Synchronization and Resource Viability in Stuart–Landau Oscillator Networks:

Certified Expenditure Bounds and Validated Enclosures

Flyxion
Independent Researcher

November 2026

Abstract

A network of coupled oscillators may converge to synchrony while the reservoir that powers it approaches exhaustion, and a network that is energetically viable may remain incoherent. This monograph treats coordination, energetic feasibility and the preservation of future admissible operation as three distinct mathematical predicates, and develops a sequence of certificates that connect them for Stuart–Landau networks with an explicit resource state. Starting from a phase-network model with nonnegative storage, bounded supply and state-dependent consumption, the text establishes positivity under inward-pointing boundary conditions, invariant operating rectangles, a conditional contraction estimate, and counterexamples showing that synchronization neither implies viability nor is implied by it. Expenditure is then derived from an explicit actuator law rather than imposed, and two plausible conventions are shown to give opposite conclusions about the cost of faster coordination. For the full complex-amplitude network, cubic monotonicity yields transverse disagreement bounds and a hierarchy of cumulative-expenditure certificates: global, decay-sensitive, modewise, edge-dissipative, and direct spectral. Persistent frequency heterogeneity is handled through forced-disagreement envelopes, graph-resolved forcing metrics, an exact signed-work identity, and prospective checkpoint-restarted enclosures. A tunable Young allocation is derived and optimized, and the resulting certificates are calibrated against integrated trajectories. For the two-oscillator real reduction, joint dissipation and Riccati comparison give uniform-in-time expenditure bounds. Fixed uncertain parameters are treated by set-membership recursions, Taylor sensitivity enclosures and continuous-cell exclusion, including a reported negative result. Finally, validated propagation by Picard tubes, outward dyadic rounding, affine generators and moving frames is compared systematically, with attention to cases in which a narrower state enclosure does not yield a tighter bound on accumulated expenditure. Every result is stated with its assumptions, the arithmetic under which it holds, and the claims it does not make. The expenditure laws are constitutive models and not calibrated physiology.

Keywords: Stuart–Landau oscillators; synchronization; viability theory; actuation expenditure; Laplacian spectral bounds; edge dissipation; validated numerics; Picard enclosures; affine generators.

Contents

I Foundations **5**

1 Scope, Conventions and Evidentiary Boundary **6**
1.1 The Problem 6
1.2 Structure of the Argument 6
1.3 Conventions 7
1.4 Reading Guide 7

2 A Resource-Constrained Oscillator Network **8**
2.1 Phase Network and Admissibility 8
2.2 Boundary Behaviour and Invariance 9
2.3 Synchronization as a Separate Obligation 9
2.4 Finite Supply and Cumulative Debt 10
2.5 Full Stuart–Landau Amplitude Dynamics with Resource Accounting 10
2.6 A First Transverse Estimate and the Open Trade-Off 11
2.7 Relation to the Published Model 12

II Expenditure Laws and Analytic Certificates **13**

3 Expenditure Laws for Coordination **14**
3.1 Charging for Coupling Capacity 14
3.2 Dissipation Accounting versus Actuator Power 15
3.3 A Resistive Realization 16
3.4 The Full Stuart–Landau Network under Actuator Accounting 16

4 Transverse Disagreement and Global Certificates **18**
4.1 Model and Stopping Convention 18
4.2 Disagreement Contraction from Cubic Monotonicity 18
4.3 Exact and Conservative Resource Criteria 19
4.4 The Decay-Sensitive Certificate 19

5 Spectral and Edge-Resolved Certificates **21**
5.1 Spectral Decomposition and an Exact Linear Benchmark 21
5.2 The Modewise Certificate 22

5.3 Time-Dependent Modal Forcing . . . 22
5.4 The Edge-Dissipative Certificate . . . 23
5.5 Comparison of Certificates . . . 23
5.6 A Certificate for Real Networks with Common Growth . . . 24

6 Persistent Frequency Heterogeneity and Network-Level Certificates 25
6.1 Forced Disagreement for the Full Complex Network . . . 25
6.2 Graph-Resolved Forcing . . . 26
6.3 The Heterogeneous Edge Certificate . . . 27
6.4 Direct Spectral Certification . . . 28

7 Signed Frequency Work and Prospective Enclosures 29
7.1 Exact Edge-Energy Accounting . . . 29
7.2 Conditional and A Priori Certificates . . . 29
7.3 A Prospective Envelope from the Initial State . . . 30
7.4 Checkpoint-Restarted Enclosures . . . 31
7.5 Degradation under Uncertain Observation and Parameters . . . 32

8 Optimal Allocation of Dissipation and Calibration against Trajectories 33
8.1 A Tunable Young Allocation . . . 33
8.2 Optimal Choices of the Parameter . . . 34
8.3 Numerical Results for the Four-Node Networks . . . 35
8.4 Calibration against Integrated Trajectories . . . 35
8.5 A Mode-Resolved Allocation . . . 37
8.6 A Time-Dependent Amplitude Envelope . . . 38
8.7 What the Calibration Implies . . . 39

III Reduced Systems and Parameter Uncertainty 40

9 Dissipation Certificates for the Two-Oscillator Reduction 41
9.1 The Reduced System and an Invariant Region . . . 41
9.2 The Joint Dissipation Certificate . . . 42
9.3 Time-Dependent Dissipation and Riccati Comparison . . . 42
9.4 Role as a Benchmark . . . 43

10 Persistent Parameters and Set-Valued Inference 44
10.1 Feasible Sets in the Product Space . . . 44
10.2 Velocity Observations and the Failure of State-Only Differences . . . 45
10.3 Joint State–Parameter Screening . . . 45
10.4 Parameter-Dependent Taylor Enclosures . . . 46
10.5 A Negative Result: Global Lipschitz Enclosures Retain Everything . . . 47

IV Validated Enclosures **48**

11 Interval Propagation and Parameter-Affine Taylor Models **49**
11.1 Global Majorants and a First-Order Validated Step 49
11.2 Picard Tubes and Higher-Order Time Coefficients 50
11.3 Cell-Centered Sensitivities 50
11.4 Parameter-Affine Taylor Models 51
11.5 What Remains Unsolved at This Level 51

12 Wrapping, Correlated Generators and Moving Frames **53**
12.1 An Exact Benchmark for Wrapping 53
12.2 Nonlinear Discrete Maps with Correlated Generators 54
12.3 Continuous-Time Validation by Picard Tubes 54
12.4 Resource-Floor Certification from Monotone Expenditure 55
12.5 Dyadic Compression and Fixed Parameter Intervals 55
12.6 Parameter-Cell Reconditioning 56
12.7 Equivalence of Coordinate Representations 56
12.8 Negative and Mixed Results on Moving Frames 57

V Synthesis **59**

13 What the Programme Establishes and What Remains Open **60**
13.1 A Taxonomy of Guarantees 60
13.2 Established Results 60
13.3 Mixed and Negative Results 61
13.4 Open Problems 62
13.5 An Extension to Intermittent Forcing 63
13.6 Closing Remarks 63

A Notation **64**

B Computational Materials **65**

Part I

Foundations

Chapter 1

Scope, Conventions and Evidentiary Boundary

1.1 The Problem

A recent computational study by Lombardo [11] couples Stuart–Landau neural oscillators with stochastic cardiac and respiratory rhythms and a scalar energy-like state, and reports model-internal associations between that state and synthetic cognitive proxies. Its neural power is $P = N^{-1} \sum_i |z_i|^2$ and its reported energy law is $\dot{E} = (0.4+0.2h(t))-0.05P(t)-\gamma E(t)$. The source limits its findings to the model, and the energy variable and its coefficients are phenomenological, so nothing below treats them as a calibrated physiological balance. What the study makes vivid is a structural question that can be posed independently of any biological interpretation: under which assumptions can coordination persist without violating storage, supply and recovery constraints?

The question separates into three predicates that are routinely conflated. Dynamical coordination asks whether the relative phases or complex amplitudes of the oscillators contract toward a common orbit. Energetic feasibility asks whether a resource trajectory stays above a stated floor under a stated supply. Preservation of future admissible operation asks whether a trajectory exists, or a control policy exists against every admissible disturbance, that remains inside an operating region for all later times. A system may satisfy the first and fail the second, as in a perfectly synchronized network that exhausts a finite reservoir at a constant basal rate. A system may satisfy the second and fail the first, as in an uncoupled network fed by a regulated supply. The monograph proves that these are independent within the model class studied, and then asks how much can nevertheless be said quantitatively when the three predicates interact through an expenditure law that charges for coordination.

1.2 Structure of the Argument

This monograph develops the subject in one sequence, with the numerical experiments that test each claim. Part I fixes the resource-constrained model, proves the separation results, and treats the full complex-amplitude network with an explicit resource state. Part II studies expenditure

laws and derives analytic certificates of cumulative actuator expenditure, beginning with crude global bounds and moving through spectral and edge-resolved refinements to network-level certificates under persistent frequency heterogeneity, signed frequency work, and the tunable allocation of dissipation between transient and persistent terms. Part III isolates the two-oscillator real reduction in which exact dissipation identities are available, and then treats fixed but uncertain parameters through set-membership recursions and Taylor sensitivity enclosures. Part IV compares validated propagation schemes, in which the enclosure of a reachable set is established by Picard tubes, outward-rounded arithmetic and correlation-preserving representations. Part V synthesizes the results and records which questions remain open. Two appendices give notation and the computational materials behind the tables.

1.3 Conventions

Throughout, $L = L^\top \succeq 0$ denotes the Laplacian of a connected undirected weighted graph on N vertices with eigenvalues $0 = \lambda_1 < \lambda_2 \leq \cdots \leq \lambda_N$, and $P = I - \mathbf{1}\mathbf{1}^\top/N$ is the projector onto the mean-zero subspace. States $z \in \mathbb{C}^N$ carry the real inner product $\mathrm{Re}\langle u, v\rangle = \mathrm{Re}\sum_j \bar{u}_j v_j$, so that $\mathbb{C}^N$ is treated as a real Hilbert space of dimension $2N$. The resource state is a scalar E with operating floor E_*, constant supply q and basal drain b unless stated otherwise. The actuator coefficient is $c_A > 0$ and coupling is $k > 0$. The first operational exit is $t_* = \inf\{t \geq 0 : E(t) < E_*\}$, and every estimate is restricted to $t < t_*$, because no claim is made about the dynamics after the operating constraints fail. The symbols $(\cdot)_+ = \max\{\cdot, 0\}$ and $\|\cdot\|$ for the Euclidean norm are used without further comment.

Three distinctions recur and are worth stating once. A *sufficient* certificate guarantees viability when its inequality holds, and its failure does not imply exhaustion. A *necessary and sufficient* criterion, such as the exact trajectory-wise integral identity for the resource, is by contrast exact but requires the trajectory. A *conservative* certificate is a sufficient one whose looseness has been quantified against an observed or exactly solvable case. Second, the arithmetic under which a statement holds matters: some results are theorems in real arithmetic, some are enclosures computed in exact rational arithmetic with outward dyadic rounding, some are enclosures computed with outward-rounded multiprecision interval arithmetic whose library is trusted but unaudited, and some are numerical integrations of a trajectory used only for comparison. Each statement below is labelled accordingly. Third, the expenditure laws are constitutive. When a law is written $c_A k^2 \|Lz\|^2$ it is a specified actuator model, and no claim is made that neurons or any other biological system obey it.

1.4 Reading Guide

Readers interested only in the certificates for the full network may proceed from Chapter 2 directly to Chapters 5 and 6. The comparison of validated propagators in Chapter 12 is self-contained given the reduced model of Chapter 9. Negative and mixed results are retained throughout, since a catalogue of bounds that reports only the best number would misrepresent how far each method can be trusted. Chapter 13 gathers them.

Chapter 2

A Resource-Constrained Oscillator Network

2.1 Phase Network and Admissibility

Let $G = (V, \mathcal{E})$ be a connected undirected weighted graph with symmetric weights $a_{ij} \geq 0$ and Laplacian L. Let $\theta_i \in \mathbb{T}$ be phases, $e \geq 0$ usable energy, and $r \in [0, R]$ a supply reservoir. On a cohesive chart in which edge differences have magnitude below $\pi/2$ consider

$$\dot{\theta}_i = \omega_i - k(e) \sum_j a_{ij} \sin(\theta_i - \theta_j) + u_i(t), \tag{2.1}$$

$$\dot{e} = s(r, t) - c_0 - c_1 D(\theta) - \ell(e), \tag{2.2}$$

$$\dot{r} = q(t) - s(r, t) - d(r), \tag{2.3}$$

where $k(e) \geq 0$ is coupling capacity, $c_0 > 0$ is basal expenditure, $c_1 \geq 0$ weights a coordination-dependent cost through $D(\theta) = \frac{1}{2} \sum_{i,j} a_{ij}[1 - \cos(\theta_i - \theta_j)]$, q is replenishment, s is transfer from reservoir to usable storage, and $\ell, d \geq 0$ are losses. A positive c_1 makes incoherence costly; the opposite empirical hypothesis is representable by another cost function, and no energetic advantage of synchrony is assumed to be a biological fact. The phase equations are invariant under a common shift, so coherence is measured by $W(\theta) = \frac{1}{2}\|\vartheta\|^2$ with $\vartheta = \theta - \bar{\theta}\mathbf{1}$ on a consistent local lift.

A state is *operationally admissible* when $(\theta, e, r) \in \mathcal{K} := \{\theta \in \Theta_\delta,\ e \in [e_*, E_*],\ r \in [r_*, R]\}$, where Θ_δ is a compact cohesive phase region. The distinction between mathematical nonnegativity $e \geq 0$ and a positive engineering reserve $e \geq e_* > 0$ is essential.

Definition 2.1 (Continuation set and viability kernel)**.** For a state x and admissible input class $\mathcal{U}$, let $\mathcal{C}_H(x)$ be the set of trajectories on $[0, H]$ that satisfy the equations, the input constraints, and $x(t) \in \mathcal{K}$ for all $t \in [0, H]$. The existential viability kernel is $\mathrm{Viab}(\mathcal{K}) = \{x \in \mathcal{K} : \mathcal{C}_\infty(x) \neq \varnothing\}$. For a disturbance class $\mathcal{D}$, the robust kernel requires one admissible control policy that keeps the state in $\mathcal{K}$ against every disturbance in $\mathcal{D}$.

The order of quantifiers is the whole content: existence of one favourable supply schedule is weaker than survival under every admissible supply interruption. The viability framework

follows Nagumo [15] and Aubin [3], and the set-invariance viewpoint of Blanchini [6].

2.2 Boundary Behaviour and Invariance

Proposition 2.2 (Positivity at the storage boundary)**.** *Suppose the vector field is locally Lipschitz and, whenever $e = 0$, the transfer satisfies $s(r, t) \geq c_0 + c_1 D(\theta) + \ell(0)$. Then every solution with $e(0) \geq 0$ stays nonnegative as long as it exists. If the inequality fails at a reachable boundary state and the field is continuous, that state belongs to no forward-invariant nonnegative region.*

Proof. At $e = 0$ the inequality gives $\dot{e} \geq 0$, and the tangent cone of $[0, \infty)$ at zero is $[0, \infty)$, so the standard tangency criterion for locally Lipschitz ordinary differential equations gives invariance. Conversely, a strictly negative derivative at the boundary produces $e(t) < 0$ for small $t > 0$. □

The proposition exposes a modelling choice. A physical system with insufficient supply must terminate operation at the boundary, shed load, or change its dynamics, and continuing the unconstrained equation into negative energy is a mathematical artifact and not a prediction.

Proposition 2.3 (Invariant energy–reservoir rectangle)**.** *Assume $\theta(t) \in \Theta$ with $D(\theta) \leq D_{\max}$ there, and let $0 < e_* < E_*$, $0 \leq r_* < R$. If uniformly in admissible t and $\theta \in \Theta$ one has $s \geq c_0 + c_1 D_{\max} + \ell(e_*)$ at $e = e_*$, $s \leq c_0 + \ell(E_*)$ at $e = E_*$, $q \geq s(r_*, t) + d(r_*)$ at $r = r_*$, and $q \leq s(R, t) + d(R)$ at $r = R$, then $[e_*, E_*] \times [r_*, R]$ is forward invariant for the energy–reservoir subsystem along every phase trajectory in Θ.*

Proof. The four inequalities give $\dot{e} \geq 0$, $\dot{e} \leq 0$, $\dot{r} \geq 0$, $\dot{r} \leq 0$ on the corresponding faces, and the tangent-cone criterion for a rectangle applies. Invariance of Θ itself is a separate phase-dynamics obligation. □

The conditions are strong and may be infeasible for a finite reservoir with strictly positive basal cost. Their value is to make the obligations explicit instead of silently assuming sustainable operation.

2.3 Synchronization as a Separate Obligation

Proposition 2.4 (Conditional contraction of relative phases)**.** *Assume identical frequencies $\omega_i = \omega$, identical forcing $u_i = u(t)$, and a solution with $|\theta_i - \theta_j| \leq \delta < \pi/2$ on every edge. If $k(e(t)) \geq k_* > 0$ then $W(t) \leq W(0) \exp[-2k_* \lambda_2(L) \cos \delta \, t]$.*

Proof. Subtracting the mean phase, symmetry of the weights cancels common forcing and frequency, and $\dot{W} = -\frac{k(e)}{2} \sum_{i,j} a_{ij} (\theta_i - \theta_j) \sin(\theta_i - \theta_j)$. On $[-\delta, \delta]$ one has $x \sin x \geq (\cos \delta) x^2$, and $\sum_{i,j} a_{ij} (\vartheta_i - \vartheta_j)^2 = 2\vartheta^\top L \vartheta \geq 2\lambda_2 \|\vartheta\|^2$. Gronwall's inequality [8] finishes the argument. □

The estimate is conditional: it neither shows that an arbitrary state enters the cohesive region nor that $e(t) \geq e_*$, and energy depletion may drive $k(e)$ to zero and void its premise.

Example 2.5 (Perfect synchrony with finite-time loss of operation)**.** Take identical oscillators initially synchronized, $k(e) \equiv k_0 > 0$, $D = 0$, $s = q = d = \ell = 0$ and $c_0 > 0$. The synchronized manifold is invariant while $e(t) = e(0) - c_0 t$, so the threshold e_* is reached at $T_* = (e(0) - e_*)/c_0$, and no continuation in $\mathcal{K}$ exists beyond T_*.

Example 2.6 (Viability without synchronization)**.** Let $k(e) \equiv 0$, take distinct constant frequencies, and let a regulated supply balance $c_0 + c_1 D(\theta) + \ell(e_0)$ while replenishment holds $r = r_0$. Then e and r stay admissible although relative phases need not converge. This is an idealized controlled counterexample and does not assert that finite reservoirs can supply indefinitely.

Together the examples show that synchronization and viability are logically independent within the model class.

2.4 Finite Supply and Cumulative Debt

Adding (2.2) and (2.3) cancels internal transfer,

$$\tfrac{d}{dt}(e + r) = q(t) - c_0 - c_1 D(\theta) - \ell(e) - d(r), \tag{2.4}$$

a bookkeeping identity of the chosen equations that implies nothing about calibration to physical units.

Proposition 2.7 (Supply-deficit obstruction)**.** *Suppose $q(t) \le \bar{q} < c_0$ and $\ell, d, c_1 D \ge 0$. Then every solution with finite initial total resource leaves $\{e \ge e_*, r \ge r_*\}$ no later than $T_{\max} = (e(0) + r(0) - e_* - r_*)/(c_0 - \bar{q})$, provided the numerator is nonnegative. Synchronization cannot prevent the exit.*

Proof. Equation (2.4) gives $\dot{e} + \dot{r} \le -(c_0 - \bar{q})$, so $e + r \le e(0) + r(0) - (c_0 - \bar{q})t$, while remaining in the rectangle requires $e + r \ge e_* + r_*$. □

Integrating (2.4) gives a finite-horizon necessary condition on cumulative supply, which is not sufficient because resource may arrive too late or transfer rates may be too small. Timing and transfer capacity are independent continuation constraints.

2.5 Full Stuart–Landau Amplitude Dynamics with Resource Accounting

Restoring complex amplitudes, let $z_i \in \mathbb{C}$ and

$$\dot{z}_i = (\alpha + \mathrm{i}\omega_i - |z_i|^2) z_i - k(e) \sum_j L_{ij} z_j + b_i(t) z_i + \xi_i(t), \tag{2.5}$$

$$\dot{e} = s(r, t) - c_0 - c_2 P(z) - c_1 D(z) - \ell(e), \tag{2.6}$$

$$\dot{r} = q(t) - s(r, t) - d(r), \tag{2.7}$$

with $P(z) = N^{-1}\sum_i |z_i|^2$, $D(z) = z^*Lz$, bounded real modulation b_i, and a bounded disturbance ξ. This is a proposed extension and not a literal reproduction of the source equations. The transfer cancels in the total balance as before.

