

Exact interfaces for modular autocatalytic reaction networks: amplification prices, catalyst-aware cores, degradation onset, and boundary traces

Anonymous manuscript

September 7, 2026

Abstract

Modular reasoning about an autocatalytic reaction network is trustworthy only when the environment cannot exploit information that a module summary discarded. We develop, and verify in Lean 4, a layered theory of *exact interfaces* for finite reaction networks with explicit catalysis. Three foundation results each identify the smallest object that composes. (i) For the maximum amplification factor (MAF), the compositional invariant is not the scalar optimum but the threshold-indexed family of strict species-price certificates; reaction addition, species aggregation, disjoint union, and shared-species coupling act on these price cones by exact set operations, and a two-species example with MAF values 2, 2, 2, 4 shows that no scalar rule can exist. (ii) Autocatalytic child-selection cores of a network with explicit catalysis are recovered exactly from its ordinary cores: a target core differs from a vertex-contained ordinary anchor by a vertex-disjoint pack of alternating matching-exchange paths, and an executable enumerator built on this fact returns precisely the child-selection cores without traversing the ambient edge powerset. (iii) For species-specific degradation at diluted onset, the growth, critical, and extinction regions of the degradation orthant are exactly the projective images $d \leq Av/v$ of positive compositions v , with the Perron mode constructed rather than assumed, and a finite actuator support can extinguish the system precisely when the unactuated principal block already admits a strict extinction certificate.

Interactions between factors of a source network are then organised on a typed factor-species incidence graph: amplification synergy forces a signed cycle, fixed-threshold feasibility is exactly a disjunction over factors on an incidence forest, and degradation loads of independent branches add at a shared root, so that two individually extinct branches can jointly grow. Finally we define the exact boundary trace of a module, prove that trace equality is precisely indistinguishability by every boundary context, characterise when a coarser summary is safe (injectivity on the realized traces), and give an exact junction-tree message theorem in which running intersection together with a scope-extensionality condition repairs otherwise invalid existential witness gluing; separator width bounds message arity. The MAF price-cone profile, a thermodynamic interval interface, and a finite Type-II passive tail are exact interfaces; the scalar MAF and two thermodynamic verdict vectors are certified information losses. Every stated result is compiled with warnings treated as errors in a frozen Lean 4.30 / Mathlib environment.

Keywords. autocatalytic networks; RAF sets; maximum amplification factor; autocatalytic cores; explicit catalysis; Metzler matrices; degradation; boundary trace; contextual equivalence; junction tree; separator width; Lean 4.

1 Introduction

1.1 The problem

Autocatalytic reaction networks are studied through several structural lenses: reflexively autocatalytic and food-generated (RAF) sets [1–4], stoichiometric autocatalytic motifs and

cores [5–7], diluted-onset dynamical autocatalysis [8, 9], growth bounds such as the maximum amplification factor [10], and thermodynamic realizability of cycles [11]. In each lens a natural engineering question arises: if a network is assembled from modules, which module data suffice to predict the behaviour of the assembly?

The question is harder than it looks because the lenses expose different boundary data. A structural verdict, a productive flux, a spectral sign, and a thermodynamic current pattern are all summaries, and a convenient summary can agree for two isolated modules while a shared partner separates them. Three concrete phenomena drove this work.

- Two two-species networks with the same maximum amplification factor 2 have parallel compositions with a common partner whose amplification factors are 2 and 4 (Theorem 3.8). The scalar loses a *direction*.
- An ordinary autocatalytic core contained in a larger autocatalytic child selection need not use the same reactant matching edges (Example 4.3). Descending along matching edges is wrong; descending along supports is right.
- A single degraded branch coupled to a shared root may be extinct while two copies of it coupled to the same root grow (Example 6.8). Degradation composition is additive, not maximum-like.

Each phenomenon is an instance of one lesson: identify the smallest certificate profile that is functorial under the operations of interest, and prove exact laws on that profile before extracting scalars.

1.2 Contributions

The paper is organised as three foundation layers, an interaction layer, and a general interface theory. All results are formalised; Section 10 lists the terminal Lean declarations and receipts.

- (I) **Amplification prices** (Section 3). Threshold feasibility $\exists x \geq 0, x \neq 0, Bx \geq qAx$ is dual, by a strict finite Gordan alternative, to the nonemptiness of a strict price cone $D_N(q)$. For an attained MAF α one has $\alpha < q \iff D_N(q) \neq \emptyset$. Reaction addition intersects $D_N(q)$ with an open half-space, aggregation retains nonnegative lifts, disjoint union gives max, and shared-species coupling intersects price cones. Equality versus strict increase of the MAF is exactly survival versus disappearance of prices above the old threshold. The scalar MAF is provably noncompositional.
- (II) **Catalyst-aware cores** (Section 4). For a finite network with explicit catalysis, every autocatalytic child-selection edge set is reconstructed from a vertex-contained ordinary core by toggling a vertex-disjoint pack of alternating matching-exchange paths with external endpoints of opposite type; unique matching of the anchor excludes alternating cycles. An executable enumerator that generates only such packs, filters with proof-carrying rational certificates, and takes strict-subset minima returns exactly the child-selection cores of Golnik et al. [6], Definition 3.4.
- (III) **Degradation onset** (Section 5). For a diluted unary source model whose reaction graph is strongly connected, and for every nonnegative species-wise degradation d , the sign of the real spectral bound of $A - \text{Diag}(d)$ is characterised by strict positive certificates, equivalently by the projective inequalities $d < Av/v$, $d = Av/v$, $Av/v < d$. The Perron mode is constructed from a finite-dimensional Schauder fixed point. Finite supported degradation on an actuator set S can extinguish the system iff the unactuated principal block admits a strict extinction certificate.
- (IV) **Factor interactions** (Section 6). For an ownership factorization of a source network, amplification feasibility that no single active factor achieves alone forces a signed cycle in

the typed factor–species incidence graph; on an incidence forest, global fixed-threshold feasibility is exactly the disjunction of factor feasibilities and the attained MAF is the maximum of factor MAFs. For a degradation star, an inverse-free Schur elimination shows that branch loads add at the root and the sign of the resulting scalar is the sign of the real spectral bound.

