Fluid Dynamics, Orders and Quantum Information

Microscopic reduction, finite complexes, quantum learning and a structural reading of the 2026 Navier–Stokes release

Prepared for Amelie Schreiber

A standalone extension of the research monograph

13 November 2026

Version 1.4: operator realizations and certified control

Purpose and scope

This supplement develops the fluid-dynamics extension of *Orders, Complexes, Cryptography and Quantum Yang–Mills*. It contains the same seven fluid, reduction, quiver and operator-realization chapters and the fluid verification appendix as the complete monograph, with independent pagination. The bibliography identifies the primary sources, and the companion repository contains the source coverage, exact formal declarations, numerical programs and interactive laboratory.

The contribution is an explicit mathematical interpretation and a collection of proved interfaces. The released continuum singularity theorem is attributed to its source. The selected local Lean declarations and finite computations have their own verification records. We do not present a stronger unforced theorem, a fully formalized version of every chapter, or a trained fluid model where none was executed.

Contents

1 Microscopic information, kinetic limits, and fluid closure **1**
- 1.1 What the microscopic perspective establishes . . . 1
- 1.2 Measurement, phase, and the limits of tomography . . . 2
- 1.3 Exterior powers are a physical state constraint . . . 4
- 1.4 Born–Huang expansion before Born–Oppenheimer truncation . . . 4
- 1.5 When a quantum packet behaves approximately as a particle . . . 6
- 1.6 Molecular dynamics and conserved microscopic information . . . 7
- 1.7 Marginals, correlations, and the kinetic approximation . . . 8
- 1.8 The collision algebra and its five invariants . . . 9
- 1.9 The weak collision identity and entropy production . . . 10
- 1.10 Exact moment balances and the closure problem . . . 12
- 1.11 DSMC as a stochastic approximation, with auditable assumptions . . . 14
- 1.12 Reduced models and learned closures . . . 15
- 1.13 What passes from complexes to a continuum theorem . . . 18

2 The Navier–Stokes release: exact scope and proof interfaces **19**
- 2.1 What the released theorem says . . . 19
- 2.2 An exact algebra for residual correction . . . 20
- 2.3 Concentration, energy, and the information lost by norms . . . 21
- 2.4 Similarity coordinates as an exact change of variables . . . 22
- 2.5 A quadratic relation that controls analytic derivative loss . . . 23
- 2.6 From formal orders to actual smooth fields . . . 24
- 2.7 The finite correction cycle as a verifiable dependency graph . . . 26
- 2.8 Localization and smooth force extension . . . 27
- 2.9 Comparison, quantitative residual checks, and proof strength . . . 28
- 2.10 Viscosity, periodization, and preservation of quantifiers . . . 29
- 2.11 Formal proof provenance and what this extension adds . . . 29

3 Orders, complexes and the algebra of a fluid singularity **31**
- 3.1 What a useful reformulation must preserve . . . 31
- 3.2 The de Rham complex behind incompressible motion . . . 31
- 3.3 A quadratic order occurs in the actual construction . . . 32
- 3.4 The number-field norm supplies an analytic estimate . . . 34
- 3.5 Symmetric tensors, covariance and the stress correction . . . 34
- 3.6 Moment determinants and exact compatibility . . . 35
- 3.7 From varieties of complexes to compatible fluid discretization . . . 36
- 3.8 Gentle algebras, path relations and quantitative interfaces . . . 38

3.9 Determinants, spectra and what the reformulation adds . . . 38

4 Quantum information, exact reduction and learned fluid models 39
4.1 Information reduction is a map, not an autonomous dynamics . . . 39
4.2 Coarse-graining and the unresolved stress . . . 40
4.3 A quantum observable can have the same closure obstruction . . . 41
4.4 An explicit heat channel on a periodic surface grid . . . 41
4.5 What the Madelung transformation does and does not derive . . . 42
4.6 Carleman lifting and the symmetric algebra . . . 43
4.7 What the quantum-algorithm literature actually supplies . . . 44
4.8 Exact finite repairs for collision block encodings . . . 47
4.9 A global compatibility test for a spin representation . . . 49
4.10 A mathematically constrained quantum ToricGT extension . . . 50
4.11 The executable examples and their proof boundaries . . . 51

