aclanthology.org/N18-1108/

R22 Lakretz et al. *The emergence of number and syntax units in LSTM language models.* NAACL 2019. https://aclanthology.org/N19-1002/

B1 Bernstein, Gelfand and Gelfand. *Differential operators on the base affine space and a study of $\mathfrak{g}$-modules.* In *Lie groups and their representations.* Original category/resolution treatment; selected definitions and resolution discussion inspected, not the entire scanned volume. https://www.math.utah.edu/~ptrapa/math-library/bernstein/bgg-resolution.pdf

B2 Beilinson, Ginzburg and Soergel. *Koszul duality patterns in representation theory.* JAMS 9 (1996), 473–527; Introduction and Theorems 1.1.1–1.1.3. https://doi.org/10.1090/S0894-0347-96-00192-0

B3 Eisenbud, Fløystad and Schreyer. *Sheaf Cohomology and Free Resolutions over Exterior Algebras.* Version 2, 2001; Transactions AMS 355 (2003), 4397–4426. Introduction and correspondence statements. https://arxiv.org/abs/math/0104203v2

B4 Humphreys. *Representations of Semisimple Lie Algebras in the BGG Category O.* AMS, 2008. Bibliographic reference; not claimed fully reread. https://doi.org/10.1090/gsm/094

B5 Elias and Williamson. *The Hodge theory of Soergel bimodules.* Annals of Mathematics 180 (2014), 1089–1136. External character/positivity theorem, not independently reproved here. https://arxiv.org/abs/1212.0791v2

B6 Braden, Licata, Proudfoot and Webster. *Hypertoric category O*. Advances in Mathematics 231 (2012), 1487–1545. Use corrected version 5 (2022), which records a correction to Lemma 8.25. https://arxiv.org/abs/1010.2001v5

Appendix F

Arithmetic and quantum-renormalization source register

E1 Roetteler, Naehrig, Svore and Lauter. *Quantum resource estimates for computing elliptic curve discrete logarithms.* ASIACRYPT 2017. https://arxiv.org/abs/1706.06752. This is a scalable circuit-resource reference, not a performance claim for the bounded compiler implemented here.

Q1 Cong, Choi and Lukin. *Quantum Convolutional Neural Networks.* Nature Physics 15, 1273–1278 (2019). https://arxiv.org/abs/1810.03787. The present unitary/trace hierarchy is an explicit finite channel inspired by hierarchical coarse graining; it is not represented as a verbatim implementation of every construction in that paper.

Q2 Girardi and De Palma. *Trained quantum neural networks are Gaussian processes.* Version 2, 2026. https://arxiv.org/abs/2402.08726v2. The light-cone, initialization, kernel and scaling assumptions must accompany any transfer of its trained-network limit.

G1 Assem, Brüstle, Charbonneau-Jodoin and Plamondon. *Gentle algebras arising from surface triangulations.* Algebra and Number Theory 4 (2010), 201–229. https://arxiv.org/abs/0903.3347. The unpunctured marked-surface hypotheses are retained.

G2 De la Peña and Redondo. *Covering functors without groups.* Journal of Algebra 321 (2009), 3816–3826. DOI 10.1016/j.jalgebra.2009.02.023. The finite voltage construction in this volume is specified directly rather than claiming all covering-functor results.

C1 NASA LAMBDA / WMAP Science Team. *WMAP DR5 ILC map description and downloads.* https://lambda.gsfc.nasa.gov/product/wmap/dr5/ilc_map_info.html; https://lambda.gsfc.nasa.gov/product/wmap/dr5/ilc_map_get.html. Product resolution, Galactic coordinates and the approximately ten-degree reliability qualification determine this study's broad scales.

C2 NASA LAMBDA / WMAP Science Team. *WMAP DR5 sky maps and masks.* https://lambda.gsfc.nasa.gov/product/wmap/dr5/skymap_info.html; https://lambda.gsfc.nasa.gov/product/wmap/dr5/masks_get.html. Native filenames, units, ordering and exact bytes are recorded in the acquisition manifest.

C3 The original GRH-research chapter *Cyclic cosmology* and its frozen source ledger. https://github.com/amelie-iska/GRH-research/tree/fd57c5251bc40c388cc25aa83e25f8a2cdf995b7/extension/review/cosmology. Original theoretical claims and published observational comparisons remain separately attributed; this continuation is not a replacement for those source papers.

Appendix G

Source-bound verification of the earlier computational continuation

The following preserved execution record refers to the pinned program-three release. This unified edition has an additional integration record; earlier test counts are not represented as expanded proof coverage.

The independent runner accepted 22 new formal declarations, 142 regression methods under each Python mode, 32 browser checks, and 9 complete JavaScript/Python sky-score replays. The maximum score discrepancy was 3.69e-14.

The published Python tests exercise the ordinary and optimized interpreters. The formal audit prints each new declaration and its axiom inventory, rejects warnings, rejects `sorryAx`, and binds the exact source to Lean 4.24.0 and the recorded mathlib revision. The finite arithmetic module checks its entire small curve table; the general Kraus and rank-count modules verify their displayed quantified statements. The whole physical interpretation is not encoded by a successful process exit.

The browser executes the emitted reversible gates and the learned four-qubit circuit. It independently recomputes all 72 fixed sky templates and all three methods for observed maps or explicitly labeled injected profiles. Its scientific JSON assets are SHA-256 checked before use. Interactive state is exportable and URL-backed; keyboard and mobile behavior are tested. A corrupted asset is rejected rather than displayed under a successful-verification badge.

The original workstation completed genuine cuStateVec checks of both $p = 7$ arithmetic circuits on a complex superposition spanning all 2,048 declared input words. The observed maximum amplitude discrepancy and dirty-ancilla probability were zero. Six trained-channel gate checks agreed with the independent Kraus implementation to 5.56×10^{-16}. The subsequent independent repository runner has no GPU and does not relabel those measurements as its own. The $p = 13$ circuits retain exhaustive basis verification without a dense-GPU claim.

Desktop Commander reached its monthly usage limit during the continuation. No retry, reconnection, service restart, credential extraction or alternate access to the paused device was used. Independent repository CI rebuilt the source and evidence on a standard public CPU runner. Remote clone synchronization after that pause is not asserted. The original Mora-3 training campaign remains outside this task.

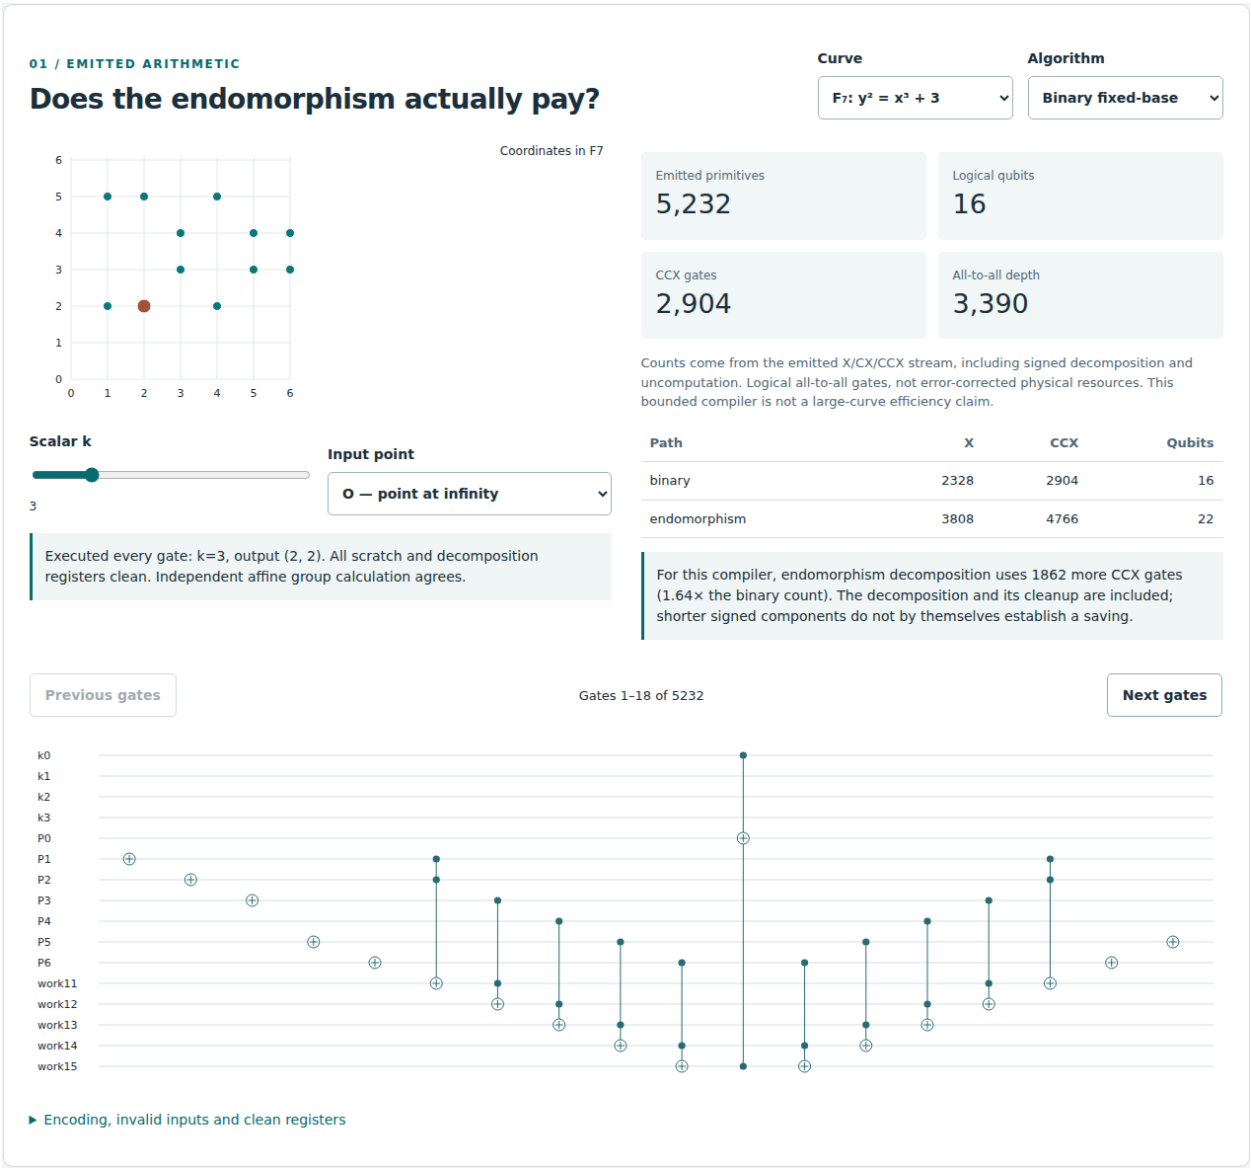


Figure G.1: The actual arithmetic explorer: finite-field points, complete emitted-gate counts, invalid-word behavior and a paginated circuit stream.