- (V) **Exact boundary traces** (Section 7). The trace of a boundary model is the existential projection of its valid internal witnesses. Trace equality is equivalent to indistinguishability by every boundary context, with singleton contexts sufficing for the converse. A summary of the trace is contextually sufficient iff it is injective on the realized traces. A loss witness refutes sufficiency. The literal source feasibility of an ownership factorization is an exact trace.
- (VI) **Bounded-width composition** (Section 8). On a rooted join tree satisfying running intersection and local scope extensionality, the recursive separator message equals the one-global-assignment semantics at every node; hence separator messages are contextually substitutable, a certificate checker is sound with message arity bounded by separator width, elimination order is irrelevant once global semantics agree, and heterogeneous exact local interfaces may be replaced simultaneously.
- (VII) **Certified instances** (Section 9). The full MAF price-cone profile is an exact interface that composes by intersection, while the scalar is a certified loss. A thermodynamic triangle’s literal current equations project exactly to an open interval. Two source-valid pairs of interacting cores certify successive losses across the hierarchy $\text{MultiPAC} \Rightarrow \text{Dir} \Rightarrow \text{LinCx} \Rightarrow \text{MultiCAC}$. A finite Type-II passive tail is exactly its affine two-port, with unique reconstruction of the private state.

1.3 Scope, non-claims, and naming

The formal statements carry their own scope, and we repeat the main boundaries here because they are part of the results rather than editorial caveats.

- Non-claim 1.1.* (a) The MAF calculus is structural: it concerns stoichiometric amplification, not realized kinetic growth, stability, persistence, or stochastic survival.
- (b) The core enumerator is exact; no polynomial-delay, fixed-parameter, or end-to-end speed-up claim over existing software is made. The number of exchange paths and packs can be exponential.
 - (c) The degradation theorem is an onset theorem for the low-concentration linearization. It says nothing about nonlinear global fate at criticality, and it does not supply a determinant-only criterion for overlapping cycles.
 - (d) Separator width bounds message *arity*, not runtime, storage, description size, or decomposition discovery. It is deliberately not treewidth: only the separator cardinalities of a supplied decomposition are asserted.
 - (e) Scalar summaries are not claimed to compose unless the realized-image criterion is established for them.

On naming: the formalization groups the results of Section 6 under an internal label (“core interaction calculus”, namespace `CoreInteraction`). We do not use that phrase in the mathematics, because the objects there are ownership factors of a source network rather than autocatalytic cores in the sense of [5, 6], and because it is not an established term in the RAF literature. We use literal names: ownership factorization, incidence forest, signed interaction cycle, degradation star. Similarly, “separator width” is used rather than treewidth.

1.4 Organisation

Section 2 fixes conventions. Sections 3–5 are the three foundations, Section 6 the interaction layer, Sections 7–8 the general interface theory, and Section 9 the certified instances. Section 10 documents the formal verification, Section 11 the relation to prior work, and Section 12 open directions.

2 Conventions

Throughout, species and reactions form finite index sets $\mathcal{S}$ and $\mathcal{R}$. A *reaction network with explicit catalysis* is a triple of matrices (A, B, K) indexed by $\mathcal{S} \times \mathcal{R}$: the input (reactant) matrix $A \geq 0$, the output (product) matrix $B \geq 0$, and a catalyst matrix $K \geq 0$ recording species that participate without net consumption. The net stoichiometric matrix is $S = B - A$. Keeping A and B separately matters: a catalyst that appears on both sides cancels in S but still constrains which species–reaction pairs are legitimate reactant edges, and Section 4 shows that erasing it changes the core structure.

For vectors we write $x \geq 0$ for componentwise nonnegativity, $x \gg 0$ for strict positivity of every coordinate, and $x \ll 0$ for strict negativity. A vector with $x \geq 0$ and $x \neq 0$ is called *semipositive*. A matrix M is *Metzler* if its off-diagonal entries are nonnegative. The real spectral bound $s(M)$ of a real square matrix is the maximum real part of its complex eigenvalues; in the formalization the statement “ λ is the real spectral bound of M ” means that λ is a real eigenvalue of M and every complex eigenvalue μ satisfies $\operatorname{Re} \mu \leq \lambda$.

A subnetwork is *stoichiometrically autocatalytic* if there is a semipositive flux x with $Sx \gg 0$ on its species [5, 7]. RAF sets [2, 4] are the food-generated, reflexively autocatalytic subsets of a catalytic reaction system; the results below are stated at the stoichiometric and kinetic level and are compatible with RAF structure through the explicit catalyst matrix, but they do not presuppose a food set.

3 Amplification prices: an exact composition and edit calculus for the MAF

The maximum amplification factor of a network is a von Neumann-type max–min quantity [12, 13] introduced for reaction networks by Gagrani et al. [10], who identify composition rules for adding reactions, merging species, and coupling motifs as an open direction. We show that no scalar rule exists and that the correct compositional state is a threshold-indexed family of strict price certificates.

3.1 Threshold feasibility and attained MAF

Definition 3.1 (Threshold feasibility). For a network $N = (A, B)$ and a real threshold q ,

$$\operatorname{Feas}_N(q) :\iff \exists x \geq 0, x \neq 0 : \quad Bx \geq qAx \text{ componentwise.}$$

A real α is an *attained MAF* of N , written $\operatorname{IsMAF}(N, \alpha)$, if $\operatorname{Feas}_N(\alpha)$ holds and every feasible q satisfies $q \leq \alpha$.

The division-free form is exactly the source ratio convention once species with zero total input under x are handled correctly; this equivalence is part of the formalization. Making attainment explicit exposes an assumption that operation theorems need, rather than hiding it in a supremum. When $A \geq 0$, feasibility is antitone in q .