5 Exact reduction contracts and the information lost by tropicalization 53
5.1 The exact closure criterion . . . 53
5.2 A lower bound on every approximate closure . . . 54
5.3 Signed cancellation is invisible to orders alone . . . 55
5.4 A current analytic obstruction to removing the force . . . 55
5.5 Formal coverage and reproducibility . . . 56

6 Corrected transport, information and quantum contracts for tropical quivers 57
6.1 Source scope and a quantitative contract . . . 57
6.2 Suffix transport is the correct cyclic coboundary . . . 58
6.3 A positive Laplacian does not itself supply a complex . . . 59
6.4 What a prefix-derived certificate can add . . . 59
6.5 Damping and stochastic limits require their own hypotheses . . . 60
6.6 Quantum and gentle-algebra versions of the transport contract . . . 61
6.7 Tropical approximation is not an automatic universality transfer . . . 62
6.8 Independent repairs to the tropical-attention reference . . . 63
6.9 A nonnormal obstruction to a Jacobian power certificate . . . 64
6.10 Path multiplicity separates tropical growth from operator stability . . . 65
6.11 Weighted small-gain certification at a fixed rate . . . 66
6.12 Proper scoring must include its probability boundary . . . 69
6.13 Predictive synergy is not determined by geometric overlaps . . . 69
6.14 Implications for the research program . . . 70

7 Quantum–tropical fluid realizations: operators, control and substrates 71
7.1 What a realization must specify . . . 71
7.2 A finite operator algebra with conservative quantum transport . . . 72
7.3 Toric features, tropical routing and the information they discard . . . 73
7.4 Nonlinear dynamics on an operator state space . . . 74
7.5 Control, observability and robustness . . . 75
7.6 An exact nonlinear quantum-state feedback example . . . 77
7.7 Candidate engineered and natural substrates . . . 78
7.8 De novo substrate design with generative models . . . 80
7.9 Connections to the wider research program . . . 81

A Fluid verification, reproducibility and proof dependencies **82**
A.1 Current verification boundary 82
A.2 What was checked and what each check means 82
A.3 The separate upstream Navier–Stokes formalization 83
A.4 Independent finite computational checks 84
A.5 Reproduction commands 84
A.6 The complete selected fluid formalization 85

Chapter 1

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.

1.1 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 [9]. 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 3 and the quantum algorithms in Chapter 4 fit into these particular stages.

1.1.1 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

$$\rho = \int f \, dv, \qquad \rho u = \int v f \, dv, \qquad 3\rho\theta = \int |v - u|^2 f \, dv. \tag{1.1}$$

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.

1.2 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 [31]. 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 [55]. Itatani's molecular-orbital tomography uses high-harmonic emission, alignment information, a reference atom, and a recollision approximation [27]. Vozzi's generalized tomography improves phase recovery and the treatment of Coulomb distortion while retaining finite-bandwidth and reconstruction assumptions [66].

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 1.1 (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 1.2 (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 [59].

1.2.1 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

$$|A(\vartheta)|^2 = |a(\vartheta)|^2 \frac{\sin^2(s\delta/2)}{\sin^2(\delta/2)}. \tag{1.2}$$

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 (1.2) 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 [30, 31].

1.3 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 1.3 (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.

1.3.1 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 [41]. 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

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle. \tag{1.3}$$

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.

1.4 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 [3, 59].

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 1.4 (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

$$-\frac{\hbar^2}{2M}\left(\Delta\chi_m + 2\sum_n d_{mn}\cdot\nabla\chi_n + \sum_n \tau_{mn}\chi_n\right) + E_m(R)\chi_m = E\chi_m \tag{1.4}$$

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 1.5 (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 (1.4); 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 1.6 (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 (1.4) and Proposition 1.6 state what must be separated.

1.5 When a quantum packet behaves approximately as a particle

For $H = P^2/(2M) + V(X)$, the formal Ehrenfest identities are

$$\frac{d}{dt}\langle X\rangle = \frac{\langle P\rangle}{M}, \qquad \frac{d}{dt}\langle P\rangle = -\langle \nabla V(X)\rangle. \tag{1.5}$$

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 1.7 (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

$$\mathrm{Var}(X(t)) = \sigma^2 + \frac{\hbar^2 t^2}{4M^2\sigma^2}. \tag{1.6}$$

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 (1.6) 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.