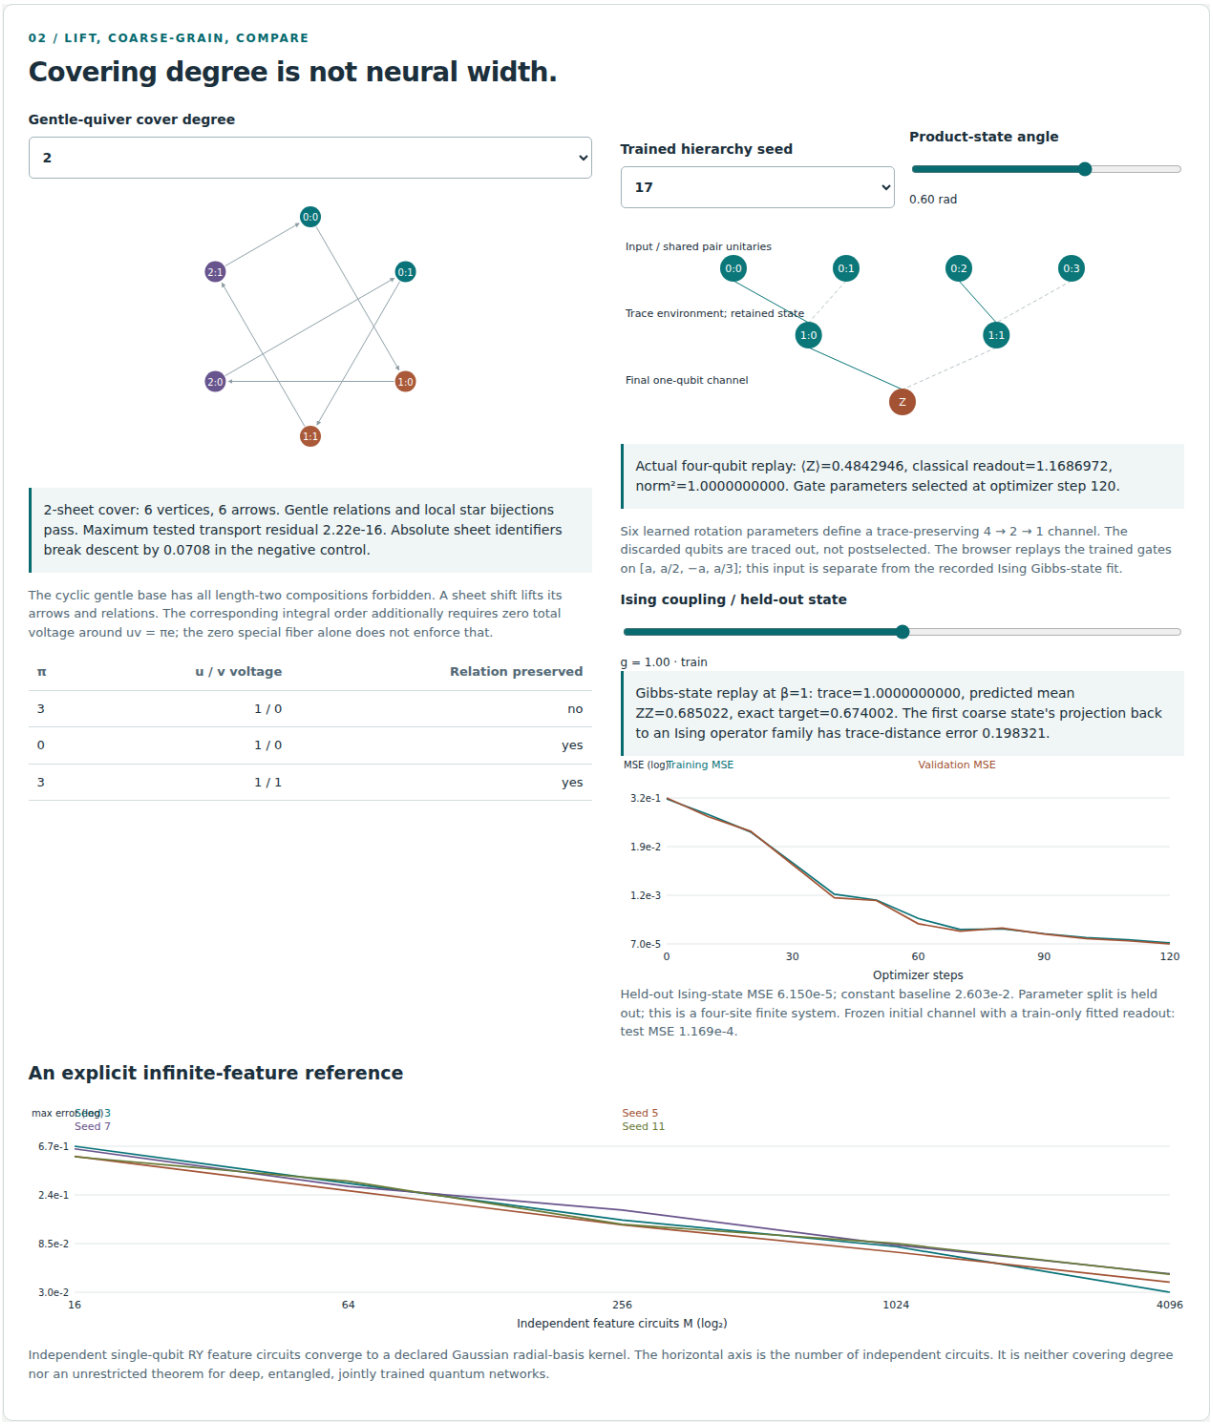


Figure G.2: Voltage covers, learned hierarchical channels, effective-Hamiltonian error and the separately scaled independent-feature width experiment.

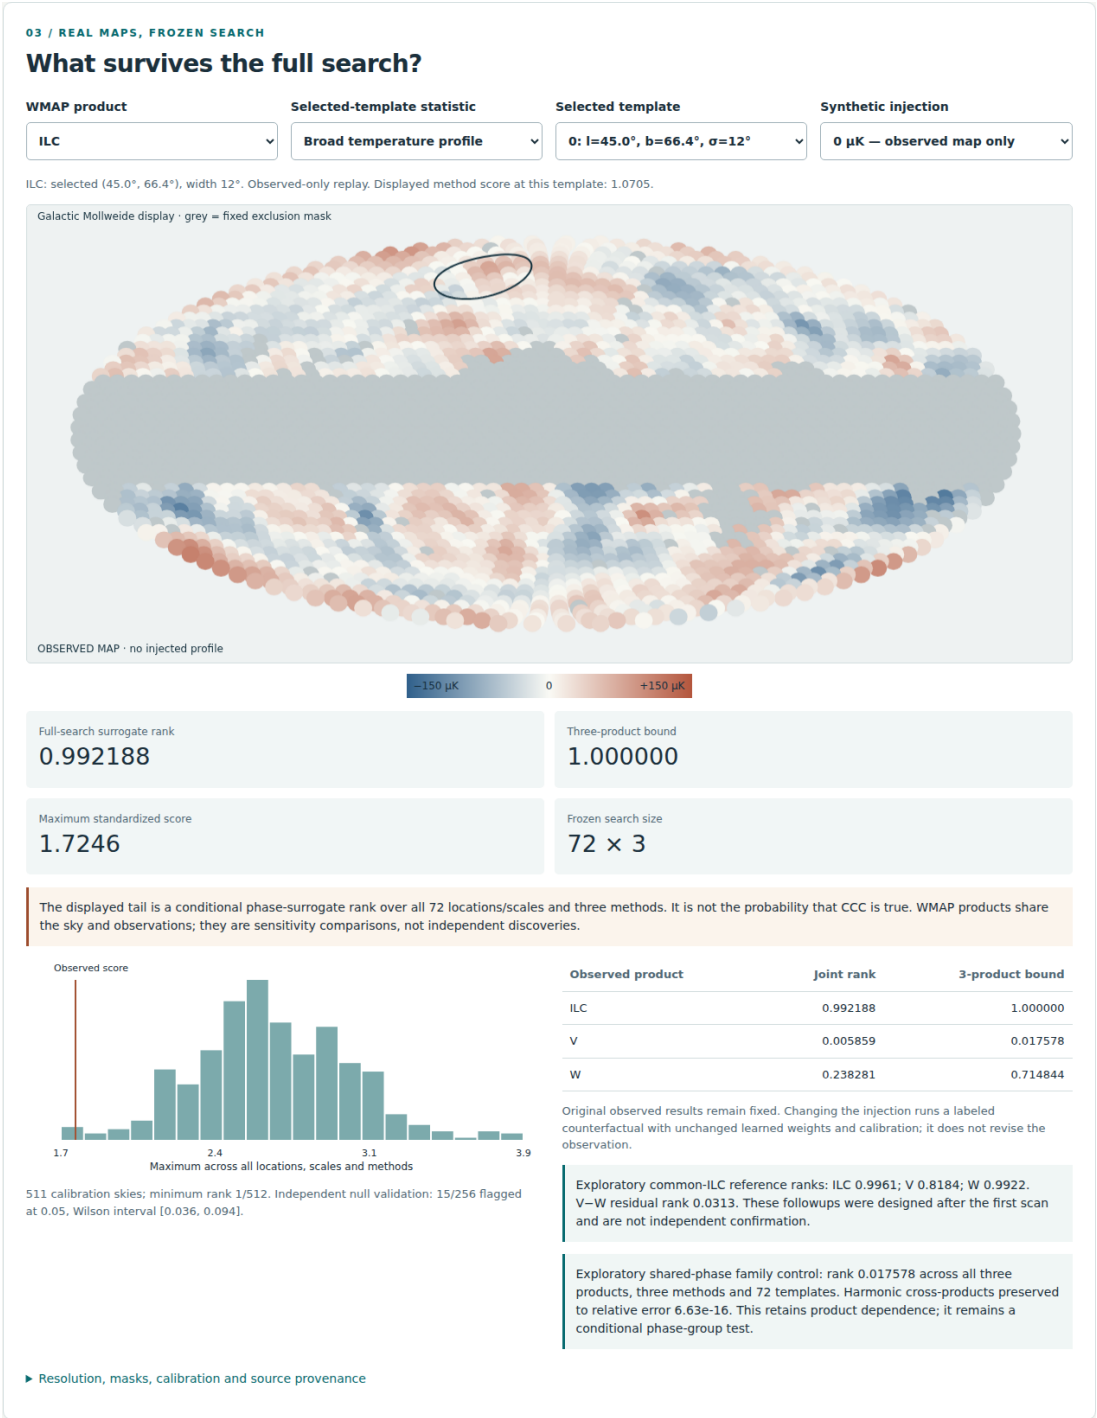


Figure G.3: Observed WMAP maps and a complete frozen search. Synthetic injections and post-inspection sensitivity tests are explicitly distinguished from the original observations.