3.2 Strict duality and the price profile

Let C be a real $\mathcal{S} \times \mathcal{R}$ matrix. Write $P(C)$ for “ $\exists x \geq 0, x \neq 0, Cx \geq 0$ ” and $D(C)$ for “ $\exists p \geq 0, C^\top p \ll 0$ ”.

Theorem 3.2 (Strict finite Gordan alternative). $P(C) \iff \neg D(C)$.

This is a classical theorem of the alternative [14, 15]; what matters here is the placement of strictness, which is dictated by the primal statement: species inequalities are weak, so the dual is strict on every reaction column. Applying it to $C(q) = B - qA$ gives the certificate of threshold infeasibility.

Definition 3.3 (Strict price cone). $D_N(q) := \{p \geq 0 : (B - qA)^\top p \ll 0\}$; equivalently $p \geq 0$ and $p^\top B_r < qp^\top A_r$ for every reaction r . The *price profile* of N is the set-valued map $q \mapsto D_N(q)$.

Corollary 3.4 (Above-threshold exactness). $\neg \text{Feas}_N(q) \iff D_N(q) \neq \emptyset$. If $A \geq 0$ and α is an attained MAF of N , then for every q ,

$$\alpha < q \iff D_N(q) \neq \emptyset.$$

The corollary is the engine of the calculus. It avoids any dual optimum at $q = \alpha$; the entire profile over $q > \alpha$ is used instead, which is both stronger in scope and simpler to formalise.

3.3 The operation calculus

Fix networks with nonnegative input stoichiometry and attained MAFs where stated.

Theorem 3.5 (Exact operation laws). (i) Reaction addition. Let $N + r$ add a reaction with input column a and output column b , and let $H_r(q) = \{p : p^\top b < qp^\top a\}$. Then $D_{N+r}(q) = D_N(q) \cap H_r(q)$ for every q . If α, β are attained MAFs of $N, N + r$ then $\alpha \leq \beta$, and

$$\alpha < \beta \iff \exists q > \alpha : D_N(q) \cap H_r(q) = \emptyset.$$

(ii) Species aggregation. Let $C \geq 0$ be an aggregation matrix and $CN = (CA, CB)$. Then $p \in D_{CN}(q) \iff p \geq 0 \wedge C^\top p \in D_N(q)$. Aggregation cannot lower the MAF, and it strictly raises it iff for some $q > \alpha$ no aggregate price lifts into $D_N(q)$.

(iii) Disjoint union. If $\text{IsMAF}(N_1, a)$ and $\text{IsMAF}(N_2, b)$ then $\text{IsMAF}(N_1 \oplus N_2, \max(a, b))$.

(iv) Shared-species coupling. For N_1, N_2 on a common species set, $D_{N_1 \parallel N_2}(q) = D_{N_1}(q) \cap D_{N_2}(q)$. If a, b, c are attained MAFs of $N_1, N_2, N_1 \parallel N_2$ then $\max(a, b) \leq c$, and

$$\max(a, b) < c \iff \exists q > \max(a, b) : D_{N_1}(q) \cap D_{N_2}(q) = \emptyset.$$

Proof sketch. Each pointwise identity is a direct computation on columns: an old flux embeds in the enlarged network by a zero new coordinate, an old strict price remains strict on old columns, and the new column contributes exactly one strict half-space. For aggregation, $p^\top (CB - qCA) = (C^\top p)^\top (B - qA)$ lifts prices, while primal inequalities are combined nonnegatively by C . For parallel coupling a dual certificate must be a single price vector strict on both column families. The global statements quantify the pointwise identities over q above the relevant MAF and apply Corollary 3.4. $\square$

The one-common-price quantifier in (iv) is essential: “there is a price for N_1 and there is a price for N_2 ” is not enough, and this is precisely where a scalar summary fails. Routine monotonicity results factor through a preorder: N *threshold-simulates* M if every q feasible in N is feasible in M ; input decrease, output increase, reaction addition, parallel inclusion, aggregation, and direct-sum embeddings are instances, and once both MAFs are attained simulation implies $\alpha(N) \leq \alpha(M)$.

Proposition 3.6 (Attainment). *If $B \geq 0$ and the feasible thresholds are bounded above, an attained MAF exists. In particular, if every reaction consumes positive total input and total product in each column is at most U times total reactant, then every feasible q satisfies $q \leq U$ and an attained MAF exists.*

Proof sketch. Normalize any semipositive flux to the standard simplex; the set of normalized feasible pairs (q, x) with $q \leq U$ is nonempty and compact, so the threshold coordinate attains a maximum. $\square$

3.4 Two counterexamples that fix the theory

Example 3.7 (The weak cone is not a certificate). Consider $A \rightarrow 2A$ and $B \rightarrow B$. The MAF is 2, but at $q = 1$ the weak price $p = (0, 1)$ satisfies $B^\top p \leq qA^\top p$: it prices only the stagnant component. Hence the natural weak semipositive cone $\{p \geq 0, p \neq 0 : B^\top p \leq qA^\top p\}$ can be nonempty far below the MAF, and the strictness in Theorem 3.2 cannot be relaxed on reducible networks.

Theorem 3.8 (Scalar MAF is not compositional). *Let $N_1 = \{A \rightarrow B, B \rightarrow 4A\}$ and $N_2 = \{A \rightarrow 4B, B \rightarrow A\}$ on species $\{A, B\}$. Then N_1, N_2 , and $N_1 \parallel N_2$ have attained MAF 2, while $N_2 \parallel N_1$ has attained MAF 4. Consequently no function of the component MAFs determines the MAF of a shared-species parallel composition.*

Proof. With fluxes (x, y) on the two reactions, N_1 is feasible at q iff $qx \leq 4y$ and $qy \leq x$ for some semipositive (x, y) ; $(x, y) = (2, 1)$ is feasible at 2, and adding the inequalities with weights 1, 2 gives $q(x + 2y) \leq 2(x + 2y)$, so $q \leq 2$. Symmetrically N_2 has MAF 2 with witness $(1, 2)$. For $N_1 \parallel N_1$ with fluxes (a, b, c, d) the constraints are $q(a + c) \leq 4(b + d)$ and $q(b + d) \leq a + c$, so the same weighting gives $q \leq 2$, attained by $(2, 1, 0, 0)$. For $N_2 \parallel N_1$ the constraints are $q(a + c) \leq b + 4d$ and $q(b + d) \leq 4a + c$; the flux $(1, 0, 0, 1)$ fires $A \rightarrow 4B$ and $B \rightarrow 4A$ and is feasible at 4, while summing the two constraints gives $q(a + b + c + d) \leq 4(a + b + c + d)$. $\square$

The missing information is directional: which species-price patterns obstruct each module. The profile $D_N(q)$ retains it; the scalar discards it. Section 6 shows that this synergy also leaves a structural footprint, a signed cycle in the factor–species incidence graph.