1.5.1 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 [15]. 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 [38, 39].

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.

1.6 Molecular dynamics and conserved microscopic information

A classical molecular Hamiltonian has the form

$$H(q,p) = \sum_{i=1}^{N} \frac{|p_i|^2}{2m_i} + U(q), \qquad \dot{q}_i = \frac{p_i}{m_i}, \quad \dot{p}_i = -\nabla_{q_i} U. \tag{1.7}$$

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 [3]. Their scope is much wider than the two-particle illustration in the video.

Proposition 1.8 (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.

1.6.1 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

$$\widetilde{H}_h(q, p) = \frac{1}{2}p^2 + \frac{1}{2}\omega^2 \left(1 - \frac{h^2\omega^2}{4}\right) q^2. \tag{1.8}$$

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.

1.7 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

$$\partial_t F_N^{(1)} + v \cdot \nabla_x F_N^{(1)} + (N-1)\nabla_v \cdot \int K(x-y) F_N^{(2)}(x, v, y, w)\, dy\, dw = 0, \tag{1.9}$$

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 1.9 (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.

1.7.1 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 [21]. 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.

1.8 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

$$v' = v - (g \cdot n)n, \qquad v'_* = v_* + (g \cdot n)n. \tag{1.10}$$

Proposition 1.10 (Elastic reflection)**.** *The map* (1.10) *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 $1, v_1, \ldots, v_d, |v|^2$ are immediately invariant. The following proof makes the converse precise under a simple regularity assumption.

Proposition 1.11 (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 1.10. □

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.

1.8.1 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 [61]. 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 [49, 58]. 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.

1.9 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

$$Q(f, f)(v) = \int_{\mathbb{R}^d}\int_{S^{d-1}} B(v - v_*, n)(f'f'_* - ff_*)\, dn\, dv_*. \tag{1.11}$$

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 1.12 (Weak collision identity)**.** *For an admissible test function ϕ,*

$$\begin{aligned}\int \phi Q(f, f)\, dv &= \frac{1}{4}\iiint B(f'f'_* - ff_*)(\phi + \phi_* - \phi' - \phi'_*)\, dn\, dv_*\, dv \\ &= \frac{1}{2}\iiint Bff_*(\phi' + \phi'_* - \phi - \phi_*)\, dn\, dv_*\, dv.\end{aligned} \tag{1.12}$$

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 1.13 (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 (1.12) with $\phi = \log f$ are integrable. Then*

$$\frac{d}{dt}\int f \log f \, dv = -D(f) \le 0, \quad D(f) = \frac{1}{4}\iiint B(f'f'_* - ff_*) \log \frac{f'f'_*}{ff_*} \, dn \, dv_* \, dv. \tag{1.13}$$

Proof. The derivative equals $\int (1+\log f) Q(f, f) \, dv$. The constant part vanishes by mass conservation. Use (1.12) 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 1.14 (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 [62]. 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.

1.9.1 Entropy minimization is not a convergence proof

Proposition 1.15 (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 1.11 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.

1.10 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

$$P = \int c\otimes c\, f\,dv, \quad e = \frac{1}{2\rho}\int |c|^2 f\,dv, \quad q = \frac{1}{2}\int |c|^2 cf\,dv, \quad E = \frac{1}{2}\rho|u|^2 + \rho e.$$

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 1.16 (Moment equations before closure)**.** *Under these hypotheses,*

$$\partial_t\rho + \nabla\cdot(\rho u) = 0, \tag{1.14}$$

$$\partial_t(\rho u) + \nabla\cdot(\rho u\otimes u + P) = 0, \tag{1.15}$$

$$\partial_t E + \nabla\cdot(Eu + Pu + q) = 0. \tag{1.16}$$

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,

$$\frac{1}{2}\int |v|^2 vf\,dv = \left(\frac{1}{2}\rho|u|^2 + \rho e\right)u + Pu + q.$$

These expansions give all three equations. □

Example 1.17 (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 [22].