Proposition 2.8 (Amplitude absorbing bound)**.** *Suppose $k(e) \geq 0$, $\alpha + b_i(t) \leq A$ and $\|\xi(t)\| \leq F$. With $V = \frac{1}{2}\|z\|^2$, $\dot{V} \leq 2AV - \frac{4}{N}V^2 + F\sqrt{2V}$, and consequently $V(t) \leq V_{\max}$ with $V_{\max}$ depending only on $A, F, N, V(0)$. No finite-time norm blow-up occurs.*

Proof. The real part of $\sum_i \bar{z}_i \dot{z}_i$ removes the frequencies, the Laplacian contribution is $-k(e)z^*Lz \leq 0$, and $\sum_i |z_i|^4 \geq N^{-1}(\sum_i |z_i|^2)^2 = 4V^2/N$. Cauchy–Schwarz bounds the forcing by $F\sqrt{2V}$, and the right side is negative for large V. □

The bound does not give a lower bound on any $|z_i|$, so phases need not remain defined and Kuramoto-style phase contraction cannot be transferred automatically to the full system. Indeed, amplitude death $z_i \to 0$ can occur while a phase-only model keeps reporting coherent phases, so a faithful simulation must record both the phase order parameter and the minimum amplitude and mark phase coherence as undefined whenever the amplitude threshold fails.

Proposition 2.9 (Synchronous manifold)**.** *Under identical frequencies, modulation and additive input, the diagonal $\mathcal{S} = \{z_1 = \cdots = z_N\}$ is forward invariant, $D = 0$ on it, and the common amplitude obeys $\dot{z} = (\alpha + \mathrm{i}\omega - |z|^2 + b)z + \xi$. If $\xi = 0$, b is constant and $\alpha + b > 0$, the nonzero synchronized orbit has radius $\sqrt{\alpha + b}$, which says nothing about resource sustainability.*

Proof. On the diagonal $L\mathbf{1} = 0$ removes the coupling, all right-hand sides coincide, and uniqueness preserves equality. The radial equation is $\dot{a} = (\alpha + b - a^2)a$. □

Proposition 2.10 (Deficit despite full-amplitude synchrony)**.** *If $q \leq \bar{q} < c_0$ and $P, D, \ell, d \geq 0$, any solution of (2.5)–(2.7) with finite initial resources leaves $\{e \geq e_*, r \geq r_*\}$ by $T_{\text{exit}} \leq (e(0) + r(0) - e_* - r_*)/(c_0 - \bar{q})$, even on $\mathcal{S}$ and with arbitrarily strong coupling.*

Proof. Integrate the total balance and use nonnegativity of the expenditure terms; the oscillator phases and the coupling coefficient play no role. □

2.6 A First Transverse Estimate and the Open Trade-Off

Write $Q = I - \mathbf{1}\mathbf{1}^\top/N$, $w = Qz$, $\bar{z} = N^{-1}\sum_i z_i$. A first transverse input-to-state estimate follows by a Lipschitz bound on the cubic.

Proposition 2.11 (Lipschitz transverse bound)**.** *Assume $\lambda_2(L) > 0$, $k(e(t)) \geq k_*$, $\|z\| \leq R$, $\alpha + b_i \leq A$. Put $\gamma = k_*\lambda_2 - A - 3R^2$ and $m(t) = \|Q((b_i + \mathrm{i}\omega_i)\bar{z} + \xi_i)_i\|$. If $\gamma > 0$ then $\|w(t)\| \leq e^{-\gamma t}\|w(0)\| + \int_0^t e^{-\gamma(t-s)}m(s)\,ds$.*

Proof. Since $QL = LQ = L$ and $\langle w, Lw\rangle \geq \lambda_2\|w\|^2$, subtract $f_i(\bar{z})$ from $f_i(z_i)$ with $f_i(u) = (\alpha + b_i + \mathrm{i}\omega_i)u - |u|^2u$. The frequency term has zero real inner product with w_i, the linear part is bounded by $A\|w\|^2$, and on $|u| \leq R$ the cubic is $3R^2$-Lipschitz. The remaining projected forcing is the vector defining m, and scalar comparison for the upper Dini derivative gives the claim. □

The threshold is conservative because it charges the cubic term $3R^2$ as a destabilizing contribution. Chapter 4 removes this term by exploiting monotonicity. A claim that raising coupling shrinks the viability kernel needs an explicit expenditure mechanism, because larger k may reduce D and thereby improve viability. The trade-off therefore remains a conjecture in this chapter.

Conjecture 2.12 (Restricted synchronization–viability trade-off)**.** For a class of coupled amplitude oscillators with nonzero coordination expenditure and bounded replenishment there are parameter regimes in which increasing coupling improves finite-time coherence while reducing the measure of the robust viability kernel.

Chapter 3 derives the expenditure laws under which this behaviour can be exhibited and shows numerically that it is parameter dependent and not universal.

2.7 Relation to the Published Model

The cited oscillator model [11] uses complex amplitudes, Laplacian coupling, a cardiac oscillator, respiratory input and multiplicative modulation by a scalar energy-like variable. The phase model above discards amplitude to isolate a tractable theorem, so its propositions do not apply automatically to the published model. A faithful extension retains the amplitudes and replaces the single energy equation by explicit storage, replenishment and dissipation equations, as in (2.5)–(2.7). That model's synthetic cognition proxies are model outputs, and a mediation effect computed from model-generated variables is neither a theorem of feasibility nor a validation of a physiological mechanism. Model-internal mathematical outcomes and independently measured biological outcomes would have to be reported in separate tables.

Part II

Expenditure Laws and Analytic Certificates

Chapter 3

Expenditure Laws for Coordination

The separation results of Chapter 2 are qualitative. To obtain quantitative statements one must say what coordination costs, and the answer is not forced by the oscillator equations. This chapter examines three specifications of that cost: a standing charge for maintaining coupling capacity, a charge proportional to the rate at which coupling removes disagreement, and a charge for quadratic actuator effort derived from a resistive transducer. They lead to different conclusions about whether faster coordination is more expensive, and the difference is itself a result.

3.1 Charging for Coupling Capacity

Consider unwrapped phases $\theta \in \mathbb{R}^N$ and available energy $E \geq E_*$ with

$$\dot{\theta} = \omega \mathbf{1} - kL\theta, \qquad \dot{E} = q - b - ck^2, \tag{3.1}$$

where $k \geq 0$ is a fixed coupling intensity and $c > 0$ is a cost coefficient. The quadratic charge represents continuously provisioned capacity, including at exact agreement. It is not the dissipative edge-disagreement cost $k\theta^\top L\theta$, which vanishes at synchrony. With $w = P\theta$ and finite-horizon viability defined as $E(t) \geq E_*$ on $[0, H]$, the target of synchronization is $\|w(\tau)\| \leq \varepsilon$ for some $\tau > 0$.

Proposition 3.1 (Contraction and exact viability threshold)**.** *For $k > 0$ one has $\|w(t)\| \leq e^{-k\lambda_2 t}\|w_0\|$. The system is viable on $[0, H]$ if and only if $E_0 - E_* + H \min\{0, q - b - ck^2\} \geq 0$. If $q - b - ck^2 < 0$ the first exit time is $T_{\text{exit}} = (E_0 - E_*)/(b + ck^2 - q)$.*

Proof. Projecting gives $\dot{w} = -kLw$ because $P\mathbf{1} = 0$ and $PL = L$, and the spectral theorem gives the norm estimate. Integrating the energy equation gives $E(t) = E_0 + (q - b - ck^2)t$, whose minimum over $[0, H]$ occurs at an endpoint. □

Proposition 3.2 (A synchronization–viability window)**.** *A sufficient synchronization threshold is $k \geq k_{\text{sync}} = \log(\|w_0\|/\varepsilon)/(\lambda_2 \tau)$, and exact finite-horizon viability is equivalent to $k \leq k_{\text{viable}} = \sqrt{(q - b + (E_0 - E_*)/H)/c}$ whenever the radicand is nonnegative. A coupling meeting both exists whenever $k_{\text{sync}} \leq k_{\text{viable}}$. If $k_{\text{sync}} > k_{\text{viable}}$ no coupling meets the sufficient synchronization guarantee together with viability; for a worst-case initial condition in the λ_2 eigenspace the threshold is also necessary, so the two goals are then genuinely incompatible.*

On an eight-node ring with $E_0 = 1$, $E_* = 0.1$, $q = 0.20$, $b = 0.10$, $c = 0.02$ and $H = 40$, one finds $k_{\text{viable}} = \sqrt{(0.10 + 0.90/40)/0.02} \approx 2.475$. The exact matrix-exponential sweep reproduces the threshold: couplings $0.5, 1, 2$ remain viable with terminal energies $4.8, 4.2, 1.8$, while $k = 3$ exits at $t = 11.25$ and $k = 4$ at $t \approx 4.09$, with terminal energies -2.2 and -7.8 obtained by continuing the unconstrained equation, which has no operational meaning past exit. The terminal disagreement is negligible at every tested coupling. The hypothesis behind (3.1) is a modelling assumption; it is not licensed by the disagreement dissipation $\mathrm{Re}(z^* L z)$ of the full model, which vanishes on exact synchrony.

3.2 Dissipation Accounting versus Actuator Power

Remove the standing charge and derive expenditure from the control force. In a rotating frame the homogeneous linearized subsystem is $\dot{x} = -kLx$ with actuation $u = -kLx$. Two nonnegative conventions are natural:

$$p_{\text{diss}} = c_D k\, x^\top L x, \qquad p_{\text{act}} = c_A \|u\|^2 = c_A k^2 \|Lx\|^2.$$

For a unit Laplacian eigenvector v with eigenvalue $\lambda > 0$ and $x(0) = av$ one has $x(t) = ae^{-k\lambda t} v$.

Proposition 3.3 (Exact transient expenditures)**.** *The cumulative costs up to H are $C_D(k, H) = \frac{c_D a^2}{2}(1 - e^{-2k\lambda H})$ and $C_A(k, H) = \frac{c_A a^2 k \lambda}{2}(1 - e^{-2k\lambda H})$. As $H \to \infty$, $C_D \to c_D a^2/2$ independently of k, while $C_A \to c_A a^2 k\lambda/2$. For fixed $H > 0$, C_A is strictly increasing in k.*

Proof. Insert the solution and integrate. For monotonicity, $\frac{d}{dk}[k(1-e^{-\beta k})] = 1 - e^{-\beta k} + \beta k e^{-\beta k} > 0$ for $\beta = 2\lambda H > 0$. □

Hence the statement that faster synchronization costs more depends on the power convention and the observation horizon. Under dissipation accounting the infinite-horizon cost is independent of the coupling rate, while under actuator accounting it grows linearly with k, because the control amplitude grows faster than its duration shrinks.

Proposition 3.4 (Viability and the synchronization window)**.** *Let $q \le b$ and $E(t) = E_0 + (q - b)t - C_A(k, t)$. Viability on $[0, H]$ is equivalent to $(b - q)H + C_A(k, H) \le E_0 - E_*$. The synchronization requirement $\|x(\tau)\| \le \varepsilon < |a|$ holds exactly when $k \ge k_{\text{sync}} = \log(|a|/\varepsilon)/(\lambda\tau)$, and a feasible constant coupling exists if and only if the viability inequality holds at $k = k_{\text{sync}}$. In particular, if $(b - q)H + C_A(k_{\text{sync}}, H) > E_0 - E_*$ no constant coupling meets both the deadline and the energy floor.*

Proof. The energy is nonincreasing because $q \le b$ and $p_{\text{act}} \ge 0$, so its minimum is at H. The cost term increases strictly with k, so the least admissible coupling minimizes expenditure. □

For a general initial state $x_0 = \sum_{j \ge 2} a_j v_j$ one has $C_A(k, H) = \frac{c_A k}{2} \sum_{j \ge 2} \lambda_j a_j^2 (1 - e^{-2k\lambda_j H})$ and $\|w(t)\|^2 = \sum_{j \ge 2} a_j^2 e^{-2k\lambda_j t}$, so the single-mode threshold is replaced by a spectral inequality.

On an eight-node cycle with its first nonzero eigenvector, $a = 1$, $E_0 = 1$, $E_* = 0.1$, $q = b = 0.1$, $c_A = 2$, $H = 40$, $\tau = 5$ and $\varepsilon = 0.1$, the threshold is $k_{\text{sync}} = \log 10/(5\lambda_2) \approx 0.7862$ and, since the exponential term is negligible at $H = 40$, viability holds for $k\lambda_2 \le 0.9$, that is $k \lesssim 1.536$. The sweep confirms a nonempty window.

Table 3.1: Actuator-power accounting on the eight-node cycle. The deadline disagreement is $\|x(5)\|$ and the energy is evaluated at $H = 40$ from the closed-form expression.

k	Disagreement at τ	Energy at H	Synchronized	Viable
0.20	0.5567	0.8829	no	yes
0.50	0.2312	0.7071	no	yes
0.7862	0.1000	0.5395	yes	yes
1.00	0.0535	0.4142	yes	yes
1.50	0.0124	0.1213	yes	yes
2.00	0.0029	−0.1716	yes	no
4.00	8.2×10^{-6}	−1.3431	yes	no

3.3 A Resistive Realization

To justify the quadratic law, suppose each complex control u_i is produced by currents $I_i \in \mathbb{R}^2$ through a linear transducer $u_i = \chi I_i$ with gain $\chi > 0$ and effective resistance $R_e > 0$. Joule heating then gives

$$p_{\mathrm{act}} = \sum_i R_e \|I_i\|^2 = \frac{R_e}{\chi^2}\|u\|^2 = c_A k^2 \|Lz\|^2, \qquad c_A = R_e/\chi^2. \tag{3.2}$$

Equation (3.2) is derived from the specified transducer and not from the oscillator equation. A chemical, neural, mechanical or regenerative actuator could have a different expenditure law, so the electrical realization is an engineering analogue and not a mechanistic identification of ion-pump energetics.

3.4 The Full Stuart–Landau Network under Actuator Accounting

Let $z \in \mathbb{C}^N$ and

$$\dot{z}_i = (a(E) + \mathrm{i}\omega_i - |z_i|^2)z_i - k(Lz)_i, \qquad \dot{E} = q - b - \frac{c_{\mathrm{osc}}}{N}\|z\|^2 - c_A k^2 \|Lz\|^2, \tag{3.3}$$

with $a(E) \le a_{\max}$ and $E(0) = E_0 > E_*$. The separate oscillator maintenance term $c_{\mathrm{osc}}\|z\|^2/N$ isolates the energetic consequences of actuation and is not asserted to be calibrated.

Proposition 3.5 (Bounded amplitude before exit)**.** *With $S = \|z\|^2$, $\dot{S} \le 2a_{\max}S - 2S^2/N$, hence $S(t) \le \max\{S(0), N\max(a_{\max}, 0)\}$ until the first energy-floor exit.*

Proof. Twice the real part of $\langle z, \dot{z}\rangle$ removes the frequency terms, $\langle z, Lz\rangle \ge 0$, and $\sum_i |z_i|^4 \ge S^2/N$. Scalar comparison closes the argument. □

Proposition 3.6 (Exact resource identity)**.** *For any solution on $[0, H]$, $E(t) = E_0 + (q-b)t - \int_0^t [\frac{c_{\mathrm{osc}}}{N}\|z\|^2 + c_A k^2 \|Lz\|^2]\, ds$. The trajectory is admissible on $[0, H]$ if and only if the right side is at least E_* for every $t \in [0, H]$. If $q \le b$ it suffices to test $t = H$.*

Proposition 3.7 (Heterogeneous frequencies obstruct exact synchrony)**.** *If $\omega_i \ne \omega_j$ for some pair and $z_1 = \cdots = z_N = \zeta \ne 0$ at an instant, then the synchronization manifold is not invariant.*

Proof. On the diagonal $Lz = 0$, so $\dot{z}_i - \dot{z}_j = \mathrm{i}(\omega_i - \omega_j)\zeta \neq 0$. □

These are respectively a genuine bound, an accounting identity and an obstruction, and none implies that cumulative cost is monotone in k along nonlinear trajectories, because $z(s)$ itself changes with k. A numerical protocol makes this concrete. On an eight-node cycle with $a(E) = 0.8 + 0.12\tanh((E-1)/0.8)$, $c_{\mathrm{osc}} = 0.012$, $q = b = 0.028$, $E_* = 0.35$, $H = 30$ and heterogeneous frequencies $(-0.16, -0.10, -0.05, -0.015, 0.015, 0.05, 0.10, 0.16)$, an ample-reserve scenario $(E_0, c_A) = (2.0, 0.045)$ and a scarce-reserve scenario $(0.9, 0.12)$ were integrated until the first downward crossing of E_*. The scarce-reserve rows display the intended tension: stronger

Table 3.2: Heterogeneous-frequency runs. A dash marks that the floor was crossed before the deadline $\tau = 8$, so deadline disagreement is not an admissible observable.

Scenario and k	Exit time	$\|Pz(8)\|$	Survives $H = 30$
Ample, $k = 0.30$	–	2.246	yes
Ample, $k = 1.00$	–	0.147	yes
Ample, $k = 4.00$	–	0.084	yes
Scarce, $k = 0.05$	–	2.342	yes
Scarce, $k = 0.30$	16.09	1.995	no
Scarce, $k = 0.60$	9.21	1.353	no
Scarce, $k = 1.00$	5.92	–	no
Scarce, $k = 2.00$	1.31	–	no
Scarce, $k = 4.00$	0.19	–	no

coupling reduces disagreement while advancing the exit. The ample-reserve rows show that every tested run survives and that final energy is not monotonically ordered in k, so the tension is parameter dependent and is not a universal law. A genuine theorem would have to control $\|Pz\|$ for heterogeneous frequencies, energy-dependent growth and the cubic term while respecting the stopping time, and Chapter 4 supplies that estimate.

Chapter 4

Transverse Disagreement and Global Certificates

4.1 Model and Stopping Convention

Consider the full network (3.3) with a locally Lipschitz, $a(E) \le a_{\max}$ on $E \ge E_*$, and $E(0) = E_0 > E_*$. Write $w = Pz$, $\bar\omega = N^{-1}\sum_i \omega_i$ and $\delta = \max_i |\omega_i - \bar\omega|$. All estimates hold up to $t_* = \inf\{t : E(t) < E_*\}$, and up to t_* inclusive when the solution extends continuously there. The common rotation $\mathrm{i}\bar\omega$ is removed by a rotating frame, which preserves all norms and costs.

Theorem 4.1 (Uniform amplitude comparison)**.** *With $S = \|z\|^2$ and*

$$B^2 = \max\{S(0), N\max(a_{\max}, 0)\},$$

one has $\|z(t)\| \le B$ for all $t < t_*$.

Proof. $\dot S = 2a(E)S - 2\sum_i |z_i|^4 - 2k\, z^*Lz \le 2a_{\max}S - 2S^2/N$, and scalar comparison gives the claim. □

4.2 Disagreement Contraction from Cubic Monotonicity

The map $g(\zeta) = |\zeta|^2\zeta$ on $\mathbb{C} \cong \mathbb{R}^2$ is the gradient of the convex function $|\zeta|^4/4$ and is therefore monotone, so $\mathrm{Re}\,(\bar\zeta - \bar\eta)(g(\zeta) - g(\eta)) \ge 0$. This replaces the destabilizing Lipschitz charge $3R^2$ of Proposition 2.11.

Theorem 4.2 (Nonlinear transverse bound)**.** *Let $\gamma = k\lambda_2 - a_{\max}$. The upper right Dini derivative of $\|w\|$ satisfies $D^+\|w\| \le -\gamma\|w\| + \delta\|z\|$. If $\gamma > 0$ then for $0 \le t < t_*$*

$$\|w(t)\| \le W(t) := e^{-\gamma t}\|w(0)\| + \frac{\delta B}{\gamma}(1 - e^{-\gamma t}). \tag{4.1}$$

Proof. Let $m = \mathbf{1}^\top z/N$, so $w_i = z_i - m$ and $\sum_i w_i = 0$. Then $\mathrm{Re}\sum_i \bar w_i g(z_i) = \mathrm{Re}\sum_i \bar w_i[g(z_i) - g(m)] \ge 0$. The common frequency contributes zero to the real inner product, $\mathrm{Re}\langle w, Lw\rangle \ge$

$\lambda_2\|w\|^2$, and $\|P\,\mathrm{diag}(\omega_i-\bar{\omega})z\| \le \delta\|z\| \le \delta B$. Therefore $\frac{1}{2}\frac{d}{dt}\|w\|^2 \le -\gamma\|w\|^2 + \delta B\|w\|$. Dividing by $\|w\|$ where positive and using the Dini derivative at zero, scalar comparison yields (4.1). □

When $\delta = 0$ the synchronized diagonal is invariant and $\gamma > 0$ gives exponential decay of disagreement. When $\delta > 0$ exact complex synchrony fails to be invariant, and the theorem guarantees only a bounded disagreement floor $\delta B/\gamma$. The condition $\gamma > 0$ is sufficient and not necessary. The mismatch term is bounded conservatively by δB, which is the first place where a graph-aware forcing metric (Chapter 6) can do better.