Bibliography

[1] Kaan Akin, David A. Buchsbaum, and Jerzy Weyman. Schur functors and schur complexes. *Advances in Mathematics*, 44(3):207–278, 1982. doi: 10.1016/0001-8708(82)90039-1. URL https://jerzyweyman.webflow.io/research/schur. Author-hosted full text retrieved; Section II.2, printed pages 223–234, read in full, including the shuffle presentation and universal-freeness proof. Diagram conventions are transposed explicitly in this monograph.

[2] Francesca Albertini and Domenico D'Alessandro. Notions of controllability for bilinear multilevel quantum systems. *IEEE Transactions on Automatic Control*, 48(8):1399–1403, 2003. doi: 10.1109/TAC.2003.815027. URL https://www.math.unipd.it/~albertin/26.pdf.

[3] Dallas Albritton, Elia Brué, and Maria Colombo. Non-uniqueness of Leray solutions of the forced Navier–Stokes equations. *Annals of Mathematics*, 196(1):415–455, 2022. doi: 10.4007/annals.2022.196.1.3. URL https://annals.math.princeton.edu/2022/196-1/p03.

[4] Michael P. Allen. Introduction to molecular dynamics simulation. In Norbert Attig, Kurt Binder, Helmut Grubmüller, and Kurt Kremer, editors, *Computational Soft Matter: From Synthetic Polymers to Proteins*, volume 23 of *NIC Series*, pages 1–28. John von Neumann Institute for Computing, 2004. URL https://dasher.wustl.edu/chem430/readings/md-intro-2.pdf. All 28 PDF pages, including the title page and references, reviewed.

[5] Alberto Amo, Jérôme Lefrère, Simon Pigeon, Claire Adrados, Cristiano Ciuti, Iacopo Carusotto, Romuald Houdré, Elisabeth Giacobino, and Alberto Bramati. Superfluidity of polaritons in semiconductor microcavities. *Nature Physics*, 5:805–810, 2009. doi: 10.1038/nphys1364. URL https://arxiv.org/abs/0812.2748.

[6] Daniel An, Krzysztof A. Meissner, Pawel Nurowski, and Roger Penrose. Apparent evidence for Hawking points in the CMB sky. *Monthly Notices of the Royal Astronomical Society*, 495(3):3403–3408, 2020. doi: 10.1093/mnras/staa1343. URL https://arxiv.org/abs/1808.01740v5. Reviewed the later arXiv version 5, 9 August 2022, seven pages; publication year and reviewed revision are distinct.

[7] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther. Finite element exterior calculus, homological techniques, and applications. *Acta Numerica*, 15:1–155, 2006. doi: 10.1017/S0962492906210018. URL https://sites.math.rutgers.edu/~falk/papers/acta.pdf.

[8] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther. Finite element exterior calculus: from Hodge theory to numerical stability. *Bulletin of the American Mathematical Society*, 47(2):281–354, 2010. doi: 10.1090/S0273-0979-10-01278-4. URL https://arxiv.org/abs/0906.4325.

[9] Michael F. Atiyah. Vector bundles over an elliptic curve. *Proceedings of the London Mathematical Society*, s3-7(1):414–452, 1957. doi: 10.1112/plms/s3-7.1.414. URL https://doi.org/10.1112/plms/s3-7.1.414.

[10] Michael F. Atiyah and Raoul Bott. A lefschetz fixed point formula for elliptic complexes i. *Annals of Mathematics*, 86(2):374–407, 1967. doi: 10.2307/1970694. URL https://annals.math.princeton.edu/1967/86-2/p07.

[11] Michael F. Atiyah and Raoul Bott. The yang–mills equations over riemann surfaces. *Philosophical Transactions of the Royal Society of London. Series A*, 308(1505):523–615, 1983. doi: 10.1098/rsta.1983.0017. URL https://www.math.utoronto.ca/mgualt/Morse%20Theory/Atiyah-Bott.pdf.

[12] Michael F. Atiyah and Isadore M. Singer. The index of elliptic operators i. *Annals of Mathematics*, 87(3):484–530, 1968. doi: 10.2307/1970715. URL https://annals.math.princeton.edu/1968/87-3/p05.

[13] Michael F. Atiyah and Isadore M. Singer. The index of elliptic operators iii. *Annals of Mathematics*, 87(3):546–604, 1968. doi: 10.2307/1970717. URL https://annals.math.princeton.edu/1968/87-3/p07.

[14] Michael F. Atiyah, Nigel J. Hitchin, Vladimir G. Drinfeld, and Yuri I. Manin. Construction of instantons. *Physics Letters A*, 65(3):185–187, 1978. doi: 10.1016/0375-9601(78)90141-X. URL https://doi.org/10.1016/0375-9601(78)90141-X.

[15] Michael F. Atiyah, Nigel J. Hitchin, and Isadore M. Singer. Self-duality in four-dimensional riemannian geometry. *Proceedings of the Royal Society of London A*, 362(1711):425–461, 1978. doi: 10.1098/rspa.1978.0143. URL https://doi.org/10.1098/rspa.1978.0143.

[16] Michael F. Atiyah, Nigel J. Hitchin, and Isadore M. Singer. Self-duality in four-dimensional riemannian geometry. *Proceedings of the Royal Society of London. Series A*, 362:425–461, 1978. doi: 10.1098/rspa.1978.0143.

[17] Stéphane Attal, Francesco Petruccione, Christophe Sabot, and Ilya Sinayskiy. Open quantum random walks. *Journal of Statistical Physics*, 147(4):832–852, 2012. doi: 10.1007/s10955-012-0491-0. URL https://arxiv.org/abs/1402.3253.

[18] Ryan Babbush, Adam Zalcman, Craig Gidney, Michael Broughton, Tanuj Khattar, Hartmut Neven, Thiago Bergamaschi, Justin Drake, and Dan Boneh. Securing elliptic curve cryptocurrencies against quantum vulnerabilities: Resource estimates and mitigations. *PRX Quantum*, 7:031001, 2026. doi: 10.1103/j3xf-bw18. URL https://journals.aps.org/prxquantum/abstract/10.1103/j3xf-bw18. Published 21 August 2026; distinguishes disclosed resource claims from withheld circuit details and their sampled verification certificate.

[19] Tadeusz Bałaban. Renormalization group approach to lattice gauge field theories: I. generation of effective actions in a small field approximation and a coupling constant renormalization in four dimensions. *Communications in Mathematical Physics*, 109:249–301, 1987. doi: 10.1007/BF01215223. URL https://doi.org/10.1007/BF01215223.

[20] Razvan Barbulescu, Mugurel Barcau, and Vicenţiu Paşol. Extending Regev's quantum algorithm to elliptic curves. In *Progress in Cryptology—LATINCRYPT 2025*, pages 234–266, 2025. doi: 10.1007/978-3-032-06754-8_9. Publisher abstract and author presentation reviewed; subscription full chapter not available in this review.

[21] Razvan Barbulescu, Mugurel Barcau, and Vicenţiu Paşol. Extending Regev's quantum algorithm to elliptic curves: Author presentation, 2025. URL https://ciencias.medellin.unal.edu.co/eventos/latincrypt/images/14.25_Extending_Regevs_quantum_algorithm_to_elliptic_curves_Barbulescu_barcau_Pasol.pdf. LATINCRYPT 2025 presentation; used only for its stated scope and hypotheses.

[22] G. Barontini, R. Labouvie, F. Stubenrauch, A. Vogler, V. Guarrera, and H. Ott. Controlling the dynamics of an open many-body quantum system with localized dissipation. *Physical Review Letters*, 110(3):035302, 2013. doi: 10.1103/PhysRevLett.110.035302. URL https://arxiv.org/abs/1212.4824.

[23] G. V. Belyi. On Galois extensions of a maximal cyclotomic field. *Mathematics of the USSR-Izvestiya*, 14(2):247–256, 1980. doi: 10.1070/IM1980v014n02ABEH001096. URL https://www.mathnet.ru/eng/im1682. English translation; original Russian publication 1979. Primary bibliographic record inspected; full PDF access was blocked.

[24] Eve Bodnia, Vlad Isenbaev, Kellan Colburn, Joe Swearngin, and Dirk Bouwmeester. The quest for CMB signatures of Conformal Cyclic Cosmology. *Journal of Cosmology and Astroparticle Physics*, 2024(05):009, 2024. doi: 10.1088/1475-7516/2024/05/009. URL https://arxiv.org/abs/2208.06021v3. Reviewed arXiv version 3, 25 January 2024, 27 PDF pages including title page; printed pages 1–26.

[25] braintruffle. Fluid dynamics feels natural once you start with quantum mechanics, May 2021. URL https://www.youtube.com/watch?v=MXs_vkc8hpY. Published 2 May 2021. All 249 English timed transcript segments, 0:00–32:36, reviewed through the public transcript interface on 13 November 2026. The transcript is not redistributed.

[26] Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLOS ONE*, 11(2):e0150171, 2016. doi: 10.1371/journal.pone.0150171. URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0150171.

[27] David A. Buchsbaum and David Eisenbud. What makes a complex exact? *Journal of Algebra*, 25:259–268, 1973. doi: 10.1016/0021-8693(73)90044-6. URL https://people.math.sc.edu/kustin/teaching/746/Buchsbaum-Eisenbud-What-makes-a-complex-exact.pdf.

[28] Tristan Buckmaster and Vlad Vicol. Nonuniqueness of weak solutions to the Navier–Stokes equation. *Annals of Mathematics*, 189(1):101–144, 2019. doi: 10.4007/annals.2019.189.1.3. URL https://annals.math.princeton.edu/2019/189-1/p03.