1.10.1 First-order constitutive laws and scaling

For a monatomic gas near local equilibrium, a formal Chapman–Enskog expansion produces the constitutive form

$$P = pI - 2\mu S_0 - \zeta(\nabla\cdot u)I, \qquad q = -\kappa\nabla T, \quad S_0 = \tfrac{1}{2}(\nabla u + \nabla u^T) - \tfrac{1}{3}(\nabla\cdot u)I. \tag{1.17}$$

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 [60, 62].

1.10.2 A precise historical hydrodynamic-limit theorem

The result of Golse and Saint-Raymond that is useful here concerns the scaled equation on $\mathbb{R}^3_x \times \mathbb{R}^3_v$

$$\varepsilon^2\partial_t F_\varepsilon + \varepsilon v\cdot\nabla_x F_\varepsilon = Q(F_\varepsilon, F_\varepsilon), \qquad M(v) = (2\pi)^{-3/2}e^{-|v|^2/2}. \tag{1.18}$$

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

$$\sup_{\varepsilon>0}\varepsilon^{-2}H(F_\varepsilon^{\mathrm{in}} \mid M) < \infty, \quad H(F\mid M) = \iint \left(F\log\frac{F}{M} - F + M\right)dx\,dv. \tag{1.19}$$

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 [22, 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 2. That chapter's precise theorem must be used with its own forcing and solution hypotheses, rather than inferred from this historical limit theorem.

1.11 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 [65]. 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 1.18 (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 1.19 (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 1.20 (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 $\Pr(|\overline{Y} - \mathbb{E}Y| \geq a) \leq \sigma^2/(na^2)$. The statistics and Monte Carlo decks are useful introductions [63, 64]; 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.

1.12 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 [18]. 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 $V^T w = 0$. If $\dot{z} = F(z)$, the exact retained coordinate satisfies $\dot{a} = V^T F(Va + w)$. Replacing the right side by $V^T F(Va)$ is a closure. The exact memory identity in Chapter 4 explains why eliminated coordinates can leave both memory and an initial-condition contribution.

Proposition 1.21 (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.

1.12.1 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 1.22 (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 [18].

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.

1.12.2 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 (1.17). This architecture has a limited but useful exact guarantee.

Proposition 1.23 (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.

1.13 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 3 proves the finite structural statements, and Chapter 2 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 2

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.

2.1 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 [13, 48].

Let $\nu > 0$ and define the momentum residual

$$\mathcal{R}_\nu(u,p) = \partial_t u + (u\cdot\nabla)u - \nu\Delta u + \nabla p. \tag{2.1}$$

The source's main conclusion can be recorded without changing its quantifiers:

$$\begin{aligned} &f \in C_c^\infty(\mathbb{R}^3\times(0,\infty);\mathbb{R}^3), \qquad u,p \in C^\infty(\mathbb{R}^3\times[0,1)),\\ &\mathcal{R}_\nu(u,p) = f, \quad \nabla\cdot u = 0, \quad u(\cdot,0) = 0, \qquad (2.2)\\ &\operatorname{supp} u(\cdot,t) \cup \operatorname{supp} p(\cdot,t) \subset K \quad (0 \le t < 1),\\ &\sup_{t<1}\|u(t)\|_2 < \infty, \qquad \limsup_{t\uparrow 1}\|u(t)\|_\infty = \infty. \qquad (2.3) \end{aligned}$$

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 [47, 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

[19]. 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.

2.2 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 2.1 (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 2.2 (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 3.

Proposition 2.3 (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.

2.3 Concentration, energy, and the information lost by norms

Proposition 2.4 (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 2.5 (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 [17]. Our comparison family proves only the logical insufficiency of the weaker norm information, independently of any existence construction.

2.4 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 2.6 (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.

2.5 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 2.7 (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 2.8 (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 2.7, 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 [47]. 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.

2.6 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 2.9 (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 2.10 (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 2.9 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.

2.7 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 [46, 47].

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.

2.8 Localization and smooth force extension

Proposition 2.11 (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 2.12 (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.

2.9 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 [47].

The following separate estimate is useful for the project's computable fluid examples and for any learned approximate field.

Proposition 2.13 (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 [2, 11, 57].