4.3 Exact and Conservative Resource Criteria

Theorem 4.3 (Exact criterion and trajectory-independent certificate)**.** *For any $H > 0$ on which the solution exists, admissibility on $[0, H]$ is equivalent to*

$$(q-b)t - \int_0^t \Big(\tfrac{c_{\mathrm{osc}}}{N}\|z\|^2 + c_A k^2 \|Lz\|^2\Big)ds \ge E_* - E_0 \quad \textit{for all } t \in [0,H].$$

A sufficient, trajectory-independent certificate is

$$E_0 - E_* \ge H\Big[(b-q)_+ + \tfrac{c_{\mathrm{osc}}}{N}B^2 + c_A k^2 \lambda_N^2 B^2\Big], \tag{4.2}$$

and if (4.2) holds strictly the floor cannot be crossed before H.

Proof. Integrate the energy equation for the identity. On any pre-exit interval $\|z\|^2 \le B^2$ and $\|Lz\|^2 \le \lambda_N^2 B^2$ while $(q-b)t \ge -(b-q)_+ H$, giving $E(t) \ge E_0 - H[\cdots]$. A strict inequality prevents a first crossing by continuity. □

Corollary 4.4 (Joint certificate)**.** *If $\gamma > 0$, (4.2) holds strictly, and $e^{-\gamma\tau}\|w(0)\| + \frac{\delta B}{\gamma}(1-e^{-\gamma\tau}) \le \varepsilon$ for a deadline $0 < \tau \le H$, then the trajectory is operational on $[0, H]$ and reaches disagreement at most ε by τ.*

The certificate is deliberately loose, since it charges worst-case amplitude and worst-case Laplacian actuation for the whole horizon. Its failure does not imply failure of viability, and it cannot be used to claim that increasing k necessarily shrinks the true viability kernel.

4.4 The Decay-Sensitive Certificate

Because $Lz = Lw$, one has $\|Lz\| \le \lambda_N\|w\| \le \lambda_N W(t)$, so the actuator charge can follow the decaying envelope instead of its initial value.

Theorem 4.5 (Decay-sensitive finite-horizon certificate)**.** *Put $w_0 = \|w(0)\|$, $\beta_0 = \delta B/\gamma$ and*

$$J(H) = \beta_0^2 H + \frac{2\beta_0(w_0-\beta_0)}{\gamma}(1-e^{-\gamma H}) + \frac{(w_0-\beta_0)^2}{2\gamma}(1-e^{-2\gamma H}), \tag{4.3}$$

$$D(H) = H\Big[(b-q)_+ + \tfrac{c_{\mathrm{osc}}}{N}B^2\Big] + c_A k^2 \lambda_N^2 J(H). \tag{4.4}$$

If $E_0 - E_* > D(H)$ the floor cannot be reached on $[0, H]$.

Proof. Direct integration gives $\int_0^H W(t)^2 dt = J(H)$, and the energy balance before the first exit gives $E(t) \geq E_0 - (b-q)_+ t - \frac{c_{\mathrm{osc}}}{N} B^2 t - c_A k^2 \lambda_N^2 \int_0^t W^2 ds$, in which all terms subtracted are nonnegative and nondecreasing in t. □

Corollary 4.6 (Identical frequencies)**.** *If $\delta = 0$ then $\beta_0 = 0$ and $J(H) = \frac{w_0^2}{2\gamma}(1 - e^{-2\gamma H})$, so the actuation contribution is bounded independently of the horizon by $c_A k^2 \lambda_N^2 w_0^2/(2\gamma)$. The basal and maintenance terms still grow linearly in H.*

When $\delta > 0$ the envelope tends to β_0 and $J(H) = \beta_0^2 H + O(1)$, so the certificate contains a persistent mismatch-driven allowance. This is a sufficient bound and not proof that actual actuation stays positive. The expression $k^2 J(H)$ is not monotone in k in general, because γ and β_0 also depend on k, so the theorem does not show that larger k always worsens viability.

For a four-node cycle ($\lambda_2 = 2$, $\lambda_N = 4$) with $a_{\max} = 0.5$, $k = 1$, $B = 2$, $w_0 = 1$, $H = 20$, $c_A = 0.01$, $c_{\mathrm{osc}} = 0$, $q = b$ and $\delta = 0$, one has $\gamma = 1.5$. The constant-amplitude allowance of (4.2) is $c_A k^2 \lambda_N^2 B^2 H = 12.8$, whereas the decay-sensitive allowance is 0.05333. With a reserve margin $E_0 - E_* = 0.1$ the earlier certificate fails and the new one succeeds. This compares sufficient certificates and is neither a simulation of depletion nor a measurement of expenditure.

Chapter 5

Spectral and Edge-Resolved Certificates

The certificate of Theorem 4.5 converts a disagreement norm into an actuator power through the worst-case factor λ_N^2, which is lossy whenever disagreement lies in low graph modes. This chapter retains the graph eigenmodes, which gives an exact identity and a modewise certificate, and then retains edge structure, which recovers the monotonicity of the cubic term that the modewise approach discards.

5.1 Spectral Decomposition and an Exact Linear Benchmark

Let v_j be orthonormal eigenvectors of L and $\xi_j = v_j^\top z$ for $j \geq 2$. For the system (3.3) with $F_i(z, E) = (a(E) + \mathrm{i}\omega_i - |z_i|^2)z_i$,

$$\dot{\xi}_j = -k\lambda_j \xi_j + f_j(t), \quad f_j = v_j^\top F(z, E), \qquad \|Lz\|^2 = \sum_{j=2}^{N} \lambda_j^2 |\xi_j|^2, \tag{5.1}$$

and the last relation is an identity, not an inequality. Replacing it by $\lambda_N^2 \|w\|^2$ loses information whenever disagreement is concentrated in low-frequency modes.

For the linear consensus subsystem $\dot{w} = -kLw$ no forcing is present and the expression is exact.

Corollary 5.1 (Exact expenditure for linear consensus)**.** $C_A(H) = c_A k^2 \int_0^H \|Lw\|^2 dt = \frac{c_A k}{2} \sum_{j\geq 2} \lambda_j |\xi_j(0)|^2 (1 - e^{-2k\lambda_j H})$.

This exposes a subtlety: high-frequency modes decay faster but can incur greater total actuator expenditure, and as $H \to \infty$ each initial mode is weighted by λ_j, not by λ_N^2/λ_2. On an eight-node cycle with an initial disagreement formed from two Laplacian eigenmodes, the exact cost and the bound obtained from global contraction, $c_A k^2 \lambda_N^2 \|w(0)\|^2 (1 - e^{-2k\lambda_2 H})/(2k\lambda_2)$, differ by a factor between 32.8 at $H = 1$ and 36.9 at $H = 20$. The factor is not an artifact of nonlinearity, since the system is linear, and it quantifies how much the worst-case graph factor can cost even in the best case.

5.2 The Modewise Certificate

Assume that until first exit $\|z\| \le B$, $|a(E)| \le A$, $|\omega_i| \le \Omega$, and that $|f_j(t)| \le M_j$.

Theorem 5.2 (Modewise finite-horizon certificate)**.** *Put $r_j = k\lambda_j$, $\beta_j = M_j/r_j$, $h_j = |\xi_j(0)| - \beta_j$ and*

$$I_j(H) = \beta_j^2 H + \frac{2\beta_j h_j}{r_j}(1 - e^{-r_j H}) + \frac{h_j^2}{2r_j}(1 - e^{-2r_j H}). \tag{5.2}$$

Then up to first exit $\int_0^H \|Lz\|^2 dt \le \sum_{j\ge 2} \lambda_j^2 I_j(H)$, and if

$$E_0 - E_* > H\Big[(b-q)_+ + \tfrac{c_{\mathrm{osc}}}{N} B^2\Big] + c_A k^2 \sum_{j\ge 2} \lambda_j^2 I_j(H)$$

the floor cannot be reached on $[0, H]$.

Proof. Variation of constants gives $|\xi_j(t)| \le e^{-r_j t}|\xi_j(0)| + (1 - e^{-r_j t})\beta_j = \beta_j + h_j e^{-r_j t}$, which is nonnegative as a convex combination of nonnegative numbers. Squaring and integrating gives (5.2), the spectral identity sums the modes, and the energy balance finishes the argument. □

The generic choice $M_j = (A + \Omega + B^2)B$ is valid but can be extremely loose. Estimating M_j from an unverified simulated trajectory does not give a rigorous certificate, and a safe practical certificate takes the minimum of separately proved global and modewise bounds.

5.3 Time-Dependent Modal Forcing

The constant envelope can be replaced by one that decays with the disagreement. Rotate so that $\bar\omega = 0$ and let $\Delta\omega_i = \omega_i - \bar\omega$. Since $v_j^\top \mathbf{1} = 0$ the spatially constant part of $g(m\mathbf{1})$ vanishes in each transverse mode, and $|g(\zeta) - g(\eta)| \le 3B^2|\zeta - \eta|$ for $|\zeta|, |\eta| \le B$, whence

$$|f_j(t)| \le \big(|a(E)| + 3B^2\big)\|w(t)\| + \delta B. \tag{5.3}$$

With $C = A + 3B^2$ and the envelope $W(t)$ of (4.1), define

$$U_j(t) = e^{-r_j t}|\xi_j(0)| + \int_0^t e^{-r_j(t-s)}[CW(s) + \delta B]\, ds.$$

Theorem 5.3 (Decay-sensitive spectral certificate)**.** *$|\xi_j(t)| \le U_j(t)$, hence $\int_0^H c_A k^2 \|Lz\|^2 dt \le c_A k^2 \sum_{j\ge 2} \lambda_j^2 \int_0^H U_j^2 dt$, and the energy floor is not reached on $[0, H]$ if $E_0 - E_* > H[(b-q)_+ + c_{\mathrm{osc}} B^2/N] + c_A k^2 \sum_j \lambda_j^2 \int_0^H U_j^2 dt$.*

With $w_0 = \|w(0)\|$, $\beta_0 = \delta B/\gamma$, $F_0 = C\beta_0 + \delta B$ and $F_1 = C(w_0 - \beta_0)$, one has for $r_j \ne \gamma$

$$U_j(t) = |\xi_j(0)|e^{-r_j t} + \frac{F_0}{r_j}(1 - e^{-r_j t}) + \frac{F_1}{r_j - \gamma}(e^{-\gamma t} - e^{-r_j t}),$$

with the last fraction replaced by $F_1 t e^{-r_j t}$ when $r_j = \gamma$. For $\delta = 0$ every U_j decays to zero and the certificate has finite infinite-horizon actuation cost; for $\delta > 0$ the envelope approaches a

plateau and the cost grows at most linearly in H. The cubic monotonicity used for the aggregate envelope does *not* make each f_j dissipative, because the nonlinearity mixes Laplacian modes. Applying the scalar forcing bound $CW(t) + \delta B$ separately to every mode therefore duplicates the worst case across modes, and the certificate can be looser than the global one.

5.4 The Edge-Dissipative Certificate

The aggregate form of monotonicity can be retained without introducing modes by working with the edge energy $Q = \mathrm{Re}\langle z, Lz\rangle = \sum_{\{i,j\}} b_{ij}|z_i - z_j|^2$, counting each undirected edge once. Consider identical frequencies and

$$\dot{z} = a(t)z - g(z) - kLz, \qquad a(t) \le a_+, \quad \gamma = k\lambda_2 - a_+ > 0, \tag{5.4}$$

with actuator power $P_A = c_A k^2 \|Lz\|^2$.

Theorem 5.4 (Edge-dissipative expenditure certificate)**.** *Every classical solution of* (5.4) *satisfies*

$$Q(t) \le Q(0)e^{-2\gamma t}, \tag{5.5}$$

$$\int_0^H P_A\, dt \le \frac{c_A k Q(0)}{2}\Big[1 + \frac{a_+}{\gamma}(1 - e^{-2\gamma H})\Big]. \tag{5.6}$$

Proof. Monotonicity of g gives, edge by edge, $\mathrm{Re}\langle Lz, g(z)\rangle = \sum_{\{i,j\}} b_{ij}\,\mathrm{Re}[(\bar{z}_i - \bar{z}_j)(g(z_i) - g(z_j))] \ge 0$. Hence $\dot{Q}/2 \le a_+ Q - k\|Lz\|^2$, and since $\|Lz\|^2 \ge \lambda_2 Q$, Gronwall gives (5.5). Integrating the same inequality gives $k\int_0^H \|Lz\|^2 \le Q(0)/2 + a_+ \int_0^H Q$, and substituting (5.5) proves (5.6). □

Corollary 5.5 (Finite-horizon reserve)**.** *If $\dot{E} = q - b - p_{\mathrm{osc}} - P_A$ with $p_{\mathrm{osc}} \ge 0$ and $\int_0^H p_{\mathrm{osc}} \le C_{\mathrm{osc}}$, and if $E(0) - E_* > (b-q)_+ H + C_{\mathrm{osc}} + \frac{c_A k Q(0)}{2}[1 + \frac{a_+}{\gamma}(1 - e^{-2\gamma H})]$, then $E > E_*$ on $[0, H]$, provided the growth bound is valid up to any putative first exit.*

Proof. The right-hand cost certificate is nondecreasing in the horizon, so applying it on every prefix interval excludes a first crossing. □

5.5 Comparison of Certificates

On the eight-node cycle with $a_+ = 0.15$, $k = 1$, $c_A = 0.02$, $H = 20$, $Q(0) = 0.0311961328$ and margin $\gamma \approx 0.435786$, with the initial state composed of a common component and two eigenmodes, all certificates were evaluated analytically and compared with the trapezoidal integral of a tightly integrated nonlinear trajectory (adaptive Runge–Kutta, relative tolerance 2×10^{-10}). The ordering is instructive and runs against the expectation that finer spectral resolution should be sharper. The modewise certificate with a decay-sensitive envelope is about five times looser than the global transverse certificate, because the uniform nonlinear forcing envelope is charged independently to each mode. The edge-dissipative certificate respects the aggregate monotonicity of the cubic term, and its improvement to 1.40 times the observed cost comes from that structure and not from any fitting to the trajectory. The comparison covers one parameter configuration and is not evidence of a universal ordering.

Table 5.1: Cumulative actuator cost at $H = 20$. The observed value is a numerical quadrature and is not a validated enclosure; the others are analytical upper bounds.

Method	Cumulative cost	Ratio to observed
Observed nonlinear trajectory	0.000299164	1.00
Edge-dissipative certificate	0.000419340	1.40
Global transverse (Theorem 4.3 route)	0.012373052	41.36
Decay-sensitive modewise	0.062016320	207.30
Constant-forcing modewise	44.613057522	149125.58

5.6 A Certificate for Real Networks with Common Growth

A cleaner statement holds for real oscillators $\dot{x}_i = ax_i - x_i^3 - k(Lx)_i$ with $a \le a_{\max}$, $J' = c_A k^2 \|Lx\|^2$, $J(0) = 0$. Let $y = Px$ and $V = \|y\|^2/2$. The identity $\sum_i (x_i - \bar{x})x_i^3 = \frac{1}{N}\sum_{i<j}(x_i - x_j)(x_i^3 - x_j^3) \ge 0$ gives $V' \le 2a_{\max}V - ky^\top Ly$, so $V(t) \le V(0)e^{-2\gamma t}$ with $\gamma = k\lambda_2 - a_{\max} > 0$. Because $L^2 \preceq \lambda_N L$ as quadratic forms, $\|Lx\|^2 \le \lambda_N y^\top Ly$.

Proposition 5.6 (Edge-energy expenditure bound)**.** *For every $T \ge 0$,*

$$J(T) \le \frac{c_A k^2 \lambda_N \lambda_2 V(0)}{k\lambda_2 - a_{\max}},$$

and the direct exponential route gives $c_A k^2 \lambda_N^2 V(0)/(k\lambda_2 - a_{\max})$, which is never smaller when $\lambda_2 \le \lambda_N$. Certified lower and upper enclosures of λ_2 and λ_N may replace the eigenvalues, because the bound decreases in λ_2 for $a_{\max} > 0$.

Proof. Integrating V' gives $k\int_0^T y^\top Ly \le V(0) + 2a_{\max}\int_0^T V \le V(0)(1 + a_{\max}/\gamma) = V(0)k\lambda_2/\gamma$, and then $J \le c_A k^2 \lambda_N \int y^\top Ly/k$. □

For $x(0) = (0.45, 0.21, 0.39, 0.27)$, $V(0) = 9/500$, $k = 0.7$, $a_{\max} = 0.19$, $c_A = 0.1$, the path has $\lambda \in \{0, 2 - \sqrt{2}, 2, 2 + \sqrt{2}\}$ (rational enclosures $\lambda_2 \ge 0.585$, $\lambda_N \le 3.415$), the cycle $\{0, 2, 2, 4\}$ and the complete graph $\{0, 4, 4, 4\}$. All three networks certify $E(t) > 0.08$ for every $t \ge 0$. If the full

Table 5.2: Edge-energy certificates for four-node real networks, with resource floor $E(0) - J_{\text{upper}}$ for $E(0) = 0.1$ and $q - b \ge 0$.

Graph	Direct bound	Edge bound	Resource floor
Path P_4	0.046861	0.008028	0.091972
Cycle C_4	0.011663	0.005831	0.094169
Complete K_4	0.005407	0.005407	0.094593

model carries further nonnegative operating costs these must be bounded separately, since the displayed guarantee covers only basal demand and actuator expenditure.

Chapter 6

Persistent Frequency Heterogeneity and Network-Level Certificates

The central question motivating this work can be put in one sentence: when does dynamical coordination preserve the capacity for continued operation? Chapters 4 and 5 answered it for identical frequencies, where disagreement decays to zero and the actuator cost of correction is a finite transient. With distinct natural frequencies the synchronized manifold is not invariant, disagreement settles at a positive floor, and actuation continues indefinitely. The relevant certificates therefore have two parts, a transient buffer and a persistent rate, and viability for all time requires the supply to cover the rate while the reserve covers the buffer. This chapter assembles four network-level certificates for the full complex network

$$\dot{z} = (a + \mathrm{i}\Omega)z - g(z) - kLz, \qquad g(z)_i = |z_i|^2 z_i, \quad J' = c_A k^2 \|Lz\|^2, \quad E(t) = E_0 + (q-b)t - J(t), \tag{6.1}$$

with common growth a, $\Omega = \mathrm{diag}(\omega_i)$, $\tilde{\omega} = \omega - \bar{\omega}\mathbf{1}$ and $\delta = \max_i |\tilde{\omega}_i|$.

6.1 Forced Disagreement for the Full Complex Network

Proposition 6.1 (Bounded modulus)**.** *Let $R = \max\{\max_i |z_i(0)|, \sqrt{\max(a,0)}\}$. Then $|z_i(t)| \le R$ for all i and $t \ge 0$.*

Proof. At a site attaining the maximal modulus the real part of the diffusive coupling is nonpositive in the radial derivative, and the cubic damping dominates the growth once the modulus exceeds $\sqrt{\max(a,0)}$, so the maximum cannot cross R upward. □

Proposition 6.2 (Forced disagreement envelope)**.** *Put $y = \|Pz\|$, $\gamma = k\lambda_2 - a$ and $\mathcal{B} = \delta\sqrt{N}R$. Then $D^+ y \le -\gamma y + \mathcal{B}$, and if $\gamma > 0$, $y(t) \le r + (y_0 - r)e^{-\gamma t}$ with $r = \mathcal{B}/\gamma$. No phase locking is assumed.*

Proof. Cubic monotonicity makes the cubic contribution to $\frac{d}{dt}\|Pz\|^2$ nonpositive, the common rotation contributes zero, and the heterogeneous term is bounded by $\|\tilde{\Omega}z\|\, y \le \delta\sqrt{N}R\, y$. □

Proposition 6.3 (Persistent expenditure and viability). *Let $\mathcal{C} = c_A k^2 \lambda_N^2$. Since $\|Lz\| \le \lambda_N y$, $J' \le \mathcal{C} y^2$, and for $y_0 \ge r$ and $s_0 = y_0 - r$,*

$$J(t) \le \mathcal{C}\Big[r^2 t + \tfrac{2rs_0}{\gamma}(1 - e^{-\gamma t}) + \tfrac{s_0^2}{2\gamma}(1 - e^{-2\gamma t})\Big].$$

If $q - b \ge \mathcal{C} r^2$ then $E(t) \ge E_0 - \mathcal{C}[2rs_0/\gamma + s_0^2/(2\gamma)]$ for all $t \ge 0$. If $y_0 < r$ the envelope $y \le r$ gives $E(t) \ge E_0$ whenever $q - b \ge \mathcal{C} r^2$.