[29] M. C. R. Butler and Claus Michael Ringel. Auslander–reiten sequences with few middle terms and applications to string algebras. *Communications in Algebra*, 15(1–2):145–179, 1987. doi: 10.1080/00927878708823416. Publisher and author bibliographic records checked; the general classification is an external theorem, while the examples in this monograph are proved directly.

[30] A. R. Calderbank and Peter W. Shor. Good quantum error-correcting codes exist. *Physical Review A*, 54:1098–1105, 1996. doi: 10.1103/PhysRevA.54.1098. URL https://math.mit.edu/~shor/papers/good-codes.pdf.

[31] Andrew T. Carroll and Calin Chindris. On the invariant theory for acyclic gentle algebras, 2013. URL https://arxiv.org/abs/1210.3579v2. Primary preprint version 2; Theorems 1–2 and

their component/stability hypotheses inspected. Published in Transactions of the American Mathematical Society 367 (2015), 3481–3508, DOI 10.1090/S0002-9947-2014-06191-6.

[32] Andrew T. Carroll and Jerzy Weyman. Semi-invariants for gentle string algebras, 2011. URL https://arxiv.org/abs/1106.0774v1. Full 42-page preprint read across the review team. Sections 3–5 reviewed for rank strata, Schur filtration, and coordinate actions. Published as Semi-invariants for gentle algebras, Contemporary Mathematics 592 (2013), DOI 10.1090/conm/592/11862. Source typographical issues are recorded in the review ledger; the monograph states explicit corrected conventions.

[33] Iacopo Carusotto and Cristiano Ciuti. Quantum fluids of light. *Reviews of Modern Physics*, 85 (1):299–366, 2013. doi: 10.1103/RevModPhys.85.299. URL https://arxiv.org/abs/1205.6500.

[34] Noam Chomsky. Noam chomsky speaks on what ChatGPT is really good for, 2023. URL https://chomsky.info/20230503-2/. Interview by C. J. Polychroniou, 3 May; author's archive.

[35] Clay Mathematics Institute. Navier–Stokes announcement, September 2026. URL https://www.claymath.org/news/navier-stokes-announcement/. 11 September 2026; retrieved 13 November 2026.

[36] Computational companion to this monograph. Quantum ToricGT token-dag experiment: pinned data, nine GPU training runs, independent references, and exact certificates, 2026. Artifacts in `extension/quantum/toricgt/`; complete protocol and hashes in `results/training-v1/benchmark.json`; finite tests and measurements are distinguished from general mathematical proofs.

[37] Iris Cong, Soonwon Choi, and Mikhail D. Lukin. Quantum convolutional neural networks. *Nature Physics*, 15:1273–1278, 2019. doi: 10.1038/s41567-019-0648-8. URL https://arxiv.org/abs/1810.03787v2. Published 26 August 2019; preprint version 2 dated 2 May 2019.

[38] Peter Constantin, Mihaela Ignatova, and Vlad Vicol. Regularity of asymptotically axisymmetric solutions to the 3d Navier–Stokes equations with analytic forcing, 2026. URL https://arxiv.org/abs/2609.20803v2. Version 2, 29 September 2026. Theorem 1.1 and equations (1.5)–(1.8); accessed 13 November 2026.

[39] Predrag Cvitanović. *Group Theory: Birdtracks, Lie's, and Exceptional Groups*. Princeton University Press, 2008. URL https://birdtracks.eu/version9.0/GroupTheory.pdf. Author's electronic version 9.0.1, 8 April 2011; diagram conventions in Chapter 4.

[40] Corrado De Concini and Elisabetta Strickland. On the variety of complexes. *Advances in Mathematics*, 41(1):57–77, 1981. doi: 10.1016/S0001-8708(81)80004-7. Geometric background; accessible original-author exposition also in De Concini–Eisenbud–Procesi, Section 16.

[41] Corrado De Concini, David Eisenbud, and Claudio Procesi. *Hodge Algebras*, volume 91 of *Astérisque*. Société mathématique de France, 1982. URL https://archive.numdam.org/item/AST_1982__91__1_0/. Section 16: The variety of complexes, pages 69–73.

[42] César R. de Oliveira and Alessandra A. Verri. Self-adjoint extensions of Coulomb systems in 1, 2 and 3 dimensions. *Annals of Physics*, 324(2):251–266, 2009. doi: 10.1016/j.aop.2008.06.001. URL https://arxiv.org/abs/0806.2764. All 23 pages of the retrieved arXiv text reviewed.

Used for operator-domain distinctions; the current example flagged in the source audit is not relied upon.

[43] Eric Dennis, Alexei Kitaev, Andrew Landahl, and John Preskill. Topological quantum memory. *Journal of Mathematical Physics*, 43(9):4452–4505, 2002. doi: 10.1063/1.1499754. URL https://arxiv.org/abs/quant-ph/0110143.

[44] Harm Derksen and Jerzy Weyman. Semi-invariants of quivers and saturation for Littlewood–Richardson coefficients. *Journal of the American Mathematical Society*, 13(3):467–479, 2000. doi: 10.1090/S0894-0347-00-00331-3. URL https://www.math.utah.edu/~ptrapa/math-library/knutson/derksen-weyman-quiver-JAMS.pdf. Primary text, exact sequence and determinant construction inspected; no claim of formalizing the global spanning and saturation theorems.

[45] Michael R. Douglas, Arthur Jaffe, Ron Nissim, Fred Rajasekaran, and Yang–Mills group. Validation of the OpenAI Yang–Mills construction: Based on lean code submitted october 28, 2026, November 2026. URL https://github.com/yang-mills-group/YM/blob/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/docs/assessment-1.pdf. Independent initial assessment dated 10 November 2026; 35 pages; accessed 13 November 2026.

[46] Frank Drewes, Annegret Habel, and Hans-Jörg Kreowski. Hyperedge replacement graph grammars. In Grzegorz Rozenberg, editor, *Handbook of Graph Grammars and Computing by Graph Transformation*, volume 1, chapter 2, pages 95–162. World Scientific, 1997. URL https://people.cs.umu.se/drewes/biblio/ps-files/hrg.pdf.

[47] Hong-Guang Duan, Valentyn I. Prokhorenko, Richard J. Cogdell, Khuram Ashraf, Amy L. Stevens, Michael Thorwart, and R. J. Dwayne Miller. Nature does not rely on long-lived electronic quantum coherence for photosynthetic energy transfer. *Proceedings of the National Academy of Sciences*, 114(32):8493–8498, 2017. doi: 10.1073/pnas.1702261114. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC5559008/.

[48] Morris Dworkin. Recommendation for block cipher modes of operation: Galois/counter mode (GCM) and GMAC. Technical Report SP 800-38D, National Institute of Standards and Technology, November 2007. URL https://csrc.nist.gov/pubs/sp/800/38/d/final.

[49] Luis Escauriaza, Gregory A. Seregin, and Vladimír Šverák. $l_{3,\infty}$-solutions of the Navier–Stokes equations and backward uniqueness. *Russian Mathematical Surveys*, 58(2):211–250, 2003. doi: 10.1070/RM2003v058n02ABEH000609. URL https://pdmi.ras.ru/~seregin/Recent%20Publications/engESS3.pdf.

[50] Charbel Farhat. Model reduction: AA216/CME345 course notes. Stanford University, Farhat Research Group, n.d. URL https://web.stanford.edu/group/frg/course_work/CME345/. All nine currently available CA-AA216-CME345-Ch1–Ch9 decks, 627 PDF pages including overlays, reviewed on 13 November 2026. Original CA-CME345 filenames linked by the 2021 video returned 404; the downloaded files are replacement versions.

[51] Charles L. Fefferman. Existence and smoothness of the Navier–Stokes equation. Clay Mathematics Institute, official Millennium Prize Problem statement and errata, n.d. URL https://www.claymath.org/wp-content/uploads/2022/06/navierstokes.pdf. Six-page file retrieved 13 November 2026.

[52] Keisuke Fujii and Kohei Nakajima. Harnessing disordered-ensemble quantum dynamics for machine learning. *Physical Review Applied*, 8(2):024030, 2017. doi: 10.1103/PhysRevApplied.8.024030. URL https://arxiv.org/abs/1602.08159.

[53] Isabelle Gallagher, Laure Saint-Raymond, and Benjamin Texier. From Newton to Boltzmann: Hard spheres and short-range potentials, 2013. URL https://arxiv.org/abs/1208.5753. Retrieved version 2, 15 January 2013. Abstract, introductory printed pages 5–23, the main convergence statement and discussion on pages 49–50, and concluding pages 119–120 reviewed; not a full review of all 134 PDF pages.

[54] François Golse and Laure Saint-Raymond. The incompressible Navier–Stokes limit of the Boltzmann equation for hard cutoff potentials, 2009. URL https://arxiv.org/abs/0808.0039. Retrieved version 2, 24 April 2009. Pages 1–15 reviewed, including Theorem 2.4 and the proof outline; not a full review of the 56-page proof.

[55] Nicolae Gonciulea and Venkatramani Lakshmibai. Schubert varieties, toric varieties and ladder determinantal varieties. *Annales de l'Institut Fourier*, 47(4):1013–1064, 1997. doi: 10.5802/aif.1590. URL https://www.numdam.org/item/AIF_1997__47_4_1013_0/. Introduction and Sections 6 and 10 read; used for the distinction between a normal toric limit and the Stanley–Reisner limit. The two-row semigroup and the explicit flat families in this monograph have independent proofs.

[56] Nikita Gourianov, Michael Lubasch, Sergey Dolgov, Quincy Y. van den Berg, Hessam Babaee, Peyman Givi, Martin Kiffner, and Dieter Jaksch. A quantum-inspired approach to exploit turbulence structures. *Nature Computational Science*, 2:30–37, 2022. doi: 10.1038/s43588-021-00181-1. URL https://arxiv.org/abs/2106.05782v3.

[57] David J. Gross and Frank Wilczek. Ultraviolet behavior of non-abelian gauge theories. *Physical Review Letters*, 30:1343–1346, 1973. doi: 10.1103/PhysRevLett.30.1343. URL https://doi.org/10.1103/PhysRevLett.30.1343.

[58] Alfredo Guevara, Alexandru Lupsasca, David Skinner, Andrew Strominger, and Kevin Weil. Single-minus gluon tree amplitudes are nonzero, 2026. URL https://arxiv.org/abs/2602.12176v2. Version 1: 12 February 2026; version 2: 9 March 2026.

[59] Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. Colorless green recurrent networks dream hierarchically. In *Proceedings of NAACL-HLT, Volume 1*, pages 1195–1205, 2018. doi: 10.18653/v1/N18-1108. URL https://aclanthology.org/N18-1108/.