2.10 Viscosity, periodization, and preservation of quantifiers

Proposition 2.14 (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 2.15 (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.

2.11 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 [46].

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 3 then develops the arithmetic, exterior-power, and filtered-algebra connections, while Chapter 4 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 3

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 [47], reviewed in Chapter 2. 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.

3.1 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.

3.2 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 3.1 (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 3.2 (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 2 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.

3.3 A quadratic order occurs in the actual construction

Section 6 of OpenAI [47] 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 3.3 (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 3.4 (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 3.5 (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.

3.4 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 3.6 (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 3.7 (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.

3.5 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 3.8 (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.

3.6 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 3.9 (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 3.10 (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.

3.7 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 [5, 6].

Proposition 3.11 (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 3.12 (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

$$\tfrac{1}{2}\frac{d}{dt}|a|^2 = -\nu a^T La - |B_\theta(a)a|^2 + a \cdot g(t).$$

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.

3.8 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 2, 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.

3.9 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 4

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.

4.1 Information reduction is a map, not an autonomous dynamics

The microscopic-to-macroscopic hierarchy in Chapter 1 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 4.1 (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 4.2 (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 [43]. 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 4.3 (Discrete memory and the resolvent determinant)**.** *For $x_{n+1} = Ax_n + Bz_n$, $z_{n+1} = Cx_n + Dz_n$,*

$$z_n = D^n z_0 + \sum_{j=0}^{n-1} D^{n-1-j} C x_j,$$

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.

4.2 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 4.4 (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 3.8 captures one finite algebraic component of that construction. Its analytic realization retains the carrier phase, exact curl correction, support separation and derivative estimates.

4.3 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 4.5 (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.

4.4 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 4.6 (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 4.7 (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.

4.5 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 4.8 (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.

4.6 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 4.9 (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.

4.7 What the quantum-algorithm literature actually supplies

Liu et al. [37] 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 4.10 (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 4.11 (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. [67] 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 4.12 (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 4.13 (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. [28] 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 [40]. Their practical result depends on optimization, ansatz approximation, measurement variance and implementation costs. Tensor-network methods such as Gourianov et al. [23] 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.

4.8 Exact finite repairs for collision block encodings

The detailed algorithm of Jennings et al. [28] 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 4.14 (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 4.15 (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 4.16 (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 4.17 (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. □

4.9 A global compatibility test for a spin representation

Meng and Yang [42] 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

$$\omega \cdot \nabla\phi = -\frac{2\nu}{\rho_0}|\nabla\psi|^2(u \cdot \omega), \qquad \omega = \nabla \times u. \tag{4.1}$$

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 4.18 (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 (4.1) 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 (4.1) 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 4.19 (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 4.18.

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.

4.10 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 3.12. The resulting dissipation identity is true for every choice of circuit parameters; it does not rely on the optimizer finding a good model.

Proposition 4.20 (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 3.12.

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 4.21 (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.

4.10.1 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 4.2 supplies a reference target and a complete exact comparison. For a nonlinear flow, the corresponding error and stability hypotheses require new analysis.

4.11 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 5

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.

5.1 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 5.1 (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 5.2 (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.

5.2 A lower bound on every approximate closure

The preceding obstruction has a quantitative version, independent of how G is trained or parameterized.

Proposition 5.3 (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.

5.3 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 [36].

Proposition 5.4 (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.

5.4 A current analytic obstruction to removing the force

Theorem 1.1 of Constantin et al. [14], 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.

5.5 Formal coverage and reproducibility

The companion module proves the two directions and uniqueness in Theorem 5.1, its discrete-closure specialization and a mismatched-fiber obstruction, Proposition 5.2, both inequalities in Proposition 5.3, and every polynomial assertion in Proposition 5.4. 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 6

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.

6.1 Source scope and a quantitative contract

The source revisions are:

- Tropical_Quivers_of_Archs [52], 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.

6.2 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 6.1 (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 6.2 (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.

6.3 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 6.3 (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 6.1 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.

6.4 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 6.4 (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 6.5 (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.

6.5 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 6.6 (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 6.7 (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. [33] 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.

6.6 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 6.1 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 6.8 (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.

6.7 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. [68]; the separate TokenGT construction is Kim et al. [32]. 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 [56] 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 [25] 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.