These are sufficient and not necessary. On four-node networks with $a = 0.19$, $k = 0.7$, $c_A = 0.1$, $E_0 = 0.1$, $q - b = 0.005$, real initial state $(0.45, 0.21, 0.39, 0.27)$, frequencies $(-0.03, -0.01, 0.01, 0.03)$, $R = 0.45$, and certified rational spectral enclosures, the path fails the supply-rate condition because its persistent ceiling 0.00864641 exceeds 0.005. This does not establish exhaustion. The cycle and complete graph certify $E(t) > 0.08$ for every $t \ge 0$. More

Table 6.1: Disagreement-route certificate (Proposition 6.3).

Graph	Rate upper	Transient buffer	Floor for all $t \ge 0$ (constant net supply)
Path	0.00864641	0.04853512	not certified
Cycle	0.00039037	0.01392210	0.08607790
Complete	0.00008390	0.00594827	0.09405173

connectivity improves the disagreement decay margin in these examples, but the factor λ_N^2 in the actuator cost can counteract that improvement in other graph families.

6.2 Graph-Resolved Forcing

The bound $\mathcal{B} = \delta\sqrt{N}R$ uses only the largest frequency offset. The spatial pattern of mismatch can be retained through a quadratic form. Write $y = Pz$ and let $a_+ \ge a$, $\gamma = k\lambda_2 - a_+ - \delta > 0$, $A = kL - (a_+ + \delta)P$, which is positive definite on the mean-zero subspace, and $h = \tilde{\omega}^\top A^\dagger \tilde{\omega}$.

Proposition 6.4 (Graph-resolved disagreement certificate). *For all $t \ge 0$, $\|Pz(t)\|^2 \le e^{-\gamma t}\|Pz(0)\|^2 + \frac{R^2 h}{\gamma}(1 - e^{-\gamma t})$, so the asymptotic ceiling is $R\sqrt{h/\gamma}$.*

Proof. Write $z = m\mathbf{1} + y$. The frequency term decomposes as $\mathrm{i}P\Omega y + \mathrm{i}m\tilde{\omega}$. Its first part contributes at most $\delta\|y\|^2$ to $\frac{d}{dt}\frac{1}{2}\|y\|^2$, and $|m| \le R$. Thus $\dot{V} \le -\langle y, Ay\rangle + R|\langle y, \mathrm{i}\tilde{\omega}\rangle|$ with $V = \|y\|^2/2$, and the weighted Young inequality $R|\langle y, \mathrm{i}\tilde{\omega}\rangle| \le \frac{1}{2}\langle y, Ay\rangle + \frac{1}{2}R^2 h$ gives $\dot{V} \le -\gamma V + R^2 h/2$. □

Table 6.2: Asymptotic disagreement ceilings (rounded upward to 40-bit dyadic rationals before display).

Graph	Original ceiling	Graph-resolved ceiling	Reduction
Path	0.12300683	0.10580117	13.99%
Cycle	0.02231405	0.01610273	27.84%
Complete	0.01034483	0.00780024	24.60%

Since $\|Lz\|^2 \le \lambda_N^2 \|Pz\|^2$, the persistent rate is $\mathcal{C}\, R^2 h/\gamma$. The graph-resolved ceiling is not guaranteed to improve every other correlation-preserving certificate, and its advantage depends on how the frequency pattern aligns with the graph modes. The pseudoinverse is a forcing metric, not an exact steady state.

6.3 The Heterogeneous Edge Certificate

The edge energy approach of Theorem 5.4 extends to heterogeneous frequencies with the cubic term still dissipative. Write $Q = \langle z, Lz\rangle$, $Y = \|Lz\|$, $P_A = c_A k^2 Y^2$, $a(t) \le a_+$, $\gamma = k\lambda_2 - a_+ > 0$, $\|z\| \le B = \max\{\|z(0)\|, \sqrt{N a_+}\}$.

Theorem 6.5 (Heterogeneous edge-dissipative certificate)**.** *Let* $h_0 = \sqrt{Q(0)}$, $r = \delta B\sqrt{\lambda_N}/\gamma$, $G(t) = r + (h_0 - r)e^{-\gamma t}$,

$$I_H = r^2 H + \tfrac{2r(h_0-r)}{\gamma}(1 - e^{-\gamma H}) + \tfrac{(h_0-r)^2}{2\gamma}(1 - e^{-2\gamma H}), \quad M_H = \tfrac{1}{2}Q(0) + a_+ I_H, \quad D_H = \delta B\sqrt{H}.$$

Then $Q(t) \le G(t)^2$ *and*

$$\int_0^H P_A\, dt \le \frac{c_A}{4}\Big(D_H + \sqrt{D_H^2 + 4kM_H}\Big)^2. \tag{6.2}$$

Proof. Edgewise monotonicity gives $\mathrm{Re}\langle Lz, g(z)\rangle \ge 0$. Removing the mean frequency and applying Cauchy–Schwarz, $\frac{1}{2}\dot{Q} \le a_+ Q - kY^2 + \delta\|z\|Y \le a_+ Q - kY^2 + \delta B Y$. From $Y^2 \ge \lambda_2 Q$ and $Y^2 \le \lambda_N Q$, $\frac{1}{2}\dot{Q} \le -\gamma Q + \delta B\sqrt{\lambda_N Q}$, so $\eta = \sqrt{Q}$ obeys $D^+\eta \le -\gamma\eta + \delta B\sqrt{\lambda_N}$ and $\eta \le G$, giving $\int_0^H Q \le I_H$. Integrating the first inequality and dropping the terminal $Q(H) \ge 0$ gives $kX^2 \le M_H + D_H X$ for $X = (\int_0^H Y^2)^{1/2}$, a quadratic inequality whose solution gives (6.2) after multiplication by $c_A k^2$. □

At $\delta = 0$ the theorem reduces to Theorem 5.4. A reserve corollary follows exactly as in Corollary 5.5. On the eight-node cycle of Chapter 5 with frequencies $\omega_j = s\cos(2\pi j/8)$ the certificate is evaluated against the integrated trajectory. The looseness grows with dispersion,

Table 6.3: Observed cost and heterogeneous edge certificate at $H = 20$, $a = 0.15$, $k = 1$, $c_A = 0.02$.

Dispersion s	Observed cost	Certified bound	Ratio
0.00	0.000299	0.000419	1.40
0.02	0.000375	0.001598	4.27
0.05	0.000772	0.006699	8.68
0.10	0.002149	0.024411	11.36
0.20	0.006863	0.094449	13.76

because B, λ_N and the Cauchy–Schwarz estimate suppress information about the evolving state, and in particular they ignore the correlation between the forcing and the edge disagreement that Chapter 7 retains.

6.4 Direct Spectral Certification

Both of the preceding routes convert a disagreement ceiling into an actuator ceiling through λ_N^2. A direct route works with the edge-weighted functional $W = \frac{1}{2} z^* L z$ and $\mathcal{Q} = \|Lz\|^2$ and avoids that conversion. Let $\tilde{\omega} = \omega - \bar{\omega}\mathbf{1}$, $\delta = \max|\tilde{\omega}_i|$, $R^2 \geq \max(a, 0)$ with $|z_i(0)| \leq R$, $\underline{\lambda}_2$ a certified lower bound on the spectral gap, and

$$\beta = k - \frac{a_+ + \delta}{\underline{\lambda}_2}, \qquad F = R\|\tilde{\omega}\|.$$

Proposition 6.6 (Direct expenditure certificate)**.** *If $\beta > 0$ then $\dot{W} \leq -\beta\mathcal{Q} + F\sqrt{\mathcal{Q}}$ and, for every $T \geq 0$,*

$$J(T) = c_A k^2 \int_0^T \mathcal{Q}\, dt \leq \frac{2c_A k^2 W(0)}{\beta} + \frac{c_A k^2 F^2}{\beta^2} T. \tag{6.3}$$

Proof. Differentiating, $\dot{W} = \mathrm{Re}\langle Lz, \dot{z}\rangle$. The growth term is $a\, z^* L z \leq (a_+/\underline{\lambda}_2)\mathcal{Q}$, because $\mathcal{Q} = \sum_j \lambda_j^2 |\xi_j|^2 \geq \underline{\lambda}_2 \sum_j \lambda_j |\xi_j|^2$. The coupling gives $-k\mathcal{Q}$. The cubic term is nonpositive by edgewise monotonicity. For the frequency term, subtract the mean frequency, write $z = Pz + m\mathbf{1}$, and use $\|Pz\| \leq \|Lz\|/\underline{\lambda}_2$: the part with Pz is at most $(\delta/\underline{\lambda}_2)\mathcal{Q}$, and the part with the mean is at most $R\|\tilde{\omega}\|\sqrt{\mathcal{Q}}$. Young's inequality $F\sqrt{\mathcal{Q}} \leq \frac{\beta}{2}\mathcal{Q} + F^2/(2\beta)$ and $W(T) \geq 0$ give (6.3). □

Corollary 6.7 (Resource floor)**.** *If $\dot{E} = s - c_A k^2 \mathcal{Q}$ with constant net supply $s = q - b$ and $s \geq c_A k^2 F^2/\beta^2$, then $E(t) \geq E(0) - 2c_A k^2 W(0)/\beta$ for all $t \geq 0$. This is a sufficient floor and does not assert that the actuator power converges to the displayed rate.*

Table 6.4: Direct spectral certificate on the four-node examples ($n = 4$, $a = 0.19$, $k = 0.7$, $c_A = 0.1$, $R = 0.45$, $s = 0.005$, $E(0) = 0.1$; spectral gap lower bounds 0.585, 2, 4). The prior rate is the graph-resolved estimate of the preceding section.

Topology	Transient buffer	Direct rate	Prior rate	Energy floor
Path	0.015792	0.0001891	0.0063967	0.084208
Cycle	0.011361	0.0000570	0.0002033	0.088639
Complete	0.010940	0.0000477	0.0000477	0.089060

The direct rate is lower than the prior rate by a factor of about 34 on the path and 3.6 on the cycle, and the two coincide on the complete graph, where all nonzero eigenvalues are equal and the conversion through λ_N^2 loses nothing. The transient buffer is conservative because the Young allocation used to obtain (6.3) splits the dissipation equally between the two terms. Chapter 8 removes that restriction.

Chapter 7

Signed Frequency Work and Prospective Enclosures

The heterogeneous edge certificate bounds the frequency forcing and the edge disagreement independently, which discards the sign of their correlation. This chapter keeps that sign. An exact balance identity isolates a signed frequency-work integral, and the question becomes how to bound that integral prospectively, from information available at the start of the horizon or at the latest checkpoint, and not retrospectively from a completed trajectory.

7.1 Exact Edge-Energy Accounting

For $\dot{z} = (a(t)I + \mathrm{i}\Omega)z - g(z) - kLz$ define $Q = z^*Lz$, $Y = \|Lz\|$, $T = \mathrm{Re}\langle Lz, g(z)\rangle \geq 0$ and the signed frequency work density $F = \mathrm{Re}\langle Lz, \mathrm{i}\Omega z\rangle$. Then

$$\tfrac{1}{2}\dot{Q} = a(t)Q - T - kY^2 + F, \tag{7.1}$$

and integrating gives the exact accounting identity for $C_A(H) = c_A k^2 \int_0^H Y^2 dt$:

$$\frac{C_A(H)}{c_A k} = \frac{Q(0) - Q(H)}{2} + \int_0^H a(t)Q\,dt - \int_0^H T\,dt + J_F(H), \qquad J_F(H) = \int_0^H F\,dt. \tag{7.2}$$

The frequency term has no fixed sign, since it depends on the state-dependent correlation between frequency offsets and edge differences, and adding a constant to every ω_j leaves F unchanged. Because L and Ω are real symmetric, $F = z^*Kz$ with the Hermitian matrix $K = \frac{\mathrm{i}}{2}[L, \Omega]$, and $\kappa = \|K\|_{\mathrm{op}}$ is a graph–frequency contrast that vanishes for identical frequencies.

7.2 Conditional and A Priori Certificates

Let $\gamma = k\lambda_2 - a_+ > 0$, $B = \max\{\|z(0)\|, \sqrt{Na_+}\}$, and let I_H be the integral envelope of Theorem 6.5.

Theorem 7.1 (Signed-work certificate)**.** *If an independently justified U_H satisfies $J_F(H) \leq U_H$, then $C_A(H) \leq c_A k[Q(0)/2 + a_+ I_H + U_H]$. For prefix-wise viability it suffices to replace U_H by*

$\max\{0, \sup_{t \le H} U(t)\}$ *for a nondecreasing envelope* $U \ge J_F$.

Proof. Drop the nonpositive terms $-Q(H)/2$ and $-\int T$ from (7.2), use $a \le a_+$, and bound $\int Q \le I_H$. □

Taking $U_H = J_F(H)$ from a completed numerical trajectory gives a valid *a posteriori* bound that cannot establish predictive viability, because it presupposes the future. A fully a priori alternative follows from $|F| \le \kappa B^2 =: \eta$, giving $U_H = \eta H$, which can be looser than the earlier estimate and must be combined with it by taking the minimum. At $s = 0.20$ the signed-work

Table 7.1: Frequency-dispersion study on the eight-node cycle ($H = 20$). The signed-work column uses the observed J_F and is therefore retrospective; the commutator column is a priori.

s	Actual cost	Prior bound	Signed-work bound	Commutator bound
0.00	0.000299	0.000419	0.000419	0.000419
0.02	0.000375	0.001598	0.001122	0.006239
0.05	0.000772	0.006699	0.004352	0.016853
0.10	0.002149	0.024411	0.015477	0.039564
0.20	0.006863	0.094449	0.058325	0.103819

bound reduces the prior estimate by about 38% and remains about 8.5 times the observed expenditure. The remaining looseness comes from I_H, which charges a worst-case amplitude and spectral envelope even when J_F is known accurately.

7.3 A Prospective Envelope from the Initial State

Suppose $\|z\| \le B$ on $[0, H]$ and $|a(t)| \le a_{\mathrm{abs}}$. Since $\|g(z)\| \le \|z\|^3$ the speed obeys $\|\dot z\| \le M := (a_{\mathrm{abs}} + \|\Omega\| + k\lambda_N)B + B^3$.

Theorem 7.2 (Initial-state frequency-work envelope)**.** *With* $F_0 = z(0)^* K z(0)$ *and* $\eta = \kappa B^2$*, for every* $t \in [0, H]$*,*

$$F(t) \le f_{\mathrm{up}}(t) := \min\{\eta,\ F_0 + 2\kappa BMt\}, \qquad J_F(H) \le U(H) = \int_0^H f_{\mathrm{up}} \le \eta H.$$

With $\tau = \min\{H, \max\{0, (\eta - F_0)/(2\kappa BM)\}\}$ *one has the closed form* $U(H) = F_0\tau + \kappa BM\tau^2 + \eta(H - \tau)$.

Proof. The quadratic form satisfies $|F| \le \kappa\|z\|^2 \le \eta$, and $\dot F = 2\,\mathrm{Re}\langle \dot z, Kz\rangle$ gives $|\dot F| \le 2\kappa BM$ almost everywhere. Integrate and split at the saturation time. □

Substituting into Theorem 7.1 gives a fully prospective actuator and reserve certificate. The envelope is no worse than the static commutator estimate by construction, but with the coarse global speed bound its gain vanishes quickly on long horizons.

Table 7.2: Predictive envelope at $s = 0.20$ against the integrated signed work and the static cap ηH.

H	Actual J_F	Predictive $U(H)$	Static ηH	Reduction
0.1	−0.000143	0.006953	0.012989	46.5%
0.5	0.000283	0.058908	0.064944	9.3%
2.0	0.010668	0.253738	0.259774	2.3%
20.0	0.323035	2.591710	2.597740	0.2%

7.4 Checkpoint-Restarted Enclosures

A sharper prospective bound restarts the enclosure at certified checkpoints. Fix a checkpoint t_j, a state estimate $\hat{z}_j$ with $\|z(t_j) - \hat{z}_j\| \le \varepsilon_j$, and a radius R_j for the ball containing the true trajectory and the comparison center on $[t_j, t_j + h_j]$. Let $\ell_j = |a| + \|\Omega\| + k\lambda_N + 3R_j^2$ (a Lipschitz constant of the vector field on the ball) and $v_j = \|f(\hat{z}_j)\|$.

Theorem 7.3 (Local signed-work enclosure)**.** *For $0 \le \tau \le h_j$ let $\rho_j(\tau) = e^{\ell_j \tau}\varepsilon_j + \frac{v_j}{\ell_j}(e^{\ell_j \tau} - 1)$. Then*

$$F(z(t_j + \tau)) \le u_j(\tau) := \min\Big\{\eta,\ F(\hat{z}_j) + \kappa\big(2\|\hat{z}_j\|\rho_j(\tau) + \rho_j(\tau)^2\big)\Big\},$$

and if the checkpoint intervals partition $[0, H]$, $J_F(H) \le U_{\mathrm{loc}}(H) = \sum_j \int_0^{h_j} u_j(\tau)d\tau$.

Proof. With $d(t) = \|z(t_j + t) - \hat{z}_j\|$ the integral equation gives $D^+ d \le \ell_j d + v_j$ and $d(0) \le \varepsilon_j$, so Gronwall gives $d \le \rho_j$. For Hermitian K, $|z^* K z - \hat{z}^* K \hat{z}| \le \kappa(2\|\hat{z}\|\rho + \rho^2)$. The static cap η supplies the minimum. □

Corollary 7.4 (Reserve with prefix supremum)**.** *Let I_Q be a certified nondecreasing bound on $\int_0^t Q$. A sufficient reserve for every prefix of $[0, H]$ is*

$$E(0) - E_* \ge H\Big[(b - q)_+ + \tfrac{c_{\mathrm{osc}}}{N}B^2\Big] + c_A k\Big[\tfrac{Q(0)}{2} + a_+ I_Q(H) + \max\{0, \sup_{t \le H} U_{\mathrm{loc}}(t)\}\Big].$$

The supremum over prefixes is essential: each local envelope is nondecreasing within its interval, but the envelope resets after a new checkpoint, so the final signed-work bound alone need not control intermediate reserve minima.

The nonlinear benchmark at $s = 0.20$ with exact simulated checkpoints every 0.1 time units and $R_j = B$ gives, against the integrated signed work J_F, a local bound exceeding it by about 1.09% at $H = 20$. These are *online-conditional* predictions: the checkpoints are exact simulated

Table 7.3: Local enclosure against static and initial-state caps ($s = 0.20$, checkpoint spacing 0.1, exact checkpoints).

H	Actual work	Static cap	Initial-state cap	Reset speed	Local bound
0.1	−0.000143	0.012989	0.006953	0.006954	0.000080
0.5	0.000283	0.064944	0.058908	0.035721	0.001169
2.0	0.010668	0.259774	0.253738	0.151548	0.012867
20.0	0.323035	2.597741	2.591706	1.698816	0.326550

states, and a deployed controller would need verified state-estimation error radii and parameter bounds. The local bound is also not an actuator-energy certificate until it is combined with a valid I_Q.

7.5 Degradation under Uncertain Observation and Parameters

Two experiments probe the cost of imperfect information. The first sweeps checkpoint spacing $\{0.05, 0.1, 0.25, 0.5, 1\}$ against certified radius $\{0, 0.001, 0.005, 0.01, 0.025, 0.05\}$ (thirty combinations) on the eight-node cycle with exact centers and an imposed allowance. The actual signed work is 0.323035 and the static cap is 2.597741. This tests the conservatism of a certified

Table 7.4: Selected combinations of the thirty-case sweep.