[60] V. G. Gurzadyan and Roger Penrose. Concentric circles in WMAP data may provide evidence of violent pre-Big-Bang activity, 2010. URL https://arxiv.org/abs/1011.3706v1. Version 1, 16 November 2010. Full eight-page preprint read for this chapter.

[61] V. G. Gurzadyan and Roger Penrose. CCC and the Fermi paradox. *European Physical Journal Plus*, 131:11, 2016. doi: 10.1140/epjp/i2016-16011-1. URL https://arxiv.org/abs/1512.00554v2. Reviewed arXiv version 2, 14 December 2015, five pages.

[62] Michael Hahn. Theoretical limitations of self-attention in neural sequence models. *Transactions of the Association for Computational Linguistics*, 8:156–171, 2020. doi: 10.1162/tacl_a_00306. URL https://aclanthology.org/2020.tacl-1.11/.

[63] Nicholas C. Harris, Gregory R. Steinbrecher, Mihika Prabhu, Yoav Lahini, Jacob Mower, Darius Bunandar, Changchen Chen, Franco N. C. Wong, Tom Baehr-Jones, Michael Hochberg, Seth Lloyd, and Dirk Englund. Quantum transport simulations in a programmable nanophotonic processor. *Nature Photonics*, 11:447–452, 2017. doi: 10.1038/nphoton.2017.95. URL https://arxiv.org/abs/1507.03406.

[64] Baran Hashemi, Kurt Pasque, Chris Teska, and Ruriko Yoshida. Tropical attention: Neural algorithmic reasoning for combinatorial algorithms, 2025. URL https://arxiv.org/abs/2505.17190v2. Version 2, 23 October 2025; definitions, Appendix C proof of Lemma 3.1, Lemma C.1 and Theorem C.5 inspected. Localized corrections are proved in the present amendment; the full proof package is not imported.

[65] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830. URL https://www.cs.rpi.edu/academics/courses/spring06/random/hoefding.pdf.

[66] Tim Hunter. Kallini et al. (2024) do not compare impossible languages with constituency-based ones. *Computational Linguistics*, 51(2), 2025. doi: 10.1162/coli_a_00554. URL https://aclanthology.org/2025.cl-2.7/.

[67] J. Itatani, J. Levesque, D. Zeidler, H. Niikura, H. Pépin, J. C. Kieffer, P. B. Corkum, and D. M. Villeneuve. Tomographic imaging of molecular orbitals. *Nature*, 432(7019):867–871, 2004. doi: 10.1038/nature03183. URL https://www.attoscience.ca/pdf/Itatani_Nature_2004.pdf. Full five-page main article reviewed from the authors' laboratory copy.

[68] Arthur Jaffe and Edward Witten. Quantum Yang–Mills theory. In James Carlson, Arthur Jaffe, and Andrew Wiles, editors, *The Millennium Prize Problems*. American Mathematical Society and Clay Mathematics Institute, 2006. URL https://www.claymath.org/wp-content/uploads/2022/06/yangmills.pdf. Official problem statement; reviewed in full, 14 PDF pages.

[69] David Jennings, Kamil Korzekwa, Matteo Lostaglio, Richard Ashworth, Emanuele Marsili, and Stephen Rolston. An end-to-end quantum algorithm for nonlinear fluid dynamics with bounded quantum advantage. *PRX Quantum*, 7:033060, 2026. doi: 10.1103/xysy-q3fp. URL https://arxiv.org/abs/2512.03758v2. Published version, arXiv version 2, 25 September 2026, 57 pages.

[70] Jiaqi Jiang and Jonathan A. Fan. Global optimization of dielectric metasurfaces using a physics-driven neural network. *Nano Letters*, 19(8):5366–5372, 2019. doi: 10.1021/acs.nanolett.9b01857. URL https://fanlab.stanford.edu/wp-content/papercite-data/pdf/jiang2019global.pdf.

[71] Claus Jönsson. Elektroneninterferenzen an mehreren künstlich hergestellten feinspalten. *Zeitschrift für Physik*, 161(4):454–474, 1961. doi: 10.1007/BF01342460. URL https://link.springer.com/article/10.1007/BF01342460. Publisher abstract and metadata reviewed; full German original unavailable in this audit. The 1974 article is an abridged translation.

[72] Claus Jönsson. Electron diffraction at multiple slits. *American Journal of Physics*, 42(1):4–11, 1974. doi: 10.1119/1.1987592. URL https://web.pdx.edu/~pmoeck/pdf/ElectronJonsson.PDF. All eight scanned pages reviewed by OCR, with selected image checks; abridged English translation of the 1961 article.

[73] Simon Josefsson and Ilari Liusvaara. Edwards-curve digital signature algorithm (EdDSA). Technical Report RFC 8032, RFC Editor, January 2017. URL https://www.rfc-editor.org/rfc/rfc8032. Informational RFC.

[74] Dylan L. Jow and Douglas Scott. Re-evaluating evidence for Hawking points in the CMB. *Journal of Cosmology and Astroparticle Physics*, 2020(03):021, 2020. doi: 10.1088/1475-7516/2020/03/021. URL https://arxiv.org/abs/1909.09672v2. Reviewed arXiv version 2, 21 January 2020, eleven pages.

[75] Julie Kallini, Isabel Papadimitriou, Richard Futrell, Kyle Mahowald, and Christopher Potts. Mission: Impossible language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, Volume 1*, pages 14691–14714, 2024. doi: 10.18653/v1/2024.acl-long.787. URL https://aclanthology.org/2024.acl-long.787/.

[76] Jinwoo Kim, Tien Dat Nguyen, Seonwoo Min, Sungjun Cho, Moontae Lee, Honglak Lee, and Seunghoon Hong. Pure transformers are powerful graph learners, 2022. URL https://arxiv.org/abs/2207.02505v2. NeurIPS 2022; node and edge tokenization and the stated equivariant-network expressivity comparison.

[77] Jinwoo Kim, Tien Dat Nguyen, Seonwoo Min, Sungjun Cho, Moontae Lee, Honglak Lee, and Seunghoon Hong. Pure transformers are powerful graph learners. In *Advances in Neural Information Processing Systems*, 2022. URL https://arxiv.org/abs/2207.02505v2.

[78] A. Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1): 2–30, 2003. doi: 10.1016/S0003-4916(02)00018-0. URL https://arxiv.org/abs/quant-ph/9707021. Preprint submitted 9 July 1997; primary preprint Section 1 inspected.

[79] Salem Lahlou, Tristan Deleu, Pablo Lemos, Dinghuai Zhang, Alexandra Volokhova, Alex Hernández-García, Léna Néhale Ezzine, Yoshua Bengio, and Nikolay Malkin. A theory of continuous generative flow networks, 2023. URL https://arxiv.org/abs/2301.12594v2. ICML 2023; continuous and hybrid state-space flow theory, distinct from a diffusion-limit theorem.

[80] Yair Lakretz, German Kruszewski, Theo Desbordes, Dieuwke Hupkes, Stanislas Dehaene, and Marco Baroni. The emergence of number and syntax units in LSTM language models. In *Proceedings of NAACL-HLT, Volume 1*, 2019. doi: 10.18653/v1/N19-1002. URL https://aclanthology.org/N19-1002/.

[81] Yair Lakretz, Théo Desbordes, Dieuwke Hupkes, and Stanislas Dehaene. Can transformers process recursive nested constructions, like humans? In *Proceedings of the 29th International Conference on Computational Linguistics*, 2022. URL https://aclanthology.org/2022.coling-1.285/.

[82] Andrew Lampinen. Can language models handle recursively nested grammatical structures? a case study on comparing models and humans. *Computational Linguistics*, 50(4), 2024. URL https://aclanthology.org/2024.cl-4.8/.

[83] Zaki Leghtas, Steven Touzard, Ioan M. Pop, Angela Kou, Brian Vlastakis, Andrei Petrenko, Katrina M. Sliwa, Anirudh Narla, Shyam Shankar, Michael J. Hatridge, Matthew Reagor, Luigi Frunzio, Robert J. Schoelkopf, Mazyar Mirrahimi, and Michel H. Devoret. Confining the state of light to a quantum manifold by engineered two-photon loss. *Science*, 347(6224): 853–857, 2015. doi: 10.1126/science.aaa2085. URL https://arxiv.org/abs/1412.4633.

[84] Martin Leslie. *Hypermap-Homology Quantum Codes*. PhD thesis, University of Arizona, 2013. URL https://arxiv.org/abs/1310.5376. Primary abstract inspected for prior work; no new code superiority claim is based on this thesis.

[85] Göran Lindblad. On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2):119–130, 1976. doi: 10.1007/BF01608499. URL https://link.springer.com/article/10.1007/BF01608499.

[86] Grigori L. Litvinov. The Maslov dequantization, idempotent and tropical mathematics: A brief introduction, 2005. URL https://arxiv.org/abs/math/0507014. Context for positive-semiring dequantization; the signed polynomial counterexample in this chapter is proved directly.

[87] Jin-Peng Liu, Herman Øie Kolden, Hari K. Krovi, Nuno F. Loureiro, Konstantina Trivisa, and Andrew M. Childs. Efficient quantum algorithm for dissipative nonlinear differential equations. *Proceedings of the National Academy of Sciences*, 118(35):e2026805118, 2021. doi: 10.1073/pnas.2026805118. URL https://arxiv.org/abs/2011.03185v4. Reviewed source: corrected arXiv version 4, 7 July 2026, 36 pages.

[88] Rodney Loudon. One-dimensional hydrogen atom. *American Journal of Physics*, 27(9):649–655, 1959. doi: 10.1119/1.1934950. URL https://doi.org/10.1119/1.1934950. Publisher-deposited abstract and metadata reviewed; full paper unavailable in this audit.

[89] Rodney Loudon. One-dimensional hydrogen atom. *Proceedings of the Royal Society A*, 472 (2185):20150534, 2016. doi: 10.1098/rspa.2015.0534. URL https://doi.org/10.1098/rspa.2015.0534. Publisher-deposited abstract and metadata reviewed; full paper unavailable in this audit.

[90] Michael Lubasch, Jaewoo Joo, Pierre Moinier, Martin Kiffner, and Dieter Jaksch. Variational quantum algorithms for nonlinear problems. *Physical Review A*, 101:010301, 2020. doi: 10.1103/PhysRevA.101.010301. URL https://arxiv.org/abs/1907.09032v3.