3.5 Strict productivity is a different alternative

MAF feasibility is weak on species and therefore has a strict dual. Strict all-species net production is strict on species and therefore has a *weak* nonzero dual:

Theorem 3.9 (Strict productivity alternative). $\exists x \geq 0, x \neq 0 : Sx \gg 0 \iff \neg \exists p \geq 0, p \neq 0 : S^\top p \leq 0$.

Applied to an augmented matrix $[S \ d]$ this decides a one-reaction transition at $q = 1$. It also blocks a tempting inference: $\alpha > 1$ in a reducible network need not mean that every species is strictly produced. The same kernel specialises, by splitting each reaction into two nonnegative orientations, to Stiemke’s alternative for sign-unconstrained flows (a strictly productive free flow exists iff no semipositive conservation vector does), which is reused elsewhere in our formalization with no further duality work.

4 Catalyst-aware cores: direct extraction from ordinary cores

Golnik, Gatter, Stadler, and Vassena [6] define autocatalytic cores for networks with explicit catalysis through *child selections* and ask whether the finer child-selection cores can be obtained from the ordinary cores. We answer this with an exact, executable enumeration theorem.

4.1 Child selections and two orders

Let $Q = (\text{reactant}, \text{product})$ with $\text{reactant}, \text{product} : \mathcal{S} \times \mathcal{R} \rightarrow \mathbb{N}$ and $\text{net} = \text{product} - \text{reactant}$.

Definition 4.1 (Child selection, autocatalyticity). A *child selection* κ consists of finite sets $X \subseteq \mathcal{S}$, $R \subseteq \mathcal{R}$ and a bijection $\sigma : X \rightarrow R$ with $\text{reactant}(x, \sigma(x)) > 0$ for all $x \in X$. Its matrix is $M_\kappa(x, y) = \text{net}(x, \sigma(y))$ for $x, y \in X$. κ is *autocatalytic* if M_κ is semipositive: $\exists v \geq 0$, $v \neq 0$ with $M_\kappa v \gg 0$.

A child selection is represented by its edge set $E_\kappa = \{(x, \sigma(x))\} \subseteq \mathcal{S} \times \mathcal{R}$, a reactant-supported matching in the bipartite reactant graph. Two partial orders are now in play.

Definition 4.2 (Ordinary cores and child-selection cores). (i) The *support* of κ is (X, R) .

A subnetwork (X, R) is autocatalytic if some child selection with that support is autocatalytic. An *ordinary core* is a subnetwork that is autocatalytic and inclusion-minimal (componentwise) among autocatalytic subnetworks.

(ii) An edge set E is autocatalytic if it is the edge set of an autocatalytic child selection. A *child-selection core* (CS core; [6], Definition 3.4) is an autocatalytic edge set none of whose proper subsets is autocatalytic.

Minimality in the support order and minimality in the edge order differ, and the difference is not bookkeeping.

Example 4.3 ([6], Example 3.3). Species x_0, x_1 and an external food species a ; reactions $r_0 : x_0 + x_1 \rightarrow 2x_1$ and $r_1 : a + x_1 \rightarrow 2x_0$. The child selection $x_0 \mapsto r_0$, $x_1 \mapsto r_1$ has matrix $\begin{pmatrix} -1 & 2 \\ 1 & -1 \end{pmatrix}$ and is autocatalytic with witness $v = (3, 2)$, giving $Mv = (1, 1)$; neither one-edge subselection is autocatalytic (each diagonal is -1), so it is a CS core. The support $(\{x_1\}, \{r_0\})$ is a strictly smaller ordinary core, autocatalytic through the *different* edge $x_1 \mapsto r_0$ with matrix $(+1)$. Its single matching edge (x_1, r_0) is not among the target's edges $\{(x_0, r_0), (x_1, r_1)\}$. Hence descent to an ordinary core along matching-edge inclusion fails; descent along supports succeeds and yields only *vertex* containment.

The formalization keeps this example as a mandatory regression (Section 10).

4.2 The exchange theorem

Fix an ordinary core with matching M_G (the *anchor*) and a target child-selection matching M_H whose support contains the anchor's vertices. The symmetric difference $M_G \triangle M_H$ has maximum degree two, so its components are paths and cycles.

Lemma 4.4 (Cycle exclusion). *If M_G is the unique matching on its vertex set, then $M_G \triangle M_H$ is acyclic.*

Proof. An alternating cycle inside the anchor support can be toggled: replacing the M_G -edges of the cycle by its M_H -edges yields a second matching with the same vertex set, contradicting uniqueness. $\square$

Uniqueness of the child-selection matching of an ordinary core is a structural property of ordinary cores in the source setting; the formal theorem takes it as an explicit hypothesis on the supplied anchors.

Theorem 4.5 (Exchange decomposition). *Let M_G be a unique matching on its vertex set, vertex-contained in the target matching M_H of a bipartite species–reaction graph. Then $M_G \triangle M_H$ is a vertex-disjoint family of alternating simple paths; a vertex of its support has*

degree one iff it is external to M_G ; every nontrivial path joins an external species to an external reaction; and every edge of $M_G \triangle M_H$ lies on a simple path between an external species and an external reaction.