Spacing	Radius	Local bound	Bound/actual
0.05	0	0.324448	1.004
0.05	0.01	0.374153	1.158
0.10	0.01	0.388549	1.203
0.25	0.01	0.471625	1.460
0.50	0.01	1.008342	3.121
1.00	0.01	1.808446	5.598

allowance and not the accuracy of a state estimator, since the centers are not perturbed. The second experiment perturbs the reported centers by 0.85ϵ in a changing complex direction and, in an uncertain-parameter case, displaces the generating parameters by $0.7d_a$, $0.65d_k$ and a frequency perturbation within d_ω, with budgets $(d_a, d_\omega, d_k) = (0.01, 0.02, 0.02)$. The enclosure uses $m = B(d_a + d_\omega + d_k\lambda_{\max})$ as an additional forcing in the Gronwall step and an extra allowance $\lambda_{\max} d_\omega B^2$ in the commutator. Measurement displacement alone costs little at short

Table 7.5: Perturbed checkpoints and uncertain parameters ($\Delta t = 0.1$, $H = 20$, sixty combinations evaluated, all enclosures exceed the actual work).

Condition	Actual work	Upper bound	Ratio
$\epsilon = 0.005$, exact parameters	0.323035	0.358343	1.11
$\epsilon = 0.010$, exact parameters	0.323035	0.390543	1.21
$\epsilon = 0.005$, uncertain parameters	0.357951	2.463582	6.88
$\epsilon = 0.010$, uncertain parameters	0.357951	2.496300	6.97

spacing, whereas parameter uncertainty is far more damaging because the disturbance estimate allows worst-case mismatch throughout each interval. A parameter-correlated enclosure would be much sharper, which motivates Chapters 10 and 11. All budgets here are imposed and the work variable is not a calibrated physiological energy.

Chapter 8

Optimal Allocation of Dissipation and Calibration against Trajectories

The direct spectral certificate (6.3) was obtained from a Young inequality with a fixed split, and nothing suggests that an equal split between the transient and persistent terms is best for a given reserve or supply rate. This chapter introduces a tunable allocation, rederives the resulting bound from first principles, optimizes it for the two resource questions that matter (the all-time floor under constant supply and the finite-horizon bound), and then calibrates the whole family of network certificates against numerically integrated nonlinear trajectories, including deliberately difficult initial conditions.

8.1 A Tunable Young Allocation

Retain the setting of Proposition 6.6: $\beta = k - (a_+ + \delta)/\underline{\lambda}_2 > 0$, $F = R\|\tilde{\omega}\|$, $W = \frac{1}{2} z^* L z$ and $\mathcal{Q} = \|Lz\|^2$, so that $\dot{W} \le -\beta\mathcal{Q} + F\sqrt{\mathcal{Q}}$. Define the two coefficients

$$\alpha = \frac{c_A k^2 W(0)}{\beta}, \qquad \rho = \frac{c_A k^2 F^2}{\beta^2}, \tag{8.1}$$

so that the equal-split certificate of Proposition 6.6 reads $J(T) \le 2\alpha + \rho T$.

Theorem 8.1 (Young-parameterized expenditure certificate)**.** *For every $\theta \in (0,1)$ and $T \ge 0$,*

$$J(T) \le \frac{\alpha}{1-\theta} + \frac{\rho T}{4\theta(1-\theta)}. \tag{8.2}$$

Proof. For $x, y \ge 0$ and $s > 0$ one has $xy \le sx^2 + y^2/(4s)$. Apply it with $x = \sqrt{\mathcal{Q}}$, $y = F$ and $s = \theta\beta$ to obtain $F\sqrt{\mathcal{Q}} \le \theta\beta\mathcal{Q} + F^2/(4\theta\beta)$. Then $\dot{W} \le -(1-\theta)\beta\mathcal{Q} + F^2/(4\theta\beta)$. Integrating over $[0,T]$ and discarding $W(T) \ge 0$ gives $(1-\theta)\beta\int_0^T \mathcal{Q} \le W(0) + F^2 T/(4\theta\beta)$. Dividing by $(1-\theta)\beta$ and multiplying by $c_A k^2$ yields (8.2). □

At $\theta = \frac{1}{2}$ the bound is $2\alpha + \rho T$, which reproduces (6.3) exactly, so the formula carries no hidden missing factor. The persistent coefficient $\rho/(4\theta(1-\theta))$ is minimized at $\theta = \frac{1}{2}$, where it equals ρ, while the transient coefficient $\alpha/(1-\theta)$ increases with θ. Hence a smaller θ trades a

larger persistent rate for a smaller transient buffer, and the useful question is when that trade is worthwhile.

8.2 Optimal Choices of the Parameter

Proposition 8.2 (Supply-matched allocation for the constant-supply floor)**.** *Let the net supply be constant, $s = q - b$, with $\dot{E} = s - c_A k^2 \mathcal{Q}$. For $\theta \in (0,1)$ the floor guarantee $E(t) \geq E(0) - \alpha/(1-\theta)$ for all $t \geq 0$ holds whenever $s \geq \rho/(4\theta(1-\theta))$. Feasible θ exist if and only if $s \geq \rho$, they form the interval $[\theta_s, 1-\theta_s]$ with*

$$\theta_s = \tfrac{1}{2}\Big(1 - \sqrt{1 - \rho/s}\Big),$$

and the best guaranteed buffer is

$$\frac{\alpha}{1-\theta_s} = \frac{2\alpha}{1+\sqrt{1-\rho/s}} \in (\alpha, 2\alpha].$$

Proof. By (8.2), $E(t) = E(0) + st - J(t) \geq E(0) - \frac{\alpha}{1-\theta} + \big(s - \frac{\rho}{4\theta(1-\theta)}\big)t$, and the last bracket is nonnegative under the stated condition. The condition $4\theta(1-\theta) \geq \rho/s$ is a quadratic inequality in θ whose solution set is $[\theta_s, 1-\theta_s]$ and is nonempty exactly when $\rho/s \leq 1$. The buffer $\alpha/(1-\theta)$ is increasing in θ, so it is minimized at θ_s. □

The improvement over the equal split is therefore at most a factor of two, approached as the supply grows without bound relative to the persistent rate, and it vanishes when $s = \rho$. For any real supply margin the saving is substantial: when $s \gg \rho$ the optimal buffer is close to α, half of the equal-split value 2α.

Proposition 8.3 (Finite-horizon optimum)**.** *Put $\kappa = \rho T/4$. For fixed $T > 0$ the bound (8.2) equals $(\alpha+\kappa)/(1-\theta) + \kappa/\theta$ and is minimized at*

$$\theta^*(T) = \frac{\sqrt{\kappa}}{\sqrt{\kappa} + \sqrt{\alpha+\kappa}}, \qquad J(T) \leq \big(\sqrt{\alpha+\kappa} + \sqrt{\kappa}\big)^2.$$

This is never larger than the equal-split value $2\alpha + 4\kappa$, with equality if and only if $\alpha = 0$.

Proof. Write $\alpha/(1-\theta) + \kappa/(\theta(1-\theta)) = \alpha/(1-\theta) + \kappa[1/\theta + 1/(1-\theta)] = (\alpha+\kappa)/(1-\theta) + \kappa/\theta$. Setting the derivative of $c/(1-\theta) + \kappa/\theta$ to zero gives $c/(1-\theta)^2 = \kappa/\theta^2$, hence $\theta/(1-\theta) = \sqrt{\kappa/c}$ for $c = \alpha+\kappa$, which is the stated θ^*, and the value is $c/(1-\theta^*) + \kappa/\theta^* = (\sqrt{c}+\sqrt{\kappa})^2$. By the arithmetic–geometric mean inequality $\sqrt{\kappa(\alpha+\kappa)} \leq \kappa + \alpha/2$, so $(\sqrt{\alpha+\kappa} + \sqrt{\kappa})^2 = \alpha + 2\kappa + 2\sqrt{\kappa(\alpha+\kappa)} \leq 2\alpha + 4\kappa$ with equality only when $\alpha = 0$. □

On short horizons ($\kappa \ll \alpha$) the bound tends to α, a factor of two below the equal split, and the optimal parameter tends to zero. On long horizons ($\kappa \gg \alpha$) one has $(\sqrt{\alpha+\kappa} + \sqrt{\kappa})^2 \approx 4\kappa + 2\alpha$, which coincides with the equal split to leading order, and $\theta^* \to \frac{1}{2}$. The allocation therefore matters for the transient and does not touch the persistent rate, which is the quantity limiting long-horizon performance.

8.3 Numerical Results for the Four-Node Networks

The experiment uses the data of the four-node examples of Chapter 6: $a = 0.19$, $k = 0.7$, $c_A = 0.1$, $R = 0.45$, $\omega = (-0.03, -0.01, 0.01, 0.03)$, $z(0) = (0.45, 0.21, 0.39, 0.27)$, $s = 0.005$, $E(0) = 0.1$ and the certified spectral-gap lower bounds 0.585, 2, 4. Rational arithmetic is used throughout, with square roots rounded upward at 2^{-80}, and the equal-split baseline reproduces the buffers 0.015792, 0.011361, 0.010940 and the rates 0.0001891, 0.0000570, 0.0000477 of Chapter 6. Because the

Table 8.1: Supply-matched allocation (Proposition 8.2) for $s = 0.005$.

Graph	β	θ_s	Buffer, $\theta = \frac{1}{2}$	Buffer, optimized	Ratio	Floor, $\theta = \frac{1}{2}$	Floor, optimized
Path	0.32393	0.009547	0.015792	0.007972	0.5048	0.084208	0.092028
Cycle	0.59000	0.002859	0.011361	0.005697	0.5014	0.088639	0.094303
Complete	0.64500	0.002391	0.010940	0.005483	0.5012	0.089060	0.094517

supply exceeds the persistent rates by more than an order of magnitude, the optimized buffers are within one percent of the theoretical limit α, which is half the equal-split value. The resource floor on the path rises from 0.0842 to 0.0920. The optimal parameter increases from near zero at

Table 8.2: Finite-horizon bounds (Proposition 8.3) against the equal split.

Graph	T	$\theta^*(T)$	Optimized bound	Equal-split bound
Path	1	0.0716	0.009216	0.015981
Path	5	0.1457	0.011142	0.016738
Path	20	0.2464	0.015570	0.019575
Path	100	0.3796	0.032804	0.034705
Cycle	1	0.0476	0.006279	0.011418
Cycle	5	0.1002	0.007104	0.011646
Cycle	20	0.1794	0.008859	0.012502
Cycle	100	0.3093	0.014896	0.017062
Complete	1	0.0446	0.006005	0.010987
Complete	5	0.0941	0.006737	0.011178
Complete	20	0.1697	0.008280	0.011894
Complete	100	0.2973	0.013492	0.015710

$T = 1$ toward $\frac{1}{2}$ as the horizon grows, as the asymptotics predicted, and the optimized bound recomputed at θ^* agrees with the closed form to machine precision. The gain falls from 42 to 45% at $T = 1$ to about 5 to 14% at $T = 100$.

8.4 Calibration against Integrated Trajectories

Conservatism can be measured only against actual expenditure. The calibration study integrates the full nonlinear complex system (6.1) together with $J' = c_A k^2 \|Lz\|^2$ (adaptive eighth-order Runge–Kutta of Dormand–Prince type, relative tolerance 10^{-11}) on the three four-node graphs, at the reference frequencies and at three times their dispersion, from nine initial states: the reference real state, the same moduli with prescribed phases, three random states with moduli at most 0.45, an alternating state $(0.45, -0.45, 0.45, -0.45)$ that places maximal disagreement

in the highest graph mode, a state along the Fiedler direction scaled to modulus 0.45, a near-synchronized state, and a small-amplitude state with all moduli 0.01. Expenditure was recorded at $T \in \{1, 5, 20, 100\}$. For each case three certificates were evaluated with the same amplitude bound $R = \max\{\max_i |z_i(0)|, \sqrt{a}\}$ and the same certified spectral enclosures: the equal-split direct bound, the optimized bound of Proposition 8.3, and the disagreement-route bound of Proposition 6.3. All 216 comparisons satisfy certificate $\geq$ actual expenditure, which is a consistency check of the analysis and not a proof. Four conclusions follow, each limited to this

Table 8.3: Median ratio of certificate to actual expenditure over the states with expenditure above 10^{-6}, at the reference frequencies. Equal-split / optimized.

Graph	$T = 1$	$T = 20$	$T = 100$
Path	4.6 / 2.4	4.1 / 2.6	4.2 / 3.2
Cycle	2.4 / 1.2	2.2 / 1.3	2.1 / 1.5
Complete	2.1 / 1.1	2.1 / 1.2	2.0 / 1.3

Table 8.4: Ratio of certificate to actual expenditure at $T = 20$, reference frequencies. For each graph: optimized direct certificate, then disagreement-route certificate.

	Path		Cycle		Complete	
Initial state	direct	route	direct	route	direct	route
Reference, real	3.9	55.2	1.7	4.2	1.5	1.4
Reference, with phases	2.6	25.6	1.3	4.1	1.2	1.1
Random 1	2.6	21.6	1.3	3.9	1.2	1.1
Random 2	2.5	32.1	1.3	4.3	1.2	1.1
Random 3	2.7	31.8	1.3	4.6	1.2	1.1
Alternating, maximal	2.4	16.2	1.3	2.4	1.1	1.1
Fiedler direction	2.4	71.6	1.3	4.9	1.2	1.1
Near-synchronized	6.1	278.4	1.6	10.8	1.3	2.3
Small amplitude	90.3	3971.1	21.0	130.1	16.0	26.1

configuration. First, on the cycle and the complete graph the optimized direct certificate lies within 1.1 to 1.7 of the actual expenditure for every well-excited state, so for these networks the certificate is already close to sharp, and the remaining gap is of the size the Young optimization can account for. Second, on the path the median ratio of 2.3 to 3.2 is larger, and at triple frequency dispersion the ratio at $T = 100$ rises to about 11 for the reference state, with the optimized and equal-split values nearly equal (10.98 against 10.99). When the two coincide the persistent term dominates, so the residual conservatism is not attributable to the allocation of dissipation. Third, the direct certificate is dramatically tighter than the disagreement route on the path (55.2 against 3.9 for the reference state), confirming that conversion through λ_N^2 is the dominant loss for graphs with spread spectrum, and the routes coincide on the complete graph. Fourth, the large ratios for the near-synchronized and small-amplitude states are caused by the amplitude bound and by frequency forcing taken at its worst case, not by the allocation. In the small-amplitude case the actual state is far inside the ball of radius $R = \sqrt{0.19}$ used by the certificate, and in the near-synchronized case the transient buffer is negligible, so equal-split and optimized values agree.

8.5 A Mode-Resolved Allocation

The Young allocation above treats the forcing as a single vector of norm F. A finer allocation lets each Laplacian mode keep its own dissipation margin. Let u_j be orthonormal eigenvectors for λ_j, $j \geq 2$, let $\tilde{\omega}_j = \langle u_j, \tilde{\omega} \rangle$, and set $y_j = \lambda_j |\xi_j|$ with $\xi_j = \langle u_j, z \rangle$, so that $\mathcal{Q} = \sum_j y_j^2$. Define the modal margins

$$\beta_j = k - \frac{a_+}{\lambda_j} - \frac{\delta}{\underline{\lambda}_2}, \qquad \beta_2 = \beta,$$

and assume $\beta_j > 0$ for all $j \geq 2$.

Theorem 8.4 (Mode-resolved direct certificate)**.** *For every $\beta_* \in (0, \min_j \beta_j)$ and every $T \geq 0$,*

$$J(T) \leq \frac{c_A k^2}{\beta_*} \Big(W(0) + T\, P(\beta_*) \Big), \qquad P(\beta_*) = \sum_{j \geq 2} \frac{R^2 \tilde{\omega}_j^2}{4(\beta_j - \beta_*)}.$$

Proof. The growth term is $a\, z^* L z = a \sum_j y_j^2 / \lambda_j \leq a_+ \sum_j y_j^2 / \lambda_j$, the coupling term is $-k \sum_j y_j^2$, the cubic term is nonpositive, and the part of the frequency work acting on the disagreement component is at most $(\delta / \underline{\lambda}_2) \sum_j y_j^2$ as in Proposition 6.6. The part acting on the mean m with $|m| \leq R$ equals $\mathrm{Re} \sum_j \lambda_j \bar{\xi}_j i m \tilde{\omega}_j \leq R \sum_j |\tilde{\omega}_j| y_j$. Hence $\dot{W} \leq -\sum_j \beta_j y_j^2 + R \sum_j |\tilde{\omega}_j| y_j$. Young's inequality applied mode by mode with weight $\beta_j - \beta_*$ gives $R|\tilde{\omega}_j| y_j \leq (\beta_j - \beta_*) y_j^2 + R^2 \tilde{\omega}_j^2 / (4(\beta_j - \beta_*))$, so $\dot{W} \leq -\beta_* \mathcal{Q} + P(\beta_*)$. Integrating and using $W(T) \geq 0$ gives $\beta_* \int_0^T \mathcal{Q} \leq W(0) + T P(\beta_*)$. □

When the forcing lies in a single mode, taking $\beta_* = (1 - \theta)\beta$ recovers the Young bound exactly. When the margins β_j coincide, as on the complete graph, no gain is possible. The gain arises when forcing weight sits in modes with larger eigenvalues and hence larger margins. The bound is minimized over β_* by a one-dimensional search, and every value of β_* gives a valid bound.

The comparison against the Young-optimized global bound used the same four-node graphs, three initial states each (aligned, alternating, random), horizons $T \in \{1, 5, 20, 100\}$, and four frequency patterns: the reference pattern at scales 1 and 3, and an alternating pattern ± 0.03 at scales 1 and 3. All 144 bounds of both kinds are at least the integrated expenditure. The

Table 8.5: Median reduction of the mode-resolved bound relative to the Young-optimized direct bound, in percent, over states and horizons.

Graph	Reference	Reference ×3	Alternating	Alternating ×3
Path	0.04	0.18	12.56	40.36
Cycle	0.87	1.36	8.09	9.62
Complete	0.00	0.00	0.00	0.00

reduction reaches 61.6% at its largest, on the path graph with the alternating pattern. For the smooth reference pattern, whose weight lies mostly in the Fiedler mode, the gain is under 3%, so the persistent slack seen in the calibration is not mostly a mode-allocation effect for that pattern. The result is therefore a statement about patterns that excite higher modes, and the

earlier remaining conservatism for smooth patterns lies in the global amplitude bound R and in the shared $\delta/\underline{\lambda}_2$ term, which this allocation leaves untouched.

8.6 A Time-Dependent Amplitude Envelope

The constant $R = \max\{\max_i |z_i(0)|, \sqrt{a_+}\}$ in the forcing term is conservative for any trajectory that starts away from the limit-cycle amplitude, and especially for a small initial state, since the mean m satisfies $|m| \le \max_i |z_i|$ and the maximum modulus is itself bounded by a decaying or growing envelope.

Proposition 8.5 (Amplitude envelope)**.** *Let $a > 0$ and $M_0 = \max_i |z_i(0)| > 0$, and set $c = 1 - a/M_0^2$. If $M_0 = 0$ the solution is identically zero, since $z = 0$ is an equilibrium of the system and solutions are unique, and the envelope is not needed. For every $t \ge 0$, $\max_i |z_i(t)|^2 \le R(t)^2 = a/(1 - ce^{-2at})$, and*

$$\int_0^T R(t)^2 \, dt = \tfrac{1}{2} \ln \frac{e^{2aT} - c}{1 - c}.$$

Proof. At a node of maximal modulus the diffusive coupling contributes $-k \sum_j (|z_i|^2 - \operatorname{Re} \bar{z}_i z_j)/|z_i| \le 0$ to the upper Dini derivative of $M = \max_i |z_i|$, so $\dot{M} \le aM - M^3$. The right side is the Bernoulli equation whose solution is $R(t)$, and comparison gives $M \le R$. The integral follows from $\frac{d}{dt} \ln(e^{2at} - c) = 2a/(1 - ce^{-2at})$. □

Substituting $R(t)$ for R in the proofs of Proposition 6.6, Theorem 8.1 and Theorem 8.4 changes only the integration of the forcing term, so $R^2 T$ is replaced by $\int_0^T R(t)^2 \, dt$ in every persistent coefficient, with all other quantities unchanged. The allocation parameter θ or β_* remains constant in time, and the closed-form optima of this chapter apply with $\kappa = c_A k^2 \|\tilde{\omega}\|^2 \int_0^T R^2/(4\beta^2)$.

The test set was the four-node graphs at frequency scales 1 and 3, five initial states with maximum modulus 0.45 (two states), 0.20, 0.05 and 0.01, and $T \in \{1, 5, 20, 100\}$. All 120 bounds of all four variants lie at or above the integrated expenditure. The gain is negligible for states

Table 8.6: Median reduction in percent of the Young-optimized direct bound when the constant R is replaced by the envelope $R(t)$, by initial maximum modulus.