[91] Han Luo, Ziyi Yang, Jingquan Luo, Ziruo Wang, Yuexin Su, Xiaoming Sun, Lvzhou Li, and Tongyang Li. Quantum algorithm for elliptic curve discrete logarithms with space-efficient point addition, 2026. URL https://arxiv.org/abs/2607.13816v3. Version 3, 5 September 2026. Resource expressions are attributed to that version; see the source ledger for exact metadata.

[92] Martin Lüscher. Construction of a selfadjoint, strictly positive transfer matrix for euclidean lattice gauge theories. *Communications in Mathematical Physics*, 54:283–292, 1977. doi: 10.1007/BF01614090. URL https://doi.org/10.1007/BF01614090.

[93] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *Advances in Cryptology—EUROCRYPT 2010*, pages 1–23, 2010. doi: 10.1007/978-3-642-13190-5_1. URL https://cims.nyu.edu/~regev/papers/ideal-lwe.pdf. Author-hosted full version reviewed selectively; search and decision statements have distinct hypotheses.

[94] Oliver Markwell and Chris Stevens. Toward fixing a framework for conformal cyclic cosmology. *General Relativity and Gravitation*, 2023. doi: 10.1007/s10714-023-03142-3. URL https://arxiv.org/abs/2212.06914v2. Reviewed arXiv version 2, 16 December 2022, fourteen pages. Author order follows the reviewed PDF; arXiv abstract metadata lists the reverse order.

[95] Samir D. Mathur. Quantum mechanics: Introductory online modules. Ohio State University online teaching material, n.d. URL https://www.asc.ohio-state.edu/mathur.16/quantummechanics27-11-17/navigation.html. All 22 available text sections of modules 1–4 reviewed on 13 November 2026; later modules were marked under construction. Interactive animations were not executed.

[96] Krzysztof A. Meissner and Roger Penrose. The physics of Conformal Cyclic Cosmology, 2025. URL https://arxiv.org/abs/2503.24263v1. Version 1, 31 March 2025, twenty pages. Full text read; proposal and its stated hypotheses distinguished from observational confirmation.

[97] Zhaoyuan Meng and Yue Yang. Quantum spin representation for the Navier–Stokes equation. *Physical Review Research*, 6:043130, 2024. doi: 10.1103/PhysRevResearch.6.043130. URL https://journals.aps.org/prresearch/pdf/10.1103/PhysRevResearch.6.043130. Published 13 November 2024. Selected model equations and stated global limitations reviewed.

[98] William Merrill and Ashish Sabharwal. The expressive power of transformers with chain of thought. In *International Conference on Learning Representations*, 2024. URL https://arxiv.org/abs/2310.07923. arXiv:2310.07923v5, 11 April 2024.

[99] Hazime Mori. Transport, collective motion, and brownian motion. *Progress of Theoretical Physics*, 33(3):423–455, 1965. doi: 10.1143/PTP.33.423. URL https://academic.oup.com/ptp/article/33/3/423/1925580.

[100] Adam Moss, Douglas Scott, and James P. Zibin. No evidence for anomalously low variance circles on the sky. *Journal of Cosmology and Astroparticle Physics*, 2011(04):033, 2011. doi: 10.1088/1475-7516/2011/04/033. URL https://arxiv.org/abs/1012.1305v3. Reviewed arXiv version 3, 7 April 2011, four pages.

[101] Kohei Nakajima, Helmut Hauser, Tao Li, and Rolf Pfeifer. Information processing via physical soft body. *Scientific Reports*, 5:10487, 2015. doi: 10.1038/srep10487. URL https://www.nature.com/articles/srep10487.

[102] Wataru Namiki, Daiki Nishioka, Yu Yamaguchi, Takashi Tsuchiya, Tohru Higuchi, and Kazuya Terabe. Experimental demonstration of high-performance physical reservoir computing with nonlinear interfered spin wave multidetection. *Advanced Intelligent Systems*, 5(12):2300228, 2023. doi: 10.1002/aisy.202300228. URL https://advanced.onlinelibrary.wiley.com/doi/full/10.1002/aisy.202300228.

[103] Mudumbai S. Narasimhan and C. S. Seshadri. Stable and unitary vector bundles on a compact riemann surface. *Annals of Mathematics*, 82(3):540–567, 1965. doi: 10.2307/1970710. URL https://doi.org/10.2307/1970710.

[104] National Institute of Standards and Technology. Module-lattice-based key-encapsulation mechanism standard. Technical Report FIPS 203, National Institute of Standards and Technology, August 2024. URL https://csrc.nist.gov/pubs/fips/203/final. Primary publication record and current update notice checked 13 November 2026.

[105] National Institute of Standards and Technology. Module-lattice-based key-encapsulation mechanism standard. Technical Report FIPS 203, U.S. Department of Commerce, 2024. URL https://csrc.nist.gov/pubs/fips/203/final. Final 13 August 2024; publication page and 17 November 2025 errata notice checked. This monograph does not implement ML-KEM.

[106] Yoav Nir and Adam Langley. ChaCha20 and Poly1305 for IETF protocols. Technical Report RFC 8439, RFC Editor, June 2018. URL https://www.rfc-editor.org/rfc/rfc8439. Informational RFC.

[107] NIST Digital Library of Mathematical Functions. Chapter 25: Zeta and related functions, n.d. URL https://dlmf.nist.gov/25. Standard normalization and analytic continuation conventions; accessed 13 November 2026.

[108] OpenAI. NavierStokesAndEuler: Lean certificates accompanying Navier–Stokes and Euler results, 2026. URL https://github.com/openai/NavierStokesAndEuler/tree/f9e8bc5b38b6e212696e8a30e3e91517af887bbd. Revision f9e8bc5b38b6e212696e8a30e3e91517af887bbd; Apache-2.0; Lean 4.34.0-rc2.

[109] OpenAI. Finite time blowup for Navier–Stokes. Research manuscript, 166 pages, September 2026. URL https://cdn.openai.com/pdf/32d9f210-8b73-45e0-91bc-82a30aef8a9a/navier-stokes.pdf. Retrieved 13 November 2026; SHA-256 0e779481c4da40bd28d1e642e1d8ca57447d129610df28dfa5a11e9af8ae228f.

[110] OpenAI. On the Navier–Stokes millennium prize problem, September 2026. URL https://openai.com/index/navier-stokes-solution/. 8 September 2026; retrieved 13 November 2026.

[111] OpenAI. Yang–Mills public joint specification: ChallengeSpec.lean, 2026. URL https://github.com/yang-mills-group/OpenAI-YM-20261028/blob/3456e5e75cf8d4cd558955a5c7638fba2d68f464/ChallengeSpec.lean. Pinned 28 October submission; read in full, 190 lines.

[112] OpenAI. Working with academic communities to build scientific understanding, October 2026. URL https://openai.com/index/building-scientific-understanding-with-academic-communities/. Announcement dated 26 October 2026; accessed 13 November 2026.

[113] OpenAI. Yang–Mills public endpoint: Solution.lean, 2026. URL https://github.com/yang-mills-group/OpenAI-YM-20261028/blob/3456e5e75cf8d4cd558955a5c7638fba2d68f464/Solution.lean. Pinned formal proof endpoint; external proof, not rebuilt in this project.

[114] OpenAI. Public gauge group and wightman definitions: PublicTheory.lean, 2026. URL https://github.com/yang-mills-group/OpenAI-YM-20261028/blob/3456e5e75cf8d4cd558955a5c7638fba2d68f464/PublicTheory.lean. Pinned 28 October submission; read in full, 325 lines.

[115] OpenAI. Public wilson lattice and ordered limit definitions: PublicWilson.lean, 2026. URL https://github.com/yang-mills-group/OpenAI-YM-20261028/blob/3456e5e75cf8d4cd558955a5c7638fba2d68f464/PublicWilson.lean. Pinned 28 October submission; read in full, 92 lines.

[116] Konrad Osterwalder and Robert Schrader. Axioms for euclidean green's functions. *Communications in Mathematical Physics*, 31:83–112, 1973. doi: 10.1007/BF01645738. URL https://doi.org/10.1007/BF01645738.

[117] Konrad Osterwalder and Robert Schrader. Axioms for euclidean green's functions II. *Communications in Mathematical Physics*, 42:281–305, 1975. doi: 10.1007/BF01608978. URL https://doi.org/10.1007/BF01608978.

[118] Konrad Osterwalder and Erhard Seiler. Gauge field theories on a lattice. *Annals of Physics*, 110(2):440–471, 1978. doi: 10.1016/0003-4916(78)90039-8. URL https://doi.org/10.1016/0003-4916(78)90039-8.

[119] J. G. Parker. Rotational and vibrational relaxation in diatomic gases. *The Physics of Fluids*, 2(4):449–462, 1959. doi: 10.1063/1.1724417. URL https://doi.org/10.1063/1.1724417. Publisher-deposited abstract and metadata reviewed; full text unavailable in this audit.

[120] Jorge Pérez, Pablo Barceló, and Javier Marinkovic. Attention is Turing-complete. *Journal of Machine Learning Research*, 22(75):1–35, 2021. URL https://jmlr.org/papers/v22/20-302.html.

[121] Cédric Pilatte. Unconditional correctness of recent quantum algorithms for factoring and computing discrete logarithms. *Forum of Mathematics, Pi*, 14:e5, 2026. doi: 10.1017/fmp.2025.10023. URL https://arxiv.org/abs/2404.16450. Published 9 February 2026; preprint version 3 dated 24 August 2026 also checked.

[122] Planck Collaboration. Planck 2018 results. VII. isotropy and statistics of the CMB. *Astronomy & Astrophysics*, 641:A7, 2020. doi: 10.1051/0004-6361/201935201. URL https://arxiv.org/abs/1906.02552v2. Selected introductory material and section 8 read in arXiv version 2, 14 September 2020; not a full sixty-six-page review.

[123] John Proos and Christof Zalka. Shor's discrete logarithm quantum algorithm for elliptic curves. *Quantum Information and Computation*, 3(4):317–344, 2003. URL https://arxiv.org/abs/quant-ph/0301141.

[124] Joan Puigcerver, Carlos Riquelme, Basil Mustafa, and Neil Houlsby. From sparse to soft mixtures of experts. In *International Conference on Learning Representations*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/79fea214543ba263952ac3f4e5452b14-Paper-Conference.pdf.

[125] Python Cryptographic Authority. Authenticated encryption: cryptography 50.0.0 documentation, 2026. URL https://cryptography.io/en/50.0.0/hazmat/primitives/aead/. Versioned primary API documentation, accessed 13 November 2026.