Proof sketch. Acyclicity is Lemma 4.4. A vertex matched by both matchings that occurs in the symmetric difference has two exchange neighbours; a vertex external to M_G has exactly one. A path whose first edge is not an M_G -edge and whose last vertex has no M_G -edge has odd length, and an odd path in a bipartite graph joins the two sides. A finite tree of maximum degree two with a leaf has a second leaf, so every external vertex lies on a path to another external vertex, and such a path exhausts its component. $\square$

Contracting each anchor edge to a point, uniqueness of the perfect assignment implies that the contracted exchange graph has no directed cycle, which is the form used by the path search.

4.3 The direct enumerator and its exactness

Construction 4.6 (Direct source-path enumeration). Inputs: the network Q ; a list of anchors (ordinary cores given by their matchings); exhaustive finite orders on species and reactions; and a *certified autocatalytic test*, i.e. a Boolean function on edge sets whose **true** answers are backed by a rational primal certificate ($v \geq 0$, $v \neq 0$, $Mv \gg 0$) for some matching realising the set, and whose **false** answers are backed by a rational Gordan–Stiemke obstruction ($y \geq 0$, $y \neq 0$, $M^\top y \leq 0$) for every matching realising the set.

1. For each anchor, run a depth-first search from every species outside the anchor to reactions outside the anchor, stepping species $\rightarrow$ reaction along reactant edges and reaction $\rightarrow$ species along anchor edges. This yields the *source boundary paths*.
2. Enumerate vertex-disjoint packs of boundary paths by include/exclude recursion, rejecting conflicts as they arise.
3. Toggle the anchor along each pack (delete its backward anchor edges, append its forward edges) and keep the results that are reactant-supported matchings.
4. Apply the certified test, then keep an edge set iff no strictly smaller generated edge set passes the test. Output the resulting family of edge sets.

Theorem 4.7 (Exact direct extraction). *Suppose the anchor list is complete (every ordinary core of Q is the support of some listed anchor) and every listed anchor is a unique matching on its vertex set. Then an edge set is emitted by Construction 4.6 if and only if it is a child-selection core of Q .*

Proof sketch. *Completeness of candidates.* Let E be any autocatalytic edge set with matching M_H . Finite descent in the support order from the support of M_H reaches an ordinary core; anchor completeness supplies a listed anchor M_G with that support, hence vertex-contained in M_H . By Theorem 4.5, $M_G \triangle M_H$ is a pack of external species-to-reaction paths, each of which the search emits (it follows reactant edges and anchor edges and ends at the first external reaction), and toggling M_G by that pack gives exactly E . Thus every autocatalytic edge set is generated, not only the minimal ones. *Exact filtering.* The certified test decides autocatalyticity of edge sets exactly, because a positive certificate proves it and a negative certificate refutes every realisation. An autocatalytic edge set is a CS core iff no strict subset is autocatalytic; since all autocatalytic subsets are generated, checking strict generated subsets is exact. $\square$

Two features deserve emphasis. First, the theorem is an equality between an executable function and a semantic predicate: its left side is computation. Second, it assumes no oracle for “is a CS core”, no hypothesis that the candidate universe contains every core, and no relation

Table 1: Deterministic benchmark: vertex-disjoint pack states of Construction 4.6 against complete reactant-edge-subset enumeration, with identical output families. Direct times exclude source-path generation and are reported only as an indication; the comparison baseline is subset enumeration, not an optimised published implementation.

vertices	reactant edges	subset states	pack states	CS cores
4	11	2 048	146	36
4	12	4 096	213	43
5	15	32 768	294	82
5	16	65 536	763	182
5	18	262 144	989	218

between anchor edges and target edges beyond vertex containment. The rational certificates are exact: strict feasibility of a rational matrix over $\mathbb{R}$ admits a rational witness by density, and the formalization proves this bridge separately.

4.4 Complexity boundary and experiments

The enumerator avoids the ambient reactant-edge powerset, but “direct” is not “polynomial”. On a deterministic benchmark family of small random networks the number of visited pack states was two to three orders of magnitude below the number of edge subsets, with exact agreement of outputs (Table 1). A separate experiment on topologically ordered exchange graphs showed that the number of boundary paths can itself grow exponentially, and an exhaustive falsifier over 190,756 labelled nested-matching pairs through size five found no alternating cycle before the proof identified matching uniqueness as the reason. None of these is evidence for the theorem; they selected the proof route.

5 Species-specific degradation at diluted onset

Unterberger and Nghe [8] derive a first-order onset model for diluted autocatalysis and prove dynamical autocatalysis for sufficiently small degradation; Nandan, Nghe, and Unterberger [9] classify diluted cores into Types I–V and state that no general result characterises dynamical autocatalysis under arbitrary species-specific degradation. We close this gap for the finite strongly connected source model.

5.1 Model

A diluted *unary reaction* $X_j \rightarrow \sum_i s_i X_i$ with rate k contributes $-k$ to entry (j, j) and $s_i k$ to (i, j) of a matrix; a finite list of such reactions with nonnegative rates and products gives the reaction part A , which is Metzler. Static species-wise degradation $d \geq 0$ subtracts $\text{Diag}(d)$:

$$\dot{x} = M(d)x, \quad M(d) = A - \text{Diag}(d).$$

The *source adjacency* carries the positive off-diagonal graph of A ; the source is *irreducible* if that graph is strongly connected. Diagonal degradation never removes an off-diagonal edge, and a large scalar shift makes $M(d) + cI$ nonnegative, so:

Lemma 5.1 (Source irreducibility survives degradation). *If the source is irreducible, then for every d there is c with $M(d) + cI$ nonnegative and irreducible.*

This is the adapter that keeps the public theorem about chemistry rather than about an assumed matrix class.

5.2 Perron existence and sign certificates

Mathlib contains matrix irreducibility but, at the pinned revision, no Perron–Frobenius theorem. We proved the needed finite-dimensional statement directly.