Initial modulus	$T = 1$	$T = 5$	$T = 20$	$T = 100$
0.45	0.12	0.83	2.33	4.30
0.20	5.52	8.52	5.14	1.99
0.05	38.63	61.24	44.85	10.74
0.01	87.02	96.19	89.14	19.80

near the limit-cycle amplitude, which is the regime of the four-node reference examples, and it is large for small initial states over moderate horizons, where the constant R greatly overstates how much the mean can contribute early on. At $T = 100$ the envelope has converged to $\sqrt{a}$ and the gain falls. Combining the envelope with Theorem 8.4 reduces the bound by up to 99.4% in the most favourable test case. These reductions should not be read as typical, because the favourable cases combine small initial amplitude with forcing in higher modes.

8.7 What the Calibration Implies

The results sort the sources of conservatism. The Young allocation can recover at most a factor of two and only on the transient term, which is substantial for reserve sizing at short horizons but irrelevant to long-horizon performance. The conversion from disagreement to expenditure has been largely removed by the direct route. What remains, and what limits the path and the small-amplitude cases, is the persistent forcing coefficient ρ, which depends on the global amplitude bound R and the full norm $\|\tilde{\omega}\|$ and uses neither the actual excitation of the graph modes nor the realized amplitude. Further theoretical refinement is therefore worthwhile only on that term. Two directions follow from the data. A mode-resolved forcing certificate that weights each Laplacian eigenmode of the frequency pattern by the available dissipation in that mode, combining the graph-resolved metric of Proposition 6.4 with the direct edge-energy method, targets exactly the path case, where the spectral spread is large. A state-dependent amplitude bound, obtained from the transverse envelope of Theorem 4.2 and a decaying amplitude estimate in place of the uniform R, targets the small-amplitude and near-collapse cases. Both are stated as open problems in Chapter 13. A third direction, treating time-varying supply and a state-dependent coupling $k(E)$, would bring the model closer to the original resource-constrained problem and requires a Lyapunov argument in which β varies with the reserve, so the equal-split and optimized formulas would have to be recomputed along the trajectory of the controller.

Part III

Reduced Systems and Parameter Uncertainty

Chapter 9

Dissipation Certificates for the Two-Oscillator Reduction

For two identical real oscillators the disagreement equation decouples enough that exact dissipation arguments are available, and uniform-in-time expenditure bounds can be obtained without any numerical integration. This chapter derives them. The same reduced system is used in Chapters 11 and 12 as the testbed on which validated propagators are compared, and the dissipation bounds here serve as an analytic benchmark for them.

9.1 The Reduced System and an Invariant Region

Take $u' = au - u^3 - k(u-v)$, $v' = av - v^3 - k(v-u)$ and $J' = 2c_A k^2 (u-v)^2$, with $a \in [0.15, 0.19]$ constant along each trajectory, $k = 0.7$, $c_A = 0.1$. In collective and disagreement coordinates $p = (u+v)/2$ and $m = (u-v)/2$,

$$p' = ap - p^3 - 3pm^2, \qquad m' = -(2k - a + 3p^2 + m^2)\, m, \qquad J' = 8c_A k^2 m^2 =: Cm^2, \tag{9.1}$$

with $C = 0.392$. Initial uncertainty is $p(0) \in [0.329, 0.331]$, $m(0) \in [0.119, 0.121]$, $J(0) = 0$. The resource is $E' = q - b - J'$ with $E(0) = 0.1$, $q = 0.01$, $b = 0.005$ and floor $E_* = 0.08$.

Proposition 9.1 (Invariant half-strip)**.** *The region* $\{p \geq 0.325,\ 0 \leq m \leq 0.121\}$ *is forward invariant for every fixed* $a \in [0.15, 0.19]$.

Proof. At $p = 0.325$ one has $p'/p \geq 0.15 - 0.325^2 - 3(0.121)^2 = 0.000452 > 0$, at $m = 0.121$ the factor $2k - a + 3p^2 + m^2$ is positive so $m' < 0$, and $m = 0$ is invariant. The polynomial vector field is locally Lipschitz, and the dissipative cubic terms with bounded m prevent finite-time escape of p. □

On the strip, $D := 2k - a + 3p^2 + m^2 \geq D_0 := 1.4 - 0.19 + 3(0.325)^2 = 1.526875 = 2443/1600$.

9.2 The Joint Dissipation Certificate

Let $V = m^2$. Then $V' = -2DV$ and $J' = CV$, so the combined observable $\mathcal{W} = J + \frac{C}{2D_0}V$ satisfies $\mathcal{W}' = CV(1 - D/D_0) \le 0$.

Theorem 9.2 (Joint dissipation bound)**.** *For every $t \ge 0$ and every admissible a and initial state,*

$$J(t) + \tfrac{C}{2D_0}m(t)^2 \le \tfrac{C}{2D_0}m(0)^2,$$
$$\textit{hence} \qquad J(t) \le \frac{0.392\,(0.121)^2}{2 \cdot 1.526875} = \frac{102487}{54531250} \approx 0.001879417765.$$

The uniform energy bound follows at once: since the net supply $q - b = 0.005$ is nonnegative, $E(t) \ge 0.1 - J(\infty) \ge 0.098120582 > 0.08$ for all $t \ge 0$. This improves the best fixed-affine validated enclosure at $H = 2$ (0.002081170556, Chapter 12) by about 9.69% and, unlike that enclosure, holds for all time. The bound is conservative because it replaces the state-dependent dissipation $D(t)$ by its lower bound.

9.3 Time-Dependent Dissipation and Riccati Comparison

The constant D_0 ignores two facts: the collective variable p increases after a transient, and D contains m^2 itself. Both can be exploited with a comparison argument.

Proposition 9.3 (Later lower bound on p)**.** *For $t \ge 0.5$ and $p \in [0.325, 0.329]$ one has $p' \ge 0.325[0.15 - 0.329^2 - 3Me^{-D_0}] > 0.0104709591$, with $M = 0.121^2$. The minimum increase over the next half-unit exceeds 0.00523547, which is greater than the strip width 0.004, so a first-crossing argument gives $p(t) \ge 0.329$ for $t \ge 1$. Consequently $D \ge D_1 := 1.534723$ for $t \ge 1$.*

Since $D = (2k - a + 3p^2) + m^2$ with the bracket bounded below by D_0 or D_1, $V' \le -2(D_\bullet + V)V$, and comparison with the scalar Riccati equation gives

$$V_D(t; v) = \frac{Dve^{-2Dt}}{D + v(1 - e^{-2Dt})}, \qquad \int_0^\infty V_D\, dt = \tfrac{1}{2}\log(1 + v/D).$$

Theorem 9.4 (Riccati and piecewise-Riccati certificates)**.** *With $V_0 = 0.121^2$,*

$$J(\infty) \le \tfrac{C}{2}\log(1 + V_0/D_0) \quad \textit{(constant Riccati)},$$

and, with $v_1 = V_{D_0}(1; V_0)$,

$$J(\infty) \le \tfrac{C}{2}\Big[\log(1 + V_0/D_0) - \log(1 + v_1/D_0) + \log(1 + v_1/D_1)\Big] \quad \textit{(piecewise Riccati)}.$$

The two-stage exponential certificate $J(\infty) \le \frac{CM}{2}[(1 - e^{-2D_0})/D_0 + e^{-2D_0}/D_1]$ is also valid, and these bounds are separately valid and must not be added.

The successive refinements differ in the fourth significant digit, the largest gain coming from retaining m^2 inside the dissipation rate, and the piecewise Riccati sharpening improves the

Table 9.1: Uniform expenditure certificates and resulting energy floors for the two-oscillator reduction. All inequality checks are evaluated with exact fractions, with exponentials and logarithms bounded by alternating rational series.

Certificate	J upper bound	Uniform E floor
Constant D_0 (joint dissipation)	0.001879417765	0.098120582235
Two-stage $D(t)$	0.001878964319	0.098121035681
Constant Riccati	0.001870464213	0.098129535787
Piecewise Riccati	0.001870015073	0.098129984927

constant-Riccati bound by about 0.024%. These are certificates for the reduced real system with its assumed parameter range, and no biological calibration is implied. The piecewise bound inherits Propositions 9.1 and 9.3 as premises.

9.4 Role as a Benchmark

The analytic bound 0.00187 is very close to the true expenditure scale of this system, since $J(\infty) = C \int_0^\infty m^2 dt$ and m decays at rate at least $D_0 \approx 1.53$. It therefore provides an independent standard against which the finite-horizon enclosures of Chapters 11 and 12 can be judged: their upper bounds on $J(2)$ near 0.00208 are valid but about 11% above the all-time analytic ceiling 0.00187, so the true $J(2)$ lies below that ceiling and the excess measures the overestimation of the enclosure itself. This is a useful reminder that a validated numerical enclosure is a different kind of object from a certificate, and the two are complementary.

Chapter 10

Persistent Parameters and Set-Valued Inference

Every certificate of the preceding chapters assumed that the growth rate, coupling and frequency pattern are known, or bounded by a worst-case interval that may vary in time. In practice the parameters are fixed but uncertain, and repeated observation of the same trajectory should narrow them. Treating the uncertainty as an independent disturbance at each checkpoint discards exactly that persistence, and Chapter 7 showed the cost. This chapter formalizes persistent parameters and examines how much a sequence of noisy observations can certify, reporting a negative result alongside the positive ones.

10.1 Feasible Sets in the Product Space

Let $z \in \mathbb{C}^N$ and $\theta = (a, k, s) \in \Theta_0 \subset \mathbb{R}^3$ with $\dot{z} = f_\theta(z) = (a + \mathrm{i}s\,\mathrm{diag}(\nu))z - |z|^2 z - kLz$, where ν is a known frequency template and the parameter is constant for the whole experiment. Observations y_j satisfy $\|z(t_j) - y_j\| \le \epsilon_j$. The exact feasible set is a subset of the product space and not a product of independently estimated intervals:

$$\mathcal{S}_0 = \{(z, \theta) : z \in Z_0,\ \theta \in \Theta_0,\ \|z - y_0\| \le \epsilon_0\}, \tag{10.1}$$

$$\mathcal{S}_{j+1} = \{(\Phi^\theta_{h_j}(z), \theta) : (z, \theta) \in \mathcal{S}_j,\ \|\Phi^\theta_{h_j}(z) - y_{j+1}\| \le \epsilon_{j+1}\}, \tag{10.2}$$

with Φ^θ_h the flow.

Theorem 10.1 (Soundness and persistence)**.** *If the true trajectory obeys the model with a fixed $\theta_\star \in \Theta_0$, begins in Z_0, and every observation error is within budget, then $(z(t_j), \theta_\star) \in \mathcal{S}_j$ for every checkpoint. The projections* $\mathrm{proj}_\theta \mathcal{S}_j$ *contain $\theta_\star$ and are nested.*

Proof. The claim holds initially. If the true pair lies in $\mathcal{S}_j$, the true flow carries it to the next state, and the observation inequality retains it under (10.2). Because the same parameter appears in every update, projection onto parameters can only remove candidates. □

Replacing $\mathcal{S}_j$ by $\mathrm{proj}_z \mathcal{S}_j \times \mathrm{proj}_\theta \mathcal{S}_j$ remains a valid outer enclosure but combines a state reached under one parameter with a parameter that never generated that state, and thereby

forfeits the temporal consistency the construction exists to exploit.

10.2 Velocity Observations and the Failure of State-Only Differences

If a measured velocity v_j satisfies $\|v_j - f_\theta(x_j)\| \le \eta_j$ at an exact state x_j, the parameter set is cut by the constraint $\|v_j + |x_j|^2 x_j - [x_j, -Lx_j, \mathrm{i}\,\mathrm{diag}(\nu)x_j]\theta\| \le \eta_j$, a convex set in θ because the model is linear in (a, k, s). Replacing the Euclidean norm by componentwise bounds yields a polytope amenable to linear programming. Persistent soundness and nesting hold exactly as in Theorem 10.1, and a state error of radius ϵ_j adds an allowance $\ell\epsilon_j$ when f_θ is ℓ-Lipschitz on the certified enclosure.

State-only checkpoints lead to a finite-difference velocity $\hat{v} = (\hat{z}(t+h) - \hat{z}(t))/h$. With $\|z\| \le B$, $\|f_\theta\| \le M$ and Lipschitz constant ℓ, the integral equation gives

$$\|\hat{v} - f_\theta(\hat{z}(t))\| \le \frac{2\epsilon}{h} + \ell\epsilon + \frac{\ell M h}{2}, \tag{10.3}$$

which is valid but imposes competing demands: the first term needs large h and the last needs small h. For the eight-node cycle with $\theta_0 = (0.15, 1, 0.20)$, prior half-widths $(0.01, 0.02, 0.02)$, generating $\theta_\star = (0.157, 1.013, 0.212)$ and forty-one observation pairs over four time units, the state-only certificate does not contract any prior interval for any tested $\epsilon \in \{0, 10^{-6}, 10^{-5}, 10^{-4}\}$ and $h \in \{0.001, 0.005, 0.02\}$. Even at $\epsilon = 0$ and $h = 0.001$ the certified coordinatewise residual radius is 0.02695, which leaves the whole prior box feasible. Independently measured velocities with a certified error budget do contract it. Because the frequency work operator scales as

Table 10.1: Parameter intervals after velocity observations. All intervals contain $\theta_\star$; the displaced velocity observations are shifted by 0.7η in changing complex directions.

η	Growth a	Coupling k	Frequency scale s
10^{-5}	$[0.156975, 0.157025]$	$[1.012904, 1.013113]$	$[0.211972, 0.212028]$
10^{-4}	$[0.156750, 0.157248]$	$[1.012040, 1.014128]$	$[0.211719, 0.212278]$
10^{-3}	$[0.154503, 0.159481]$	$[1.003402, 1.020000]$	$[0.209189, 0.214783]$

$K(s) = sK(1)$ for the template $\Omega = s\,\mathrm{diag}(\nu)$, an admissible interval $[s_-, s_+]$ can replace the prior frequency interval in any robust work envelope of Chapter 7. No improved resource figure is claimed here, since the experiments test identifiability and enclosure tightness and not a recomputed energy certificate.

10.3 Joint State–Parameter Screening

A finite-candidate screening on a 7^3 Cartesian grid of θ values, augmented by $\theta_\star$ as an audit control, evolves each candidate under its own unchanging parameter and discards it when its predicted checkpoint state differs from the observation by more than the certified radius. With the initial state known exactly, observations every 0.2 time units through $t = 4$, perturbation norm 0.0065 and radius 0.01: At $t = 4$ the surviving frequency-scale range is about $[0.2067, 0.2133]$,

Table 10.2: Retained candidates of 344 after joint state–parameter screening. The generating parameter is retained at every time.

Time	Retained candidates	Retained growth interval
0.0	344	$[0.140000, 0.160000]$
0.8	162	$[0.146667, 0.160000]$
1.6	59	$[0.153333, 0.160000]$
2.8	15	$[0.156667, 0.157000]$
4.0	12	$[0.156667, 0.157000]$

while the coupling range stays $[0.98, 1.02]$, which indicates weak excitation of coupling-sensitive directions. These extrema describe sampled candidates and are not certified bounds on the continuous true parameter set, and retention of the audit control demonstrates consistency and not that the method cannot falsely exclude an unsampled admissible parameter. The comparison with the state-only failure above shows that retaining the state–parameter correlation recovers discrimination that the finite-difference route cannot, but it does so only at sampled parameters.

10.4 Parameter-Dependent Taylor Enclosures

For one scalar parameter $a \in A = [a_-, a_+]$ fixed along the trajectory, let $Z(t, a)$ be the flow, $S = \partial_a Z$ and $T = \partial_a^2 Z$. Differentiating the real $2N$-dimensional equation gives

$$\dot{S} = D_z f(Z, a)S + Z, \quad \dot{T} = D_z f(Z, a)T + D_z^2 f(Z, a)[S, S] + 2S, \quad S(0) = T(0) = 0,$$

where $D_z f(Z, a)v = (a + \mathrm{i}s\nu)v - 2|Z|^2 v - Z^2\bar{v} - kLv$ componentwise. The second variation does not follow from treating $|z|^2 z$ as holomorphic. Set $B = \max\{\|z_0\|, \sqrt{N(a_+)_+}\}$, $\Lambda = a_{\mathrm{abs}} + |s|\|\nu\|_\infty + k\lambda_N + 3B^2$, and let U, V solve $U' = \Lambda U + B$ and $V' = \Lambda V + 6BU^2 + 2U$ with zero initial values.

Theorem 10.2 (Parameter-dependent Taylor enclosure)**.** *For $a, a_0 \in A$, $\|Z(t, a) - Z(t, a_0) - (a - a_0)S(t, a_0)\| \le \frac{1}{2}V(t)|a - a_0|^2$. A parameter cell $C = [m - \varrho, m + \varrho] \subset A$ has the outer enclosure $\mathcal{E}_C(t) = \overline{B}(Z(t, a_0) + (m - a_0)S(t, a_0),\ \|S(t, a_0)\|\varrho + \frac{1}{2}V(t)(|m - a_0| + \varrho)^2)$, and the cell may be discarded when $\mathcal{E}_C(t_j)$ is disjoint from $\overline{B}(y_j, \epsilon_j)$.*

Proof. On the ball $\|D_z f\| \le \Lambda$ and $\|D_z^2 f[v, v]\| \le 6B\|v\|^2$. Gronwall comparison bounds $\|S\| \le U$ and $\|T\| \le V$ uniformly in a, and Taylor's integral remainder in the real parameter gives the first claim. Splitting $a - a_0 = (m - a_0) + (a - m)$ gives the cell enclosure, and disjoint closed balls cannot contain a common exact state. □

Expanding at the cell midpoint instead of a common anchor makes the remainder scale with ϱ^2 and removes the anchor distance.

Theorem 10.3 (Cell-centered enclosure)**.** *For $C = [m - \varrho, m + \varrho]$ and every $a \in C$, $\|Z(t, a) - Z(t, m) - (a - m)S(t, m)\| \le \frac{1}{2}V(t)\varrho^2$, so the image of C lies in $\overline{B}(Z(t, m),\ \|S(t, m)\|\varrho + \frac{1}{2}V(t)\varrho^2)$. A cell is excluded when $\|Z(t_j, m) - y_j\| > \|S(t_j, m)\|\varrho + V(t_j)\varrho^2/2 + \epsilon_j$.*

On the three-node path with $k = 0.7$, $s = 0.16$, $\nu = (-1, 0, 1)$, $z_0 = (0.45 + 0.10\mathrm{i}, 0.21 - 0.06\mathrm{i}, 0.33 + 0.04\mathrm{i})$, $A = [0.12, 0.22]$, $a_\star = 0.173$ and observation radius 0.003, the common-anchor enclosure of Theorem 10.2 retains $68, 36, 28, 122$ of 128 cells at times $0.2, 0.4, 0.6, 0.8$, with the last number showing loss of discrimination because the global majorant V reaches 43.9. Cell-centering retains $67, 35, 24, 19$, and intersecting cells across all checkpoints (persistent retention) keeps the same counts. The generating cell survives in every test. Cell-centering therefore removes the failure at the longest time, although V itself remains loose. The floating-point integration of the nominal flows is not outward rounded, so these counts are numerical demonstrations and not formally certified exclusions.

10.5 A Negative Result: Global Lipschitz Enclosures Retain Everything

A Euler-based continuous-cell scheme gives an enclosure in exact real arithmetic under explicit bounds. Let the cell have center $\bar{a}$ and half-width ϱ, let $x_{n+1} = x_n + h f_{\bar{a}}(x_n)$ and

$$r_{n+1} = (1 + h\Lambda) r_n + h\varrho B + \tfrac{1}{2} h^2 M,$$

with $\|\ddot{z}\| \le M = \Lambda F_\star$ and $F_\star$ the speed bound. Then $\|z(t_n) - x_n\| \le r_n$ for every a in the cell, and a cell is discarded only when $\|x_n - y_j\| > r_n + \epsilon_j$. On the same three-node system with the continuous prior $[0.12, 0.22]$, known initial state, observations at $t = 0, 0.2, \ldots, 0.8$ of norm 0.35ϵ with $\epsilon = 0.003$ and step 0.002, the constants are $B \approx 0.8124$, $\Lambda = 4.46$, $F_\star \approx 2.551$ and $M \approx 11.38$. The entire prior survives every partition, because the compounding factor

Table 10.3: Continuous-cell Euler enclosure: subdivision does not discriminate.