[126] Python Cryptographic Authority. Ed25519 signing: cryptography 50.0.0 documentation, 2026. URL https://cryptography.io/en/50.0.0/hazmat/primitives/asymmetric/ed25519/. Versioned primary API documentation, accessed 13 November 2026.

[127] Python Cryptographic Authority. ML-KEM key encapsulation: cryptography 50.0.0 documentation, 2026. URL https://cryptography.io/en/50.0.0/hazmat/primitives/asymmetric/mlkem/. Versioned primary API documentation, accessed 13 November 2026.

[128] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):34:1–34:40, 2009. doi: 10.1145/1568318.1568324. URL https://cims.nyu.edu/~regev/papers/qcrypto.pdf.

[129] Oded Regev. An efficient quantum factoring algorithm, 2023. URL https://arxiv.org/abs/2308.06572. Reviewed version 3.

[130] Martin Roetteler, Michael Naehrig, Krysta M. Svore, and Kristin Lauter. Quantum resource estimates for computing elliptic curve discrete logarithms. In *Advances in Cryptology—ASIACRYPT 2017*, 2017. URL https://arxiv.org/abs/1706.06752. Reviewed preprint version 3.

[131] Klaus W. Roggenkamp. On the irreducible lattices of orders. *Canadian Journal of Mathematics*, 21:970–976, 1969. doi: 10.4153/CJM-1969-106-x. URL https://www.cambridge.org/core/services/aop-cambridge-core/content/view/25660C441608D384D2564A931E2D692D/S0008414X00042346a.pdf/on_the_irreducible_lattices_of_orders.pdf.

[132] Klaus W. Roggenkamp. A remark on separable orders. *Canadian Mathematical Bulletin*, 12 (4):453–455, 1969. doi: 10.4153/CMB-1969-055-4. URL https://www.cambridge.org/core/services/aop-cambridge-core/content/view/54C0AFAE48BBD8E21854E750B6EF49BE/S000843950005387Xa.pdf/a-remark-on-separable-orders.pdf.

[133] Klaus W. Roggenkamp. *Lattices over Orders II*, volume 142 of *Lecture Notes in Mathematics*. Springer, 1970. doi: 10.1007/BFb0058955. User-supplied scanned volume; printed pages 1–387. PDF pages 6–392 correspond to printed pages 1–387.

[134] Klaus W. Roggenkamp. Grothendieck groups of bass orders. *Canadian Journal of Mathematics*, 23(1):103–115, 1971. doi: 10.4153/CJM-1971-011-9.

[135] Klaus W. Roggenkamp and Verena Huber-Dyson. *Lattices over Orders I*, volume 115 of *Lecture Notes in Mathematics*. Springer, 1970. doi: 10.1007/BFb0068796. URL https://link.springer.com/book/10.1007/BFb0068796. User-supplied scanned copy: 309 PDF pages, xix + 290 printed pages.

[136] Pradeep Sarvepalli. Relation between surface codes and hypermap-homology quantum codes. *Physical Review A*, 89:052316, 2014. doi: 10.1103/PhysRevA.89.052316. URL https://arxiv.org/abs/1312.6344.

[137] Amelie Schreiber. GRH-research: source repository and linked research program. GitHub repository, frozen commit dbd954eede479a5129831e4e7c30537389ae973d, 2026. URL https://github.com/amelie-iska/GRH-research/tree/dbd954eede479a5129831e4e7c30537389ae973d. Accessed 13 November 2026. Historical and provisional claims retain their original status.

[138] Amelie Schreiber. Surface-order determinants, monomial Artin L-functions, and the general Hilbert–Pólya extension problem. Research manuscript in the GRH repository, 2026. URL https://github.com/amelie-iska/GRH/tree/0c77122d3cff5b737fc56730f54045ef185e3df1. Pinned linked source; 223-page PDF at the reviewed snapshot. Research status is assessed statement by statement.

[139] Amelie Schreiber. Emergent cyclic time and the Hilbert–Pólya research program. Surface-Orders-Hilbert-Polya research repository, 2026. URL https://github.com/amelie-iska/Surface-Orders-Hilbert-Polya/tree/212ff730cdfe5d9ab2dee17c11c7f83061fa0bf2. Descriptive citation to the pinned repository and its 265-page compiled manuscript; not a claim that GRH is proved.

[140] Amelie Schreiber. ToricGT III: Sheafified MCP memory, graph-structured vector retrieval, and long-context tropical ring attention, 2026. URL https://github.com/

amelie-iska/ToricGT-Program/blob/f91ba1811b5a4f70d53c7120acf8ca2a0049f85b/ToricGT-III/assets/toricgt_iii_long_context_tropical_mcp_memory.tex. Pinned primary manuscript; prefix-safe information and the final exact-context utility decomposition inspected directly.

[141] Amelie Schreiber. ToricGT IV: Oracle trajectory classifiers for behavior, reasoning quality, and bits-per-byte control, 2026. URL https://github.com/amelie-iska/ToricGT-Program/blob/f91ba1811b5a4f70d53c7120acf8ca2a0049f85b/ToricGT-IV/assets/toricgt_iv_oracle_trajectory_classifiers.tex. Pinned primary manuscript; proper-scoring endpoint hypotheses and selected theory statements inspected directly.

[142] Amelie Schreiber. Composable tropical quivers for learned operators: Transformer-general architectures, ergodic embedding flows, and hybrid GFlowNet search, 2026. URL https://github.com/amelie-iska/Tropical_Quivers_of_Archs/blob/2b6ab79a685e2c6233fb8892e59589951feeddad/tropical_quiver_research_program.tex. Pinned primary manuscript; selected transport, contraction, universality and stochastic-flow claims audited, 13 November 2026.

[143] Amelie Schreiber. ToricBLM: Thought alcoves, tropical attention, and small models with serious geometry, June 2026. URL https://huggingface.co/blog/AmelieSchreiber/toricblm. Author's proposal, published June 19, 2026; selected conceptual sections reviewed, with no independent trained model reproduction claimed.

[144] Amelie Schreiber. Toric algebraic training objectives for ToricGT: Combinatorial commutative algebra, reasoning trajectories, and bits-per-byte compression, June 2026. URL https://github.com/amelie-iska/ToricGT-Program/blob/f91ba1811b5a4f70d53c7120acf8ca2a0049f85b/ToricGT-V/assets/toricgt_toric_algebra_training_paper.tex. Selected sections reviewed directly in the pinned primary TeX source; localized corrections are recorded in this monograph's ToricGT review.

[145] Amelie Schreiber. ToricGT: graph tokenizer, tropical attention, and soft expert source modules, 2026. URL https://github.com/amelie-iska/ToricGT/tree/a5199bc989a6fb170e21220e5e87bf95bc2fc0b4/src/toricgt. Reviewed at commit a5199bc989a6fb170e21220e5e87bf95bc2fc0b4; module-level content review, not certification of all repository claims.

[146] Amelie Schreiber. ToricGT Research Program: program overview and evaluation discipline, 2026. URL https://github.com/amelie-iska/ToricGT-Program/blob/f91ba1811b5a4f70d53c7120acf8ca2a0049f85b/README.md. Reviewed at commit f91ba1811b5a4f70d53c7120acf8ca2a0049f85b.

[147] Amelie Schreiber. ToricGT: finite weyl matrix implementation, 2026. URL https://github.com/amelie-iska/ToricGT/blob/a5199bc989a6fb170e21220e5e87bf95bc2fc0b4/src/toricgt/toric_algebra.py. Pinned source module, with explicit clock and backward-shift conventions.

[148] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. Evaluating analytic gradients on quantum hardware. *Physical Review A*, 99:032331, 2019. doi: 10.1103/PhysRevA.99.032331. URL https://arxiv.org/abs/1811.11184. Published 21 March 2019; the generator and parameter-occurrence assumptions are retained.

[149] Peter Selinger. A survey of graphical languages for monoidal categories. In *New Structures for Physics*, volume 813 of *Lecture Notes in Physics*, pages 289–355. Springer, 2011. doi: 10.1007/978-3-642-12821-9_4. URL https://arxiv.org/abs/0908.3347.

[150] Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x. URL https://web.mit.edu/6.441/spring05/reading/shannon1948.pdf. Foundational attribution; bibliographic metadata and opening material checked, not a full source review. Finite information results used here are proved in the chapter.

[151] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997. doi: 10.1137/S0097539795293172. URL https://arxiv.org/abs/quant-ph/9508027.

[152] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, 2 edition, 2009. doi: 10.1007/978-0-387-09494-6.

[153] Simons Observatory Collaboration. The Simons Observatory: Science goals and forecasts. *Journal of Cosmology and Astroparticle Physics*, 2019(02):056, 2019. doi: 10.1088/1475-7516/2019/02/056. URL https://arxiv.org/abs/1808.07445v2. Selected section 2.6 read in version 2, 1 March 2019; cited for forecasting methodology, not present operating status or a CCC-specific forecast.

[154] Andrew Steane. Multiple-particle interference and quantum error correction. *Proceedings of the Royal Society A*, 452:2551–2577, 1996. doi: 10.1098/rspa.1996.0136. URL https://arxiv.org/abs/quant-ph/9601029.

[155] A. S. Stodolna, A. Rouzée, F. Lépine, S. Cohen, F. Robicheaux, A. Gijsbertsen, J. H. Jungmann, C. Bordas, and M. J. J. Vrakking. Hydrogen atoms under magnification: Direct observation of the nodal structure of Stark states. *Physical Review Letters*, 110(21):213001, 2013. doi: 10.1103/PhysRevLett.110.213001. URL https://www.physics.purdue.edu/~robichf/papers/prl110.213001.pdf. Full five-page main article reviewed; supplementary material not reviewed.

[156] Ye Su and Yong Liu. Geometric capacity of transformers: A tropical geometry perspective, 2026. URL https://arxiv.org/abs/2604.14727v2. Version 2, 21 July 2026; primary abstract and version metadata checked for the fixed-key, zero-temperature and finite-temperature distinctions.

[157] Andrew V. Sutherland. 18.783 elliptic curves: Lecture 17, the CM torsor. MIT lecture notes, 2023. URL https://math.mit.edu/classes/18.783/2023/LectureNotes17.pdf. November 9, 2023.

[158] Andrew V. Sutherland. 18.783 elliptic curves: Lecture 12, endomorphism algebras. MIT lecture notes, 2023. URL https://math.mit.edu/classes/18.783/2023/LectureNotes12.pdf. October 24, 2023.

[159] Andrew V. Sutherland. 18.783 elliptic curves: Lecture 7, point counting. MIT lecture notes, 2023. URL https://math.mit.edu/classes/18.783/2023/LectureNotes7.pdf. October 3, 2023.