Theorem 5.2 (Perron mode for irreducible Metzler matrices). *If M is Metzler and some scalar shift of M is nonnegative and irreducible, then there are $\lambda \in \mathbb{R}$ and $v \gg 0$ with $\sum_i v_i = 1$, $Mv = \lambda v$, and λ is the real spectral bound of M . Moreover λ is the real spectral bound iff such a normalized positive eigenvector exists, and there is a positive left eigenvector $w \gg 0$ with $M^\top w = \lambda w$.*

Proof sketch. For an entrywise positive P , the map $x \mapsto Px / \sum_i (Px)_i$ is continuous and maps the standard simplex to itself; a finite-dimensional Schauder fixed point [16] gives a normalized positive eigenvector. For nonnegative A , regularize $A + \varepsilon_n J$ with $\varepsilon_n \rightarrow 0$, and pass to a convergent subsequence by compactness of the simplex. If A is irreducible, a zero coordinate of a nonnegative eigenvector would be contradicted by a positive entry of some power A^k along a path from a positive coordinate. Shift back to M . For dominance, conjugate M by $\text{Diag}(v)$: every row of the conjugate sums to λ , so Gershgorin discs give $\text{Re } \mu \leq \lambda$ for every complex eigenvalue. The left vector is obtained by the same construction on $M^\top$, and the pairing $\langle w, v \rangle > 0$ forces its eigenvalue to equal λ . $\square$

Theorem 5.3 (Exact sign certificates). *Under the hypotheses of Theorem 5.2,*

$$\begin{aligned} s(M) > 0 &\iff \exists v \gg 0 : Mv \gg 0, \\ s(M) = 0 &\iff \exists v \gg 0 : Mv = 0, \\ s(M) < 0 &\iff \exists v \gg 0 : Mv \ll 0. \end{aligned}$$

Proof. Forward directions use the Perron vector. Conversely, pair a strict sub- or supersolution z with the positive left vector: $w^\top Mz = \lambda w^\top z$ with $w^\top z > 0$, so the sign of λ is the sign of $w^\top Mz$. $\square$

These certificate characterisations are classical positive-systems theory [17, 18]; the contribution is their formal derivation from the source model and the consequences below. The irreducibility hypothesis is necessary: $\text{Diag}(1, -1)$ has $s = 1$ but no strict positive subsolution, and $\text{Diag}(0, -1)$ has $s = 0$ but no positive nullvector.

5.3 The degradation orthant in projective coordinates

For $v \gg 0$ define the projective degradation $\Phi_A(v)_i = (Av)_i / v_i$. Solving $(A - \text{Diag}(d))v = \lambda v$ coordinatewise gives $d_i = \Phi_A(v)_i - \lambda$.

Theorem 5.4 (Arbitrary degradation regions). *Let the source be irreducible with nonnegative reactions, A its reaction part, and $d \geq 0$. Then*

$$\begin{aligned} s(M(d)) > 0 &\iff \exists v \gg 0 : d < \Phi_A(v) \text{ componentwise}, \\ s(M(d)) = 0 &\iff \exists v \gg 0, \sum_i v_i = 1 : d = \Phi_A(v) (\geq 0), \\ s(M(d)) < 0 &\iff \exists v \gg 0 : \Phi_A(v) < d \text{ componentwise}. \end{aligned}$$

More generally, for every real λ : λ is the real spectral bound of $M(d)$ iff there is a normalized $v \gg 0$ with $d = \Phi_A(v) - \lambda \cdot \mathbf{1}$ and $d \geq 0$.

The spectral predicates on the left contain no eigenvector; the Perron mode is a conclusion. The growth region is a union of open axis-aligned boxes below $\Phi_A(v)$, the extinction region a union of boxes above, and the critical surface is the envelope where equality is attained

by a positive mode. Under a positive-kernel-ray uniqueness hypothesis the normalized map $v \mapsto \Phi_A(v)$ is injective onto the critical surface; the corresponding bijection statement is proved by hand (Remark 5.8). The general projective formula $d_i = (Av)_i/v_i - \lambda$ and minimum-effort diagonal control of Metzler matrices are prior art [19]; what is new is the source-faithful composition and the chemistry-specific consequences.

5.4 Actuator placement

Split the species as $U \sqcup S$, with degradation allowed only on S , and let $A[U, U]$ be the unactuated principal block.

Theorem 5.5 (Partial actuation). *For any Metzler A on $U \sqcup S$: there exists a finite $d \geq 0$ supported on S such that $A - \text{Diag}(d)$ admits a strict extinction certificate iff $A[U, U]$ admits a strict extinction certificate. Under irreducible-shift hypotheses on A and $A[U, U]$, the same holds with “negative real spectral bound” in place of “extinction certificate”.*

Proof sketch. Necessity restricts a global positive supersolution to U ; the discarded cross terms from S are nonnegative, which only helps. Sufficiency is constructive: extend a strict supersolution u on U by a common small $\varepsilon > 0$ on S , chosen so that cross-couplings cannot consume the negative margin on U , then set each actuated degradation above its projective ratio. $\square$

Placement and magnitude separate: a support S can ever stabilise the system only if the uncontrolled principal subsystem is already strictly extinguishable. Explicit 2×2 families illustrate the three cases: $\begin{pmatrix} -(1+\alpha) & 2 \\ 2 & -1 \end{pmatrix}$ crosses at $\alpha = 3$, while $\begin{pmatrix} -(1+\alpha) & 1 \\ 1 & 0 \end{pmatrix}$ and $\begin{pmatrix} -(1+\alpha) & 1 \\ 1 & 1/2 \end{pmatrix}$ never cross.

5.5 Topology-specific compressions and a branch trap

Proposition 5.6 (Type I and Type II diagnostics). *(i) For the Type-I core $X_0 \rightarrow mX_1$ (rate k_1), $X_1 \rightarrow X_0$ (rate k_2) with degradations d_1, d_2 , the onset boundary is $(k_1 + d_1)(k_2 + d_2) = m k_1 k_2$.*

(ii) For the Type-II₁ matrix $\begin{pmatrix} -a & 0 & r \\ p & -b & 0 \\ q & u & -c \end{pmatrix}$, $\det = -abc + bqr + pr u$; with gains $g_S = rq/(ac)$, $g_L = rpu/(abc)$ the Perron root t solves $t^3 = g_S t + g_L$, and $t > 1 \iff g_S + g_L > 1$ via $g_S + g_L - 1 = (t - 1)(t^2 + t + 1 - g_S)$.