Parameter cells	Retained cells	Final enclosure radius	True cell retained
8	8	0.125031	yes
16	16	0.105741	yes
32	32	0.096095	yes
64	64	0.091273	yes
128	128	0.088861	yes

$1 + h\Lambda$ and the global second-derivative remainder dominate the observational radius. The result is evidence that this enclosure is too coarse, and not evidence of intrinsic non-identifiability, since the sampled screening above discriminates sharply. It separates mathematical inclusion from useful identification and from machine certification, and it motivates correlation-preserving propagation, the subject of Chapters 11 and 12.

Part IV

Validated Enclosures

Chapter 11

Interval Propagation and Parameter-Affine Taylor Models

The analytic certificates of Part II are valid for all time but are only as tight as their global constants permit. A complementary route computes an enclosure of the reachable set of the augmented state (z, J) directly, by validated numerical integration, and reads cumulative expenditure from the J component. The guarantee is then conditional on the arithmetic and on the correctness of the enclosure scheme, and it holds for finite horizons only, but the bounds can be tight over those horizons. This chapter develops the interval and Taylor schemes in order of increasing dependency preservation, on the three-node path with $k = 0.7$, $c_A = 0.1$, $\Omega = 0.16\,\mathrm{diag}(-1, 0, 1)$, $a \in [0.12, 0.22]$ fixed along each trajectory, and initial state $z_0 = (0.45+0.10\mathrm{i}, 0.21-0.06\mathrm{i}, 0.33+0.04\mathrm{i})$, with $\dot{J} = c_A k^2 \|Lz\|^2$. All statements about computed enclosures below refer to outward-rounded multiprecision interval arithmetic (35 to 40 decimal digits), whose library is relied upon but has not been independently audited, and none constitutes a machine-checked proof.

11.1 Global Majorants and a First-Order Validated Step

The cubic radial damping keeps $\|z(t)\| \le B = 1$ for the three-node system, because at $\|z\|^2 = 1$ the upper derivative $2a_{\max}\|z\|^2 - 2\|z\|^4/N$ is negative. On this ball the speed and curvature are bounded by $F = (a_{\max} + \|\Omega\| + B^2 + k\lambda_{\max})B = 3.48$ and $M = \Lambda F = 19.0704$ with $\Lambda = a_{\max} + \|\Omega\| + 3B^2 + k\lambda_{\max} = 5.48$, and $|\ddot{J}| \le M_J = 2c_A k^2 \lambda_{\max}^2 BF$.

Theorem 11.1 (First-order interval step with remainder)**.** *Let X_n and I_n be interval enclosures of $z(t_n; C)$ and $J(t_n; C)$ for the parameter cell C, with one parameter interval used in every evaluation. Then*

$$X_{n+1} = X_n + hf(X_n, C) + [-1, 1]\tfrac{h^2 M}{2}, \qquad I_{n+1} = I_n + hp(X_n) + [-1, 1]\tfrac{h^2 M_J}{2}$$

enclose the true state and cumulative expenditure at t_{n+1} for every $a \in C$, provided all interval operations are outward rounded and the majorants hold along the step. Here $p(z) = c_A k^2 \|Lz\|^2$.

Proof. For fixed $a \in C$, Taylor's theorem gives $z(t+h) = z(t) + hf + r$ with each coordinate of r bounded by $h^2 M/2$, and likewise for J. Induction closes the argument. □

Since J is nondecreasing, an upper bound on $J(H)$ controls actuation at every earlier time, but other signed resource terms require their own prefix treatment. With $h = 0.002$ and 8, 16, 32 cells the largest coordinate width is 0.0049 at $H = 0.1$ for 8 cells and 0.0046 for 32, so subdivision helps little, because wrapping and the step remainder dominate.

11.2 Picard Tubes and Higher-Order Time Coefficients

A tighter step builds a tube Y with $X_n + [0,h]F(Y,C) \subset \operatorname{int} Y$, which by the Picard argument encloses every solution starting in X_n for the fixed parameter, and then evaluates Taylor remainders over Y instead of using global majorants.

Proposition 11.2 (Picard-tube Taylor update)**.** *If X_n contains every admissible augmented state, the fixed parameter lies in C, interval operations enclose their exact counterparts, and the Picard inclusion holds, then both $X_{n+1}^{(1)} = X_n + hP_1(X_n, C) + h^2 P_2(Y, C)$ and $X_{n+1}^{(2)} = X_n + hP_1(X_n, C) + h^2 P_2(X_n, C) + h^3 P_3(Y, C)$ enclose every admissible state at t_{n+1}, where P_j are the time-Taylor coefficients computed by truncated polynomial arithmetic.*

With $h = 0.005$ and all twelve tested configurations passing the Picard test, first and second order give at $H = 0.1$ with 8 cells widths 0.000623 and 0.000612, and at $H = 0.2$ widths 0.001362 and 0.001340 with J upper endpoints near 0.002026. The second-order gain is slight, consistent with interval wrapping dominating the time truncation on this short benchmark.

11.3 Cell-Centered Sensitivities

To reduce parameter overestimation, expand about the midpoint m of each cell with analytic majorants for the second sensitivities, as in Theorem 10.3, augmented by expenditure sensitivities $K = \partial_a J$ and $W = \partial_a^2 J$ obeying $\dot{K} = 2c_A k^2 \operatorname{Re}\langle Lz, LS\rangle$ and $\dot{W} = 2c_A k^2(\|LS\|^2 + \operatorname{Re}\langle Lz, LT\rangle)$. With $U(t) = B(e^{\Lambda t} - 1)/\Lambda$, $V(t) = \frac{e^{\Lambda t}-1}{\Lambda}(6BU^2 + 2U)$ and $W_\star(t) = 2c_A k^2 \|L\|^2 t(U^2 + BV)$, one has $\|S\| \le U$, $\|T\| \le V$ and $|W| \le W_\star$.

Theorem 11.3 (Centered enclosure for state and expenditure)**.** *If outward-rounded intervals Z_m, S_m, I_m, K_m enclose the state, its first sensitivity, the expenditure and its first sensitivity at the midpoint, then for every $a \in [m - \varrho, m + \varrho]$,*

$$z(t,a) \in Z_m + [-\varrho,\varrho]S_m + [-R_z, R_z]^6, \qquad J(t,a) \in I_m + [-\varrho,\varrho]K_m + [-R_J, R_J],$$

with $R_z = V\varrho^2/2$ and $R_J = W_\star \varrho^2/2$.

Three high-accuracy numerical trajectories (the two parameter endpoints and the midpoint) fall inside the computed enclosures for every cell, which supplements and does not replace the analytic inclusion proof. The gain is under 2%, because the nominal interval trajectories of Z_m and S_m still wrap and that wrapping is shared by both schemes.

Table 11.1: Direct against centered enclosures, three-node path. Widths are maxima over the six coordinates and all cells.

Cells	H	Direct width	Centered width	Direct J^+	Centered J^+
8	0.10	0.00492123	0.00486705	0.00156433	0.00156393
8	0.20	0.01143580	0.01128266	0.00277231	0.00277134
32	0.10	0.00455461	0.00454730	0.00156321	0.00156309
32	0.20	0.01078766	0.01074527	0.00276744	0.00276676

11.4 Parameter-Affine Taylor Models

The sensitivity expansion can be pushed one order further and the second coefficient propagated over the whole cell, writing $X(t,a) = X_0(t) + uX_1(t) + u^2R(t,u)$ with $u = a - m$.

Proposition 11.4 (Cell-centered affine enclosure)**.** *If $\mathcal{X}_0, \mathcal{X}_1$ enclose $X(t,m)$ and $\partial_a X(t,m)$ and $\mathcal{X}_2$ encloses $\frac{1}{2}\partial_a^2 X(t,a)$ for every $a \in C$, then $X(t,a) \in \mathcal{X}_0 + [-\varrho,\varrho]\mathcal{X}_1 + [0,\varrho^2]\mathcal{X}_2$ for every $a \in C$, and its seventh coordinate encloses cumulative expenditure.*

Proof. Taylor's theorem in the scalar parameter with Lagrange remainder applies coordinatewise to the smooth solution map, and the intermediate parameter lies in C. □

Time propagation uses a Picard tube and the second-order step $X_{n+1} \subset X_n + hF(X_n, C) + \frac{h^2}{2}DF(\mathcal{B})F(\mathcal{B})$ with $\mathcal{B}$ the tube, computed by polynomial automatic differentiation in the parameter. Strict inclusion is tested with interval-native comparisons, and steps are built from decimal strings, so that no binary floating-point conversion enters the verification. The conversion to floating point remains only for reporting. The reduction grows with horizon, to about 22%

Table 11.2: Strict parameter-affine enclosure against direct interval integration.

Cells	H	Direct width	Affine width	Reduction	Affine J^+
4	0.05	0.00059168	0.00055627	5.98%	0.00067422
4	0.10	0.00123378	0.00108684	11.91%	0.00122036
4	0.20	0.00269847	0.00208698	22.66%	0.00202078
8	0.05	0.00029854	0.00028236	5.42%	0.00067421
8	0.10	0.00062259	0.00055227	11.30%	0.00122034
8	0.20	0.00136216	0.00106309	21.96%	0.00202076

at $H = 0.2$, while the expenditure upper bounds are almost unchanged, a first sign of a theme developed in Chapter 12: width reduction in state coordinates need not tighten a bound on an accumulated observable.

11.5 What Remains Unsolved at This Level

All the schemes of this chapter retain interval wrapping in the sensitivity dynamics, so none controls long-horizon growth. None uses observations to reject parameter cells with machine-certified enclosures. None establishes the energy floor at every intermediate time beyond the

monotonicity argument for J. The next chapter replaces coordinatewise hulls of the nominal and sensitivity trajectories by representations that transport correlations explicitly.

Chapter 12

Wrapping, Correlated Generators and Moving Frames

Coordinatewise interval hulls discard the correlations between state variables created by the flow, so the enclosure of a stable system can grow even when the true reachable set does not. This chapter isolates that effect on an exactly solvable benchmark, then develops correlated-generator enclosures for the nonlinear reduced oscillator pair of Chapter 9, validates them in continuous time by Picard tubes and outward dyadic rounding, and compares fixed and moving frames. The comparisons are deliberately controlled, with identical initial sets, step sizes and rounding, and they include several outcomes that run against expectation.

12.1 An Exact Benchmark for Wrapping

Let $p \in [9/10, 11/10]$, $v_0 = (p, 0)^\top$ and $v_{n+1} = Rv_n$ with the rotation-scaling $R = \frac{1}{5}\left(\begin{smallmatrix} 3 & -4 \\ 4 & 3 \end{smallmatrix}\right)$, which satisfies $R^\top R = I$. The affine representation $v_n = R^n(1, 0)^\top p$ retains the same p at every step, whereas a coordinatewise box computes $X_{n+1} = \mathrm{hull}(RX_n)$.

Proposition 12.1 (Exponential wrapping)**.** *The sum of widths of the box after $n \geq 1$ steps is $W_n^{\mathrm{box}} = \frac{7}{25}(\frac{7}{5})^{n-1} = \frac{1}{5}(\frac{7}{5})^n$, whereas the affine representation has width $W_n^{\mathrm{aff}} = \frac{1}{5}(|(R^n)_{11}| + |(R^n)_{21}|) \leq \frac{\sqrt{2}}{5}$. For the expenditure surrogate $J_n = \sum_{j<n} \|v_j\|^2$ orthogonality gives $J_n = np^2 \in [\frac{81n}{100}, \frac{121n}{100}]$ exactly.*

Proof. After the first rotation the coordinate radii are $(3/50, 4/50)$, and under independent interval evaluation their sum is multiplied by $7/5$ at every later step. The affine widths are $\frac{1}{5}$ times the absolute entries of the first column of R^n, and Cauchy–Schwarz with orthogonality gives the bound. □

The benchmark is a discrete linear surrogate, so it proves that retaining a common uncertain variable can prevent catastrophic wrapping in a stable map but does not quantify the improvement for the nonlinear oscillator, where parameter–state interactions create genuine remainders. A related tangent model makes the same point for the oscillator itself: freezing the Jacobian at a reference point of the three-oscillator system and transporting a shared affine frame

Table 12.1: Exact rational benchmark: box against affine width.

Steps	Box width	Affine width	Ratio
4	0.768320	0.276160	2.78
8	2.951578	0.265716	11.11
12	11.338782	0.224578	50.49
16	43.559067	0.281801	154.57
20	167.336511	0.250652	667.61
24	642.839940	0.245251	2621.15

$A_n = M^n A_0$ gives the exact image of the initial set, with coordinate half-widths $\sum_j |(A_n)_{ij}| \le b_n$ componentwise, where b_n is the repeated-hull recursion $b_{n+1} = |M| b_n$. The width ratio of repeated boxes to the affine frame grows from 2.41 at $t = 0.5$ to 5.57 at $t = 4$. This too is a linear surrogate, not a certification of the nonlinear flow.

12.2 Nonlinear Discrete Maps with Correlated Generators

For the reduced real pair in physical coordinates, a nonlinear Euler map with persistent growth $a = 0.17 + 0.02\xi$, $\xi \in [-1, 1]$, can be enclosed by $y = c + A\xi + e$ with the center following the midpoint map, A following the exact first derivative of the map at the center including its explicit parameter derivative, and the remainder bounded in exact rationals by $d\,U_i + 3|c_i|U_i^2 + U_i^3$ for the state and $2c_A k^2 (U_1 + U_2)^2$ for expenditure, where $U_i = |A_i| + r_i$ and $d = 0.02$. With $h = 0.025$, eight steps give at $t = 0.2$ an affine state-width sum 0.00518828 against 0.00630983 for direct interval evaluation of the same map, a 17.8% reduction, while the expenditure upper bounds are 0.00087020 and 0.00087735. In collective–disagreement coordinates, three independent generators encoding the initial p, initial m and the persistent parameter, with the Jacobian at the center transporting G and a mean-value bound controlling the Jacobian variation, give at $t = 0.4$ a 37.7% narrower state enclosure than interval evaluation (0.009272 against 0.014893). These are inclusion theorems for the discrete polynomial map in exact rational arithmetic, and they certify neither the continuous flow nor expenditure beyond the discrete accumulator.

12.3 Continuous-Time Validation by Picard Tubes

To bound the continuous flow, each step first constructs a tube $\mathcal{T}$ with $X + hf(\mathcal{T}, A) \subset \operatorname{int} \mathcal{T}$, so every solution from X stays in $\mathcal{T}$ for the fixed parameter. Because the system is autonomous and the parameter constant, $x'' = D_x f(x, a) f(x, a)$, and the flow endpoint lies in $X + hf(X, [a]) + [-\mathcal{D}, \mathcal{D}]$ with $\mathcal{D}_i = \frac{h^2}{2} \sum_j \sup_{\mathcal{T}} |(D_x f)_{ij} f_j|$. The same defect is added to the correlated affine propagator. All endpoints are exact rational fractions rounded outward to a dyadic grid, here 40 fractional bits, with inclusion comparisons carried out in exact arithmetic.

Proposition 12.2 (Continuous-time inclusion)**.** *If X_n encloses the exact states for every admissible parameter and the Picard inclusion holds, then the exact trajectory over $[t_n, t_n + h]$ lies in $\mathcal{T}_n$ and its endpoint lies in X_{n+1}; if the lower energy endpoint of every tube is at least E_*, the floor holds at every time through the final step.*

For the reduced pair with $a \in [0.15, 0.19]$ fixed on each trajectory, $p(0) \in [0.329, 0.331]$, $m(0) \in [0.119, 0.121]$ and $h = 0.01$, the correlated-generator enclosure represents $x = c + \sum_\ell g_\ell \xi_\ell + [-r, r]$ with three generators, the third encoding $a = 0.17 + 0.02\xi_3$. A longer run, H up to 2, with

Table 12.2: Continuous-time validated enclosures, interval against correlated generators.

H	Interval width	Affine width	Affine J upper
0.1	0.006825	0.005899	0.00051106
0.2	0.009859	0.007678	0.00088686
0.4	0.016705	0.011075	0.00136804
0.6	0.024889	0.014504	0.00163178
0.8	0.034876	0.018200	0.00177920

the same fixed-generator representation shows the effect compounding: the reduction in state width relative to direct intervals grows from 22.1% at $H = 0.2$ to 47.8% at 0.8 and 67.2% at 2.0 (widths 0.193859 against 0.063604), while the J upper bound increases slowly from 0.00088686 to 0.00208117. The affine J bound at $H = 2$ is a validated statement for $t \in [0, 2]$, in contrast to the analytic ceiling 0.00187 of Chapter 9, which holds for all time.

12.4 Resource-Floor Certification from Monotone Expenditure

The passage from an enclosure of J to a statement about E rests on one lemma, used repeatedly in the tables above and below.

Lemma 12.3 (Prefix certification)**.** *Let $E(t) = E(0) + (q - b)t - J(t)$ with J nondecreasing, $J(0) = 0$, and $J(H) \le J^+$. If $q \ge b$ then $E(t) \ge E(0) - J^+$ for every $t \in [0, H]$. If $q < b$ then $E(t) \ge E(0) - (b - q)H - J^+$ for every $t \in [0, H]$. In particular the endpoint lower bound $E(H) \ge E(0) + (q - b)H - J^+$ is a different statement and must not be reported as uniform in time.*

Proof. Monotonicity gives $J(t) \le J(H) \le J^+$, and $(q-b)t \ge \min\{0, (q-b)H\}$ for $t \in [0, H]$. □

With $E(0) = 0.1$, $q = 0.01$, $b = 0.005$ and floor 0.08, the correlated enclosure gives $E(t) \ge 0.0979188294$ for all $t \in [0, 2]$ and $E(2) \ge 0.1079188294$, both above the floor. The conservative variant that discards positive supply uses $E(0) - bH - J^+$.

12.5 Dyadic Compression and Fixed Parameter Intervals

Exact rational endpoints grow in denominator under repeated polynomial operations, so outward rounding to a bounded dyadic grid keeps the computation finite. For the reduced real pair with fixed uncertain $a \in [0.15, 0.19]$, a rational interval step with outward dyadic rounding at $p = 24$, 40 and 64 bits gives at $H = 0.4$ state-width sums 0.02356875, 0.02356170, 0.02356170, upper expenditure estimates 0.00141549, 0.00141426, 0.00141426, and energy lower bounds 0.10058439, 0.10058574, 0.10058574, so denominator control does not dominate enclosure error beyond 40 bits. This is a rigorous continuous-time inclusion for the reduced real subsystem, conditional on

the Picard test, and it does not avoid interval wrapping, because the state widths continue to grow.

Combining dyadic compression with the parameter-affine decomposition $y = c + s\xi + e$ ($a = 0.17 + 0.02\xi$, $h = 0.025$) adds the rounding errors of center and sensitivity to the remainder before its upper endpoint is rounded outward. At $H = 0.4$ the conservative all-prefix bound

Table 12.3: Dyadic compression with a parameter-affine decomposition (40 bits).

H	Direct width	Affine width	Reduction	Affine J upper
0.1	0.004641	0.004389	5.4%	0.00052934
0.2	0.010038	0.008804	12.3%	0.00091958
0.3	0.016387	0.013316	18.7%	0.00120858
0.4	0.023918	0.017988	24.8%	0.00142392

$E(0) - bH - J^+$ is 0.09657608, above 0.08, and the endpoint-only bound crediting supply is 0.10057608.

12.6 Parameter-Cell Reconditioning

Subdividing the growth parameter into one, two or four cells gives each cell its own nominal trajectory, first sensitivity and remainder. This is parameter-cell reconditioning and not a moving-frame state algorithm. Four cells reduce the maximum per-cell width by 44.4% while

Table 12.4: Cell subdivision at $H = 0.4$. The reported width is the maximum per-cell width and not the width of the union.

Cells	Max cell width	Max J upper	Min energy lower
1	0.017988	0.00142392	0.096576
2	0.012651	0.00142003	0.096580
4	0.010006	0.00141817	0.096582

the global J upper bound, the maximum over cells, decreases by only 0.4%. A valid global expenditure bound is the maximum of the cell bounds, so width reductions per cell must not be reported as reductions of the union.

12.7 Equivalence of Coordinate Representations

An apparent improvement from collective–disagreement coordinates (p, m) appears when the enclosure is compared with a physical-coordinate representation that preserves exactly the same initial parallelogram. The direct baseline starts from the coordinate hull of that set and thereby loses the initial correlation.

Proposition 12.4 (Representation equivalence)**.** *Let Y be a validated enclosure in (p, m, J) coordinates and $\mathcal{C}(Y) = \{(p + m, p - m, J)\}$. If a physical affine propagator evaluates the conjugated vector field with the same interval operations and rounding decisions, its output at every checkpoint equals $\mathcal{C}(Y)$.*

Proof. $\mathcal{C}$ is an invertible linear map conjugating the physical field to the modal one, so transporting the shared generators through either field and applying $\mathcal{C}$ at output gives the same set. The statement concerns the specified computational constructions and not every interval algorithm in the two coordinate systems. □

Table 12.5: Direct hull against the correlated representation (identical initial uncertainty set, $h = 0.01$).