[160] Andrew V. Sutherland. 18.783 elliptic curves: Lecture 18, riemann surfaces and modular curves. MIT lecture notes, 2023. URL https://math.mit.edu/classes/18.783/2023/LectureNotes18.pdf. November 14, 2023.

[161] Andrew V. Sutherland. 18.783 elliptic curves: Lecture 13, ordinary and supersingular elliptic curves. MIT lecture notes, 2023. URL https://math.mit.edu/classes/18.783/2023/LectureNotes13.pdf. October 26, 2023.

[162] Richard G. Swan. Vector bundles and projective modules. *Transactions of the American Mathematical Society*, 105:264–277, 1962. URL https://www.math.unl.edu/~jkettinger2/swan.pdf.

[163] Terence Tao. Finite time blowup for an averaged three-dimensional Navier–Stokes equation. *Journal of the American Mathematical Society*, 29(3):601–674, 2016. doi: 10.1090/jams/838. URL https://doi.org/10.1090/jams/838.

[164] John Tate. Endomorphisms of abelian varieties over finite fields. *Inventiones Mathematicae*, 2:134–144, 1966. doi: 10.1007/BF01404549. URL https://pazuki.perso.math.cnrs.fr/index_fichiers/Tate66.pdf.

[165] Jan Tempus, Philip Whittington, Craig W. Schmidt, Dennis Komm, and Tiago Pimentel. Tokenisation via convex relaxations, 2026. URL https://arxiv.org/abs/2605.22821v1. Primary token-DAG and colored-flow formulation; the companion implements an original bounded example.

[166] Paul Tod. Conformal methods in general relativity with application to Conformal Cyclic Cosmology: a minicourse given at the IXth IMLG, Warsaw 2018, 2021. URL https://arxiv.org/abs/2102.02701v1. Selected sections 1.2–1.2.4, 1.3.1, 3.6 and 3.7 read; not a full sixty-page review.

[167] Paolo Valentini, Thomas E. Schwartzentruber, Jason D. Bender, Ioannis Nompelis, and Graham V. Candler. Direct molecular simulation of nitrogen dissociation based on an ab initio potential energy surface. *Physics of Fluids*, 27(8):086102, 2015. doi: 10.1063/1.4929394. URL https://experts.umn.edu/en/publications/direct-molecular-simulation-of-nitrogen-dissociation-based-on-an-/. University-hosted and publisher-deposited abstracts and metadata reviewed; full paper unavailable in this audit.

[168] Leon van Dommelen. Quantum mechanics for engineers. Online textbook, Florida A&M University–Florida State University, n.d. URL https://web1.eng.famu.fsu.edu/~dommelen/quantum/style_a/. Current online sections 3.1, 3.2, 3.4, 4.6, 5.6, 5.7, 5.9, 6.14, 9.2 and derivation D.51 reviewed on 13 November 2026; not a review of the entire book.

[169] John Voight. Supersingular elliptic curves. In *Quaternion Algebras*, volume 288 of *Graduate Texts in Mathematics*, pages 783–797. Springer, 2021. doi: 10.1007/978-3-030-56694-4_42. URL https://link.springer.com/chapter/10.1007/978-3-030-56694-4_42.

[170] Alexey Volkov. Elementary kinetic theory of gases. ME591/491 Non-equilibrium Gas Dynamics, Chapter 1, University of Alabama, 2017. URL https://volkov.eng.ua.edu/ME591_491_NEGD/2017-Spring-NEGD-01-ElemKineticTheory.pdf. Readable extracted text of all 88 pages reviewed; formula-font extraction limitations recorded.

[171] Alexey Volkov. Molecular models. ME591/491 Non-equilibrium Gas Dynamics, Chapter 2, University of Alabama, 2017. URL https://volkov.eng.ua.edu/ME591_491_NEGD/2017-Spring-NEGD-02-MolecularModels.pdf. Readable extracted text of all 46 pages reviewed; selected formula pages visually checked.

[172] Alexey Volkov. Kinetic theory of dilute gases. ME591/491 Non-equilibrium Gas Dynamics, Chapter 3, University of Alabama, 2017. URL https://volkov.eng.ua.edu/ME591_491_NEGD/2017-Spring-NEGD-03-KineticTheory.pdf. Readable extracted text of all 100 pages reviewed; selected entropy and closure formula pages visually checked.

[173] Alexey Volkov. Elements of statistics. ME591/491 Non-equilibrium Gas Dynamics, Chapter 4, University of Alabama, 2017. URL https://volkov.eng.ua.edu/ME591_491_NEGD/2017-Spring-NEGD-04-Statistics.pdf. Readable extracted text of all 34 pages reviewed; formula-font extraction limitations recorded.

[174] Alexey Volkov. Monte carlo method. ME591/491 Non-equilibrium Gas Dynamics, Chapter 5, University of Alabama, 2017. URL https://volkov.eng.ua.edu/ME591_491_NEGD/2017-Spring-NEGD-05-MonteCarloMethod.pdf. Readable extracted text of all 43 pages reviewed, including illustrative code.

[175] Alexey Volkov. Direct simulation monte carlo method. ME591/491 Non-equilibrium Gas Dynamics, Chapter 6, University of Alabama, 2017. URL https://volkov.eng.ua.edu/ME591_491_NEGD/2017-Spring-NEGD-06-DSMC.pdf. Readable extracted text of all 60 pages reviewed; selected flux and collision-sampling formulas visually checked. Illustrative source code was reviewed, not executed.

[176] C. Vozzi, M. Negro, F. Calegari, G. Sansone, M. Nisoli, S. De Silvestri, and S. Stagira. Generalized molecular orbital tomography. *Nature Physics*, 7(10):822–826, 2011. doi: 10.1038/nphys2029. URL https://www.dg.dk/filer/Nyheder/2011/dyre_naturephysics_okt.pdf. Full five-page main article reviewed; supplementary material not reviewed.

[177] Kenneth G. Wilson. Confinement of quarks. *Physical Review D*, 10:2445–2459, 1974. doi: 10.1103/PhysRevD.10.2445. URL https://doi.org/10.1103/PhysRevD.10.2445.

[178] William K. Wootters and Wojciech H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982. doi: 10.1038/299802a0. URL https://www.nature.com/articles/299802a0. Foundational attribution; publisher metadata and abstract checked. The trace-distance proof used here is supplied in full.

[179] Hsuan-Cheng Wu, Jingyao Wang, and Xiantao Li. Quantum algorithms for nonlinear dynamics: Revisiting Carleman linearization with no dissipative conditions. arXiv:2405.12714v2, 2025. URL https://arxiv.org/abs/2405.12714v2. Version 2, 31 January 2025; the exact nonresonance and bounded-solution assumptions are discussed in the text.

[180] Kevin Xia, Kai-Zhan Lee, Yoshua Bengio, and Elias Bareinboim. The causal-neural connection: Expressiveness, learnability, and inference. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://arxiv.org/abs/2107.00793. Reviewed arXiv:2107.00793v3.

[181] Chen Ning Yang and Robert L. Mills. Conservation of isotopic spin and isotopic gauge invariance. *Physical Review*, 96:191–195, 1954. doi: 10.1103/PhysRev.96.191. URL https://doi.org/10.1103/PhysRev.96.191.

[182] Yang–Mills group. Combined acceptance specification: AcceptanceSpec.lean, 2026. URL https://github.com/yang-mills-group/YM/blob/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/submissions/openai-20261028/checks/AcceptanceSpec.lean. Read in full, 178 lines; independent target retaining one construction witness.

[183] Yang–Mills group. Historical clean-build and axiom-audit evidence, 2026. URL https://github.com/yang-mills-group/YM/tree/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/submissions/openai-20261028/evidence/historical/build. Historical upstream records, not a fresh build by this project.

[184] Yang–Mills group. Proof of combined acceptance: CombinedAcceptance.lean, 2026. URL https://github.com/yang-mills-group/YM/blob/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/submissions/openai-20261028/checks/CombinedAcceptance.lean. Read in full, 122 lines; external verified result.

[185] Yang–Mills group. Historical combined-acceptance comparator evidence and controls, 2026. URL https://github.com/yang-mills-group/YM/tree/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/submissions/openai-20261028/evidence/historical/acceptance-comparator. Built-in kernel checks; recorded macOS execution did not provide Linux isolation.

[186] Yang–Mills group. Worked leading trace check: LeadingTrace20261028.lean, 2026. URL https://github.com/yang-mills-group/YM/blob/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/submissions/openai-20261028/docs/examples/LeadingTrace-20261028.lean. Commented source explaining the fixed normalization and scope.

[187] Yang–Mills group. Independent operator-product statement check: OPECheck.lean, 2026. URL https://github.com/yang-mills-group/YM/blob/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/submissions/openai-20261028/checks/OPECheck.lean. Check of field-degree asymptotics, not all-orders perturbative coefficient matching.

[188] Yang–Mills group. Scalar osterwalder–schrader interface check: CompleteOS.lean, 2026. URL https://github.com/yang-mills-group/YM/blob/55ba9fc36293b3f4d52c6b1f813169a0be76ee98/submissions/openai-20261028/checks/CompleteOS.lean. Selected centered trace; diagonal-flat Schwartz test domain.

[189] Chulhee Yun, Srinadh Bhojanapalli, Ankit Singh Rawat, Sashank J. Reddi, and Sanjiv Kumar. Are transformers universal approximators of sequence-to-sequence functions?, 2020. URL https://arxiv.org/abs/1912.10077v2. ICLR 2020; Theorems 2 and 3 and the ambient Lebesgue Lp metric inspected directly.

[190] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Zilong Wang, Aliaksandra Shysheya, Jonathan Crabbé, Shoko Ueda, Roberto Sordillo, Lixin Sun, Jake Smith, Bichlien Nguyen, Hannes Schulz, Sarah Lewis, Chin-Wei Huang, Ziheng Lu, Yichi Zhou, Han Yang, Hongxia Hao, Jielan Li, Chunlei Yang, Wenjie Li, Ryota Tomioka, and Tian Xie. A generative model for inorganic materials design. *Nature*, 639:624–632, 2025. doi: 10.1038/s41586-025-08628-5. URL https://www.nature.com/articles/s41586-025-08628-5.

[191] Elinor Zerah Harush and Yonatan Dubi. Do photosynthetic complexes use quantum coherence to increase their efficiency? Probably not. *Science Advances*, 7(8):eabc4631, 2021. doi: 10.1126/sciadv.abc4631. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC7888942/.