(iii) For Type-II₂ with normalized gains (g_1, g_2, g_L) and $z = t^2$, $(z - g_1)(z - g_2) = g_L$, and $z < 1 \iff g_1 < 1, g_2 < 1, g_L < (1 - g_1)(1 - g_2)$. At $(g_1, g_2, g_L) = (2, 3, 2)$ the polynomial vanishes at $z = 1$ and at $z = 4$: the determinant is zero while the Perron root is $t = 2$ and the system grows.

The last item is the *branch trap*: $\det M(d) = 0$ is an algebraic zero set that can select a non-Perron eigenvalue. Determinant and cycle-cover formulas are therefore used after, not instead of, the positive-system characterisation. For every rationally tagged classified core, the characteristic polynomial of its next-generation matrix has the exact signed cycle-cover expansion $\sum_{\sigma} \text{sign}(\sigma) \prod_i (tI - K)_{\sigma(i), i}$, which is a finite algebraic atlas; certified real-root isolation is not part of the formal package.

Proposition 5.7 (Cycle control and robust corners). *(i) For a three-cycle with losses a_i and gains b_i , $\det = -a_0 a_1 a_2 + b_0 b_1 b_2$, and the AM-GM/KKT certificate holds: if $x_i > 0, y_i \geq 0, \mu > 0, (c_i x_i - \mu)(y_i/x_i - 1) \geq 0$ for each i , and $x_0 x_1 x_2 \leq y_0 y_1 y_2$, then $\sum c_i x_i \leq \sum c_i y_i$.*

(ii) If $B \leq B^+$ entrywise, $d^- \leq d$, $v \geq 0$ and $(B^+ - \text{Diag } d^-)v \ll 0$, then $(B - \text{Diag } d)v \ll 0$; dually a lower-corner witness certifies growth over the whole interval box.

Remark 5.8 (Hand-proved extensions). Three statements are proved in the campaign notes but not compiled, and we flag them as such: the n -cycle water-filling law $d_i^* = \max(0, \mu/c_i - k_i)$ with $\prod_i \max(k_i, \mu/c_i) = m \prod_i k_i$; the condensation rule that a strict global growth certificate exists iff every source strongly connected component has positive spectral bound; and the positive-ray limit $\lim_{\alpha \rightarrow \infty} s(A - \alpha \text{Diag } q) = s(A_{ZZ})$ with $Z = \{i : q_i = 0\}$. The projective bijection onto the critical surface is compiled only as surjectivity plus injectivity under an explicit uniqueness hypothesis.

6 Interactions between factors of a source network

The three foundations share a pattern: the compositional object is a certificate profile, and composition acts on it by exact set operations. This section makes the *interaction structure* explicit for a source network cut into factors.

6.1 Ownership factorizations and the typed incidence graph

Definition 6.1 (Ownership factorization). Let $Q = (A, B, K)$ be a source network. A *factorization* of Q assigns to every reaction an owner factor $f = \text{owner}(r)$ and to every factor a support $\text{supp}(f) \subseteq \mathcal{S}$ containing every species with a nonzero input, output, or catalyst entry in some owned reaction. The factor input, output and catalyst matrices are the owned columns of A, B, K , and the *factor residual* at threshold q under flux x is

$$\rho_f(q, x)_s = \sum_{r: \text{owner}(r)=f} (B_{sr} - qA_{sr}) x_r.$$

Ownership is a function, so every reaction lives in exactly one factor and the literal matrices are reconstructed entrywise: $\sum_f A_f = A$, and likewise for B, K .

Lemma 6.2 (Residual split). $\sum_f \rho_f(q, x)_s = ((B - qA)x)_s$ for every species s , with no duplicated or omitted contribution.

The *typed incidence graph* is the bipartite graph on factors and species with an edge $f - s$ iff $s \in \text{supp}(f)$. A species is *private* to f if f is its only supporting factor and an *interface* species if two distinct factors support it; every supported species is one or the other. A factor is *active* under x if some owned reaction has positive flux; it is *locally feasible* at q if it is active and $\rho_f(q, x) \geq 0$; the flux is *globally feasible* if $x \geq 0$, $x \neq 0$ and $(B - qA)x \geq 0$.

6.2 Amplification synergy forces a signed cycle

Theorem 6.3 (Signed interaction cycle). Let x be globally feasible at q and suppose no active factor is locally feasible at q . Then there are a map next on active factors and a species $\text{via}(f)$ for each active f such that $\text{next}(f) \neq f$, both f and $\text{next}(f)$ support $\text{via}(f)$, $\rho_f(q, x)_{\text{via}(f)} < 0$ and $\rho_{\text{next}(f)}(q, x)_{\text{via}(f)} > 0$, and some iterate of next repeats. Consequently the typed incidence graph contains a cycle: it is not a forest.

Proof sketch. A private species of f has $\rho_f = \rho_{\text{global}} \geq 0$, so an active factor that is not locally feasible has a strictly negative residual at some *interface* species it supports. By Lemma 6.2 a negative summand in a nonnegative sum has a distinct positive compensator at the same species. This defines a total signed step on the finite set of active factors, hence a repeated state by pigeonhole. The signs are what make the closed walk a genuine cycle: the walk

$f \rightarrow \text{via}(f) \rightarrow \text{next}(f) \rightarrow \text{via}(\text{next}(f)) \rightarrow \dots$ cannot immediately retrace an edge, because ρ is positive at $\text{via}(f)$ for $\text{next}(f)$ and negative at $\text{via}(\text{next}(f))$ for $\text{next}(f)$, so these species differ. $\square$

The first attempted proof used only the repeated state and failed: a closed walk in an undirected graph may retrace an edge. The sign labels are the missing evidence, and the Lean proof records exactly this repair.