6.8 Independent repairs to the tropical-attention reference

Version 2 of Hashemi et al. [25], 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 6.9 (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 5.4. 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 6.10 (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 6.11 (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.

6.9 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 6.12 (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 $J_M^T J_M = \mathrm{diag}(0, M^2)$, whereas $J_M J_M^T = \mathrm{diag}(M^2, 0)$. 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.

6.10 Path multiplicity separates tropical growth from operator stability

The max-plus cocycle in Schreiber [52] 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 6.13 (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 6.14 (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.

6.11 Weighted small-gain certification at a fixed rate

The weighted small-gain criterion in Schreiber [52] 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 6.15 (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.

6.11.1 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 6.16 (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

$$d_{n,i} \le R_n w_i, \qquad R_n = q^n r_0 + \sum_{k=0}^{n-1} q^{n-1-k} \delta_k. \tag{6.1}$$

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 (6.1) 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.

6.11.2 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.

6.12 Proper scoring must include its probability boundary

The cross-entropy theorem in Schreiber [51] 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 6.17 (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.

6.13 Predictive synergy is not determined by geometric overlaps

The last context-utility theorem of Schreiber [50] 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 6.18 (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.

6.14 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 7

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.

7.1 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 6. We retain that pinned meaning of "ToricGT-style"; the name alone imposes neither a hydrodynamic limit nor a physical quantum representation.[52–54]

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.

7.2 A finite operator algebra with conservative quantum transport

7.2.1 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

$$(\mathcal{T}\rho)_i = \sum_j P_{ij} \Phi_{ij}(\rho_j). \tag{7.1}$$

This finite construction is related to open quantum walks, whose operator-valued transition rules combine quantum internal states with classical position records.[7]

Proposition 7.1 (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.

7.2.2 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.[35] 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.[4, 12]

7.3 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 7.2 (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 6. 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.

7.4 Nonlinear dynamics on an operator state space

7.4.1 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. [8, 12]

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.

7.4.2 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.[10] The residual belongs in an error recurrence such as Theorem 6.16. This connects operator learning to a checked control budget without declaring a finite closure exact.

7.5 Control, observability and robustness

7.5.1 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 6.16. 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.[20]

7.5.2 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.[1] 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.[26]

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.

7.6 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

$$F_i(\rho) = a\rho_i + (1 - a)\tau(u_i). \tag{7.2}$$

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 7.3 (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.

7.7 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. [24]

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. [4, 12]
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.[8]
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.[34]
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.[45]
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. [44]
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.[16, 70]

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.

7.8 De novo substrate design with generative models

7.8.1 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.[29, 69]

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.

7.8.2 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.[69] Generative photonic inverse design directly targets an optical response, but its particular demonstrated response still needs an additional mapping to this graph-fluid architecture.[29] 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.

7.9 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.

Appendix A

Fluid verification, reproducibility and proof dependencies

A.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.

A.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.

A.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.

A.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 4; 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.

A.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 7.

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.

A.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
```

Bibliography

[1] 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.

[2] 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.

[3] 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.

[4] 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.

[5] 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.

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] Clay Mathematics Institute. Navier–Stokes announcement, September 2026. URL https://www.claymath.org/news/navier-stokes-announcement/. 11 September 2026; retrieved 13 November 2026.

[14] 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.

[15] 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.

[16] 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/.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] 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.

[24] 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.

[25] 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.

[26] 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.

[27] 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.

[28] 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.

[29] 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.

[30] 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.

[31] 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.

[32] 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.

[33] 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.

[34] 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.

[35] 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.

[36] 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.

[37] 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.

[38] 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.

[39] 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.

[40] 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.

[41] 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.

[42] 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.

[43] 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.

[44] 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.

[45] 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.

[46] 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.

[47] 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.

[48] 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.

[49] 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.

[50] 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.

[51] 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.

[52] 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.

[53] 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.

[54] 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.

[55] 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.

[56] 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.

[57] 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.

[58] 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.

[59] 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.

[60] 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.

[61] 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.

[62] 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.

[63] 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.

[64] 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.

[65] 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.

[66] 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.

[67] 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.

[68] 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.

[69] 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.

[70] 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/.