H	Direct width	Correlated width	Direct J upper	Correlated J upper
0.1	0.015778	0.015767	0.000551	0.000537
0.2	0.024832	0.024019	0.000974	0.000940
0.4	0.048213	0.042502	0.001573	0.001477
0.6	0.081786	0.064646	0.001995	0.001804
0.8	0.130889	0.091977	0.002362	0.002023

At $H = 0.8$ the correlated width is 29.73% below the direct hull. The improvement therefore measures the cost of losing correlation relative to independent physical-coordinate boxes, and it is not a property intrinsic to the collective–disagreement coordinates.

12.8 Negative and Mixed Results on Moving Frames

Three attempts to improve on the fixed correlated representation by changing the frame give results that bear directly on which claims are warranted.

A Jacobian-guided shear. The moving frame uses $v = p - sm$ with s updated after each step by a midpoint-Jacobian heuristic, and between steps $v_{\text{new}} = v_{\text{old}} + (s_{\text{old}} - s_{\text{new}})m$. The physical-coordinate enclosure becomes *wider*: 0.120353 against 0.091977 at $H = 0.8$, a 30.8% increase, with the J bounds nearly equal (0.0020247 against 0.0020227). The interval hull taken at each basis change can overwhelm any reduction in local linear coupling, so a dynamically changing basis is not intrinsically wrapping resistant, and no general theorem that fixed frames dominate moving frames follows.

A rescaling aligned with linear decay. Writing $m = s_n v$ with $s_{n+1} = (1 - 2kh)s_n$ is an invertible rescaling that follows the linear disagreement decay. The improvement is negligible (0.091905 against 0.091977 at $H = 0.8$, and 0.0020222 against 0.0020227 for J), because the coordinatewise enclosure of the nonlinear remainder terms is unchanged.

Matrix transport of the remainder. A genuine matrix-valued frame $B_{n+1} = (I + hDf(c_n, a_0))B_n$ with the interval remainder carried in the transported coordinates through the exact rational inverse is a first-order validated moving-frame enclosure. At $H = 2$ matrix transport reduces the state width by 55.1% (0.0636 to 0.0285) but *raises* the expenditure upper bound from 0.00208117 to 0.00231344. The outcome is mixed: state-space wrapping improves, while the accumulated nonlinear observable J is sensitive to the residual forcing mapped through $|B^{-1}|$.

Table 12.6: Matrix transport against the fixed correlated representation.

H	Direct width	Fixed width	Moving width	Fixed J upper	Moving J upper
0.2	0.009859	0.007678	0.007501	0.000887	0.000893
0.4	0.016705	0.011075	0.010357	0.001368	0.001386
0.8	0.034876	0.018200	0.015194	0.001779	0.001833
1.2	0.062950	0.027313	0.019648	0.001918	0.002019
1.6	0.109760	0.040877	0.024067	0.001990	0.002154
2.0	0.193859	0.063604	0.028549	0.002081	0.002313

Observable-aware combination. One might hope to take the best of both. A third certificate intersects the disagreement projections of the fixed and moving Picard tubes on each step and accumulates $U_{n+1} = \lceil U_n + hC \sup\{m^2 : m \in M_n\} \rceil$, so each of the three enclosures independently bounds J and the minimum is valid. The independently accumulated tube bound is weaker

Table 12.7: Expenditure bounds from three independent enclosures and the validated minimum.

H	Fixed J upper	Moving J upper	Quadrature J upper
0.2	0.0008869	0.0008927	0.0010418
0.4	0.0013680	0.0013864	0.0016390
0.8	0.0017792	0.0018328	0.0021963
1.2	0.0019185	0.0020192	0.0024060
1.6	0.0019904	0.0021536	0.0024976
2.0	0.0020812	0.0023134	0.0025458

than both, so taking the minimum recovers 0.00208117 with no new improvement. Narrower state intervals do not automatically tighten a bound on a path-dependent quadratic observable, because interval quadrature discards temporal dependence and intersecting projected tubes does not recover the correlation between $m(t)$ and the accumulated J. A stronger scheme would transport a joint correlated representation of (p, m, J), possibly with generators for integrated quadratic terms. All of these results concern the reduced real system, and the matrix-transport frame is a first-order construction, not a QR-based Lohner integrator.

Part V

Synthesis

Chapter 13

What the Programme Establishes and What Remains Open

The preceding chapters pursued one question, namely when dynamical coordination preserves the capacity for continued operation, through three families of guarantees: analytic bounds valid for all time under stated constants, validated enclosures valid on finite horizons under stated arithmetic, and numerical experiments that calibrate the former against trajectories. This chapter states how those families relate, separates what is established from what is merely suggested, and lists the problems that the programme leaves open.

13.1 A Taxonomy of Guarantees

The scope of a guarantee matters as much as its numerical value, since a tighter bound with narrower scope is not a better result of the same kind. Table 13.1 records for each family the object bounded, the horizon, the assumptions it leans on, and the kind of evidence offered. Only the first four rows are theorems whose conclusion holds for the continuous system without reference to a computation. The enclosure rows are theorems about inclusion, but their application to a specific number depends on an implementation of interval arithmetic that is trusted and not machine-verified. The calibration row is an experiment, and its content is a set of measured ratios on a stated sample of graphs, frequency scales, initial states and horizons.

13.2 Established Results

The resource floors below hold for a constant positive net supply, a fixed coupling gain and the stated actuator law. They do not certify viability for a finite reservoir or for a coupling gain that depends on the resource. Within the stated model, the programme establishes the following. A resource floor E_* is preserved on a horizon whenever cumulative actuator expenditure stays below the initial surplus plus net supply, and the edge-dissipative structure of the cubic nonlinearity yields an expenditure bound that does not grow with network size through the nonlinearity itself. The direct spectral certificate gives a transient term proportional to initial edge energy divided by $\beta = k - (a_+ + \delta)/\lambda_2$ and a persistent term that grows linearly in time with a rate set by the

Table 13.1: Guarantee types used in this monograph.

Family	Object bounded	Horizon	Principal assumption
Edge-dissipative and direct spectral certificates (Chapters 4–6)	Cumulative expenditure J	All time	Constants λ_2, k, global amplitude bound, forcing size
Signed work and checkpoint envelopes (Chapter 7)	J and its prospective growth	Finite, restartable	Lipschitz constants on the enclosure
Young allocation (Chapter 8)	Same as direct certificate	All time	Choice of θ, which only redistributes slack
Reduced-system dissipation (Chapter 9)	J for two oscillators	All time	Invariant strip, parameter interval
Interval and affine enclosures (Chapters 11–12)	State and J	Finite, fixed horizon	Outward rounding, Picard inclusion
Set-valued inference (Chapter 10)	Feasible parameters	Observation window	Persistent parameter, noise budget
Calibration (Chapter 8)	Ratio of bound to trajectory	Sampled	Numerical integration accuracy

frequency work F. The Young parameter allocates the same dissipation budget between the two terms, the closed forms of Chapter 8 are exact, and the achievable gain on the transient is a factor of at most two. The reduced pair admits a joint dissipation functional that reduces certification to a scalar Riccati comparison. Persistence of parameters through set-valued recursion is sound, and correlated generators prevent the exponential wrapping that coordinatewise hulls suffer on a stable linear map.

13.3 Mixed and Negative Results

Several outcomes were not what the preceding analysis suggested, and they bound the claims that may responsibly be made.

Calibration. The optimized direct certificate lies within roughly 1.1 to 1.7 times the true expenditure on cycle and complete graphs, and a median of about 2.3 to 3.2 times on the path graph. The remaining slack lies in the persistent term, which uses the global amplitude bound and the full frequency-gradient norm. The Young parameter cannot recover that slack because it acts on the transient only.

State-only inference. Finite-difference velocities derived from state checkpoints fail to contract any prior interval at the tested noise levels, while direct velocity observations and sampled joint screening succeed.

Global Lipschitz cells. Euler enclosures with global constants retain every parameter cell at every tested subdivision, so a valid inclusion is not necessarily informative.

Moving frames. Matrix transport narrows the state enclosure by 55.1% at $H = 2$ yet raises the expenditure bound from 0.00208117 to 0.00231344, and an observable-aware minimum of three valid bounds recovers the smaller value without improving it.

Coordinate choice. The apparent advantage of collective–disagreement coordinates over physical coordinates disappears when the physical-coordinate baseline keeps the same initial correlation (Proposition 12.4). These results share a pattern: narrowing an enclosure in state space does not by itself narrow a bound on a path-dependent accumulated quantity, and a bound that is valid for every parameter in a set may be uninformative about which parameter is true.

13.4 Open Problems

Mode-resolved forcing. Theorem 8.4 allocates dissipation mode by mode and reduces the bound by up to 62% for forcing concentrated in higher modes, but by under 3% for the reference pattern. What remains open is the shared cross-mode term $\delta/\underline{\lambda}_2$, which is not resolved by mode, and a modal replacement for the global amplitude bound R, which would have to control cross-mode coupling through the cubic term.

State-dependent amplitude bound. Proposition 8.5 replaces the constant R by a closed-form time-dependent envelope, which removes most of the amplitude conservatism for small initial states and little of it near the limit-cycle amplitude. A bound that also uses restart information, such as the observed maximum modulus at a checkpoint, and a bound on the mean m that is sharper than the maximum modulus, remain open. The latter matters most for the four-node reference examples, where the initial amplitude is close to $\sqrt{a}$.

Time-varying supply and actuator control. Throughout, q, b and k are constants. If the coupling gain depends on the resource, $k = k(E)$, then expenditure and synchronization interact, and Lemma 12.3 no longer applies as stated because J ceases to be a prescribed monotone functional of the trajectory alone. A viability analysis with feedback from E to k would address the central question directly, since it asks when coordination can be sustained while it consumes the resource it relies upon.

Mixed-state and stochastic extension. Intermittent or random supply, and noise in the frequency pattern, would require replacing deterministic floors by probabilistic ones. A natural route uses the dissipation inequality of Chapter 6 as a supermartingale drift condition, but no such calculation has been carried out.

Machine-checked enclosures. All validated enclosures reported here rely on a rational or multiprecision interval library whose correctness is assumed. Re-running the Picard-tube step in a verified environment, or in exact rational arithmetic throughout, would convert the finite-horizon tables into machine-checked statements.

13.5 An Extension to Intermittent Forcing

The same formalism extends naturally to environmental forcing that switches on and off, as in field studies of plants whose resource supply follows an irregular photoperiod or precipitation cycle. If the supply is $q(t) \in \{q_{\mathrm{on}}, q_{\mathrm{off}}\}$ with dwell times bounded below, the prefix certificate of Lemma 12.3 applies segment by segment, and the question becomes whether the expenditure accumulated during an off interval can be recovered during the next on interval before the floor is reached. Whether a periodic schedule keeps the surplus above the floor depends on the ratio of mean supply to mean expenditure rate, and on the burst expenditure that follows each switch, because transients restart the edge energy term. This direction is proposed and not developed, and it is the main reason the checkpoint construction was designed to allow restarts.

13.6 Closing Remarks

The central finding is modest and exact. A coordination mechanism can be priced in advance by a bound that depends only on graph spectrum, coupling strength and the size of the frequency pattern, and that price can be made reasonably tight on well-connected graphs. The remaining conservatism is concentrated in the persistent forcing term, so progress on the question posed at the start depends on resolving that term mode by mode, and on allowing supply and coupling to depend on the resource state.

Appendix A

Notation

Graphs are undirected and connected with N nodes, Laplacian L with eigenvalues $0 = \lambda_1 < \lambda_2 \le \cdots \le \lambda_N = \lambda_{\max}$, and projector $P = I - \frac{1}{N}\mathbf{1}\mathbf{1}^\top$ onto the disagreement subspace. The oscillator state is $z \in \mathbb{C}^N$ with transverse part $w = Pz$. The parameters a, k and $\Omega = \mathrm{diag}(\omega_i)$ denote growth rate, coupling strength and natural frequencies, with $\tilde{\omega} = P\omega$. The resource is E with floor E_*, supply q and basal drain b, and the cumulative actuator expenditure is $J(t) = \int_0^t c_A k^2 \|Lz\|^2 \, ds$. Edge energy is $Q = z^* L z$ and $W = Q/2$. The disagreement rate is $\beta = k - (a_+ + \delta)/\lambda_2$, the amplitude bound is R, and the frequency work is $F = R\|\tilde{\omega}\|$. The Young parameter is $\theta \in (0, 1)$, and the transient and persistent coefficients are $\alpha = c_A k^2 W(0)/\beta$ and $\rho = c_A k^2 F^2/\beta^2$. For the two-oscillator reduction, p and m denote collective and disagreement coordinates. Intervals are written in capitals, $\overline{B}(c, r)$ is the closed ball, and outward rounding is understood wherever an enclosure is called validated.

Appendix B

Computational Materials

Every table in this monograph was produced by a short Python program, and each program writes its results to a comma-separated file. Table B.1 lists the principal programs behind each chapter. Programs whose names begin with `make` only format results as tables and are omitted.

Conventions. Enclosure computations use exact rational arithmetic or multiprecision interval arithmetic as stated in the relevant chapter, and reference trajectories use floating-point integration and are marked numerical wherever reported. The calibration and allocation studies of Chapter 8 are described here in more detail. The program `young_parameter_optimization.py` evaluates the Young-parameterized certificate in exact rational arithmetic, rounds square roots upward at 2^{-80}, and writes the supply-matched and finite-horizon optima. The program `certificate_calibration.py` integrates the complex four-node system on path, cycle and complete graphs at two frequency scales and nine initial states with an eighth-order Dormand–Prince method at relative tolerance 10^{-11}, evaluates three certificates at $T \in \{1, 5, 20, 100\}$, and writes 216 rows in each of which every bound is at least the integrated expenditure. The program `mode_resolved_certificate.py` evaluates Theorem 8.4 against the global Young bound (144 rows), and `amplitude_envelope_certificate.py` evaluates Proposition 8.5 with and without the modal allocation (120 rows) and imports helper functions from the modal program, so the two must share a directory.

Evidence kind by table. The guarantee behind a number depends on how it was computed, and Table B.2 separates the four kinds used here. Agreement between a bound and an integrated trajectory is a consistency check and not an enclosure proof, and no independent machine verification of the interval-arithmetic implementations was performed.

Reproduction. The command `python3 reproduce.py` in the experiments directory re-runs the four programs of Chapter 8 in a scratch directory, compares every regenerated table with the reference tables to a relative tolerance of 10^{-9}, and checks the headline figures quoted in the text. Library versions are pinned in `requirements.txt`, and the random seeds are 20261118 (calibration), 7 (modal allocation) and 11 (amplitude envelope).

Table B.1: Computational materials by chapter.

Chapter	Principal scripts
2, 3	synchronization_experiments, stuart_landau_actuation_experiment, actuation_derived_viability_experiment, coupling_viability_experiment
4, 5	spectral_actuation_experiment, nonlinear_certificate_comparison, edge_dissipation_comparison, time_dependent_dissipation
6	network_dissipation_certificate, heterogeneous_edge_experiment, heterogeneous_network_certificate, direct_spectral_expenditure, graph_aware_frequency, frequency_correlation_experiment
7	local_frequency_work, predictive_frequency_work, perturbed_checkpoint_experiment, uncertain_checkpoint_experiment
8	young_parameter_optimization, certificate_calibration, mode_resolved_certificate, amplitude_envelope_certificate
9	joint_dissipation_certificate, piecewise_riccati
10	parameter_set_extension, joint_state_parameter_extension, taylor_parameter_extension, local_taylor_cells, continuous_cell_propagation
11	validated_joint_propagation, nonlinear_taylor_validated, validated_centered_taylor, affine_nonlinear_parameter, strict_affine_validation
12	affine_wrapping_experiment, moving_frame_benchmark, nonlinear_moving_affine, lohner_generator_experiment, continuous_affine_validated, continuous_lohner_validated, dyadic_viability, dyadic_affine_combined, long_horizon_affine_validation, adaptive_affine_reconditioning, modal_frame_validated, fair_affine_comparison, matrix_frame_test, moving_frame_reconditioning, adaptive_lohner_frame, observable_aware_dual_enclosure

Table B.2: Evidence kind and arithmetic backend by chapter.

Chapters	Evidence kind	Backend and check
2–7	Theorems with numerical illustration	Closed-form bounds evaluated in double precision, compared with integrated trajectories (consistency only)
8	Theorems, exact optima, calibration	Exact rationals with upward-rounded square roots at 2^{-80} for the optima; DOP853 integration at relative tolerance 10^{-11} for calibration (consistency only)
9	Theorems with exact-rational evaluation	Rational arithmetic for the tabulated Riccati values
10	Set-membership soundness theorems; sampled screening	Floating-point integration of nominal flows, so retained-cell counts are demonstrations and not certified exclusions
11	Validated enclosures	Outward-rounded multiprecision interval arithmetic, trusted but not machine-verified; Picard inclusion tested in every step
12	Validated enclosures	Exact rationals with outward dyadic rounding at 40 bits unless stated; Picard inclusion tested in every step

Bibliography

[1] J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort and R. Spigler, The Kuramoto Model: A Simple Paradigm for Synchronization Phenomena, *Reviews of Modern Physics* 77 (2005), 137–185.

[2] I. S. Aranson and L. Kramer, The World of the Complex Ginzburg–Landau Equation, *Reviews of Modern Physics* 74 (2002), 99–143.

[3] J.-P. Aubin, *Viability Theory*, Birkhäuser, Boston, 1991.

[4] J.-P. Aubin, A. M. Bayen and P. Saint-Pierre, *Viability Theory: New Directions*, second edition, Springer, Berlin, 2011.

[5] M. Berz and K. Makino, Verified Integration of ODEs and Flows Using Differential Algebraic Methods on High-Order Taylor Models, *Reliable Computing* 4 (1998), 361–369.

[6] F. Blanchini, Set Invariance in Control, *Automatica* 35 (1999), 1747–1767.

[7] M. Fiedler, Algebraic Connectivity of Graphs, *Czechoslovak Mathematical Journal* 23 (1973), 298–305.

[8] T. H. Gronwall, Note on the Derivatives with Respect to a Parameter of the Solutions of a System of Differential Equations, *Annals of Mathematics* 20 (1919), 292–296.

[9] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence*, Springer, Berlin, 1984.

[10] R. J. Lohner, Enclosing the Solutions of Ordinary Initial and Boundary Value Problems, in *Computer Arithmetic: Scientific Computation and Programming Languages*, Teubner, Stuttgart, 1987, 255–286.

[11] D. Lombardo, Energy-Driven Structure in Coupled Brain–Body Dynamics, *Scientific Reports* 16 (2026), 38175, doi:10.1038/s41598-026-72942-9.

[12] M. Milanese and A. Vicino, Optimal Estimation Theory for Dynamic Systems with Set Membership Uncertainty: An Overview, *Automatica* 27 (1991), 997–1009.

[13] R. E. Moore, *Interval Analysis*, Prentice-Hall, Englewood Cliffs, 1966.

[14] R. E. Moore, R. B. Kearfott and M. J. Cloud, *Introduction to Interval Analysis*, SIAM, Philadelphia, 2009.

[15] M. Nagumo, Über die Lage der Integralkurven gewöhnlicher Differentialgleichungen, *Proceedings of the Physico-Mathematical Society of Japan* 24 (1942), 551–559.

[16] N. S. Nedialkov, K. R. Jackson and G. F. Corliss, Validated Solutions of Initial Value Problems for Ordinary Differential Equations, *Applied Mathematics and Computation* 105 (1999), 21–68.

[17] A. Neumaier, The Wrapping Effect, Ellipsoid Arithmetic, Stability and Confidence Regions, *Computing Supplementum* 9 (1993), 175–190.

[18] R. Olfati-Saber and R. M. Murray, Consensus Problems in Networks of Agents with Switching Topology and Time-Delays, *IEEE Transactions on Automatic Control* 49 (2004), 1520–1533.

[19] A. Pikovsky, M. Rosenblum and J. Kurths, *Synchronization: A Universal Concept in Nonlinear Sciences*, Cambridge University Press, Cambridge, 2001.

[20] J. T. Stuart, On the Non-Linear Mechanics of Wave Disturbances in Stable and Unstable Parallel Flows, *Journal of Fluid Mechanics* 9 (1960), 353–370.

[21] W. Tucker, *Validated Numerics: A Short Introduction to Rigorous Computations*, Princeton University Press, Princeton, 2011.