Theorem 6.4 (Exact composition on incidence forests). *If the typed incidence graph of a factorization is a forest, then for every q*

$$\exists x \text{ globally feasible at } q \iff \exists f \exists x \geq 0 \text{ with } f \text{ locally feasible at } q \text{ under } x.$$

If moreover each factor has an attained threshold q_f and $f_{\max}$ attains $\max_f q_f$, then $q_{f_{\max}}$ is the attained threshold of the whole network.

Proof. Forward is the contrapositive of Theorem 6.3, and it exposes a locally feasible factor at the same q and the same x . Backward, restrict x to the owner fibre of f ; the global residual of the restriction equals ρ_f . $\square$

Example 6.5 (The $2 + 2 \rightarrow 4$ synergy is a cycle). Realise Theorem 3.8 as a source network on species $\{A, B\}$ with four reactions $A \rightarrow 4B$, $B \rightarrow A$, $A \rightarrow B$, $B \rightarrow 4A$, factor 0 owning the first two (N_2) and factor 1 the last two (N_1), both supporting both species. The flux $(1, 0, 0, 1)$ is globally feasible at $q = 4$, neither factor is locally feasible at 4 (factor 0 fails at A , factor 1 at B), and the incidence graph is $K_{2,2}$, whose 4-cycle is the signed interaction cycle. The synergy of Section 3 is not an accident of the scalar; it necessarily leaves this footprint.

6.3 Degradation stars and additive Schur loads

Definition 6.6 (Degradation star). A *star* consists of a root diagonal entry $m \in \mathbb{R}$ and, for each branch i in a finite set I , an internal matrix A_i on a finite coordinate set J , a row of couplings $b_i : J \rightarrow \mathbb{R}$ from the branch to the root, and a column $c_i : J \rightarrow \mathbb{R}$ from the root to the branch. Its matrix on $\{*\} \sqcup (I \times J)$ is the arrowhead matrix with root entry m , off-root blocks b_i and c_i , and no branch–branch coupling. A *solve witness* is $u = (u_i)_{i \in I}$ with $A_i u_i = -c_i$ for every i ; the *Schur load* is $\kappa(u) = m + \sum_i \sum_j b_{ij} u_{ij}$.

Theorem 6.7 (Schur load trichotomy). *Let u be a solve witness with $u \gg 0$ and $c_i \gg 0$ for all i , and let $v = (1, u)$. Then $Mv = \kappa(u) e_*$. Hence $\kappa > 0$, $= 0$, < 0 yield a strict growth certificate, a positive nullvector, and a strict extinction certificate respectively, by rescaling the root coordinate. If moreover M is Metzler with an irreducible nonnegative shift, then the sign of $\kappa(u)$ equals the sign of the real spectral bound of M . The same holds when M is presented as $A - \text{Diag}(d)$ from a list of nonnegative unary reactions with irreducible source graph.*

Proof sketch. Direct multiplication: the solve witness cancels every internal row, and the root row is $m + \sum b_{ij} u_{ij}$. Pairing $Mv = \kappa e_*$ with the positive left Perron vector w of Theorem 5.2 gives $\lambda \langle w, v \rangle = \kappa w_*$ with both weights positive. No inverse is formed. $\square$

Example 6.8 (Local extinction, global growth). Take $m = -1$, one internal coordinate per branch with $A_i = (-15)$, $b_i = c_i = 3$, so the solve witness is $u_i = 1/5$. One branch: $M = \begin{pmatrix} -1 & 3 \\ 3 & -15 \end{pmatrix}$, $\kappa = -1 + 3/5 = -2/5 < 0$, and $v = (100, 21)$ gives $Mv = (-37, -15) \ll 0$: extinction. Two identical branches on the same root: $\kappa = -1 + 2 \cdot 3/5 = 1/5 > 0$, and $v = (100, 19, 19)$ gives $Mv = (14, 15, 15) \gg 0$: growth. Duplicating an extinct branch at a shared root flips the sign. Degradation composition is additive in branch loads, not maximum-like, and this exact rational falsifier ruled out a max-type law before any formal work.

The interaction layer thus supplies, for each foundation, a composition law on the certificate profile: intersection of price cones on shared species, disjunction over factors on forests, and addition of Schur loads at a root. What remains open here is the general case: simple-cycle gain formulas, cactus-shaped incidence graphs, and degradation trees of depth greater than one (Section 12).

7 Exact boundary traces and contextual sufficiency

We now abstract the pattern. A module exposes a boundary; the environment sees only the boundary; the question is which boundary data license substitution in every environment.

7.1 Boundary models and traces

Definition 7.1 (Boundary model, trace, context). Fix a parameter type Θ , an internal-state type I , and a boundary type $\mathcal{B}$. A *boundary model* is a predicate $V_M(\theta, i, b)$. Its *trace* is

$$\text{Tr}_M(\theta) = \{b \in \mathcal{B} : \exists i \in I, V_M(\theta, i, b)\}.$$

A *context* is any relation $C(\theta, b)$. M is *compatible* with C at θ if some $b \in \text{Tr}_M(\theta)$ satisfies $C(\theta, b)$.

No finiteness, topology, measure, decidable equality, or enumeration hypothesis occurs anywhere in this section.

Theorem 7.2 (Exact substitution). *If $\text{Tr}_{M_1}(\theta) = \text{Tr}_{M_2}(\theta)$ for every θ , then M_1 and M_2 have equal compatibility verdicts for every context and every θ .*

Theorem 7.3 (Canonical trace characterisation). *$\text{Tr}_{M_1} = \text{Tr}_{M_2}$ pointwise if and only if no context distinguishes M_1 from M_2 . The converse direction uses only singleton contexts $C_b(\theta, b') : \Leftrightarrow b' = b$.*

Proof. Forward is Theorem 7.2. Conversely, fix θ and b ; compatibility with C_b is exactly $b \in \text{Tr}(\theta)$, so equal verdicts give equal traces by extensionality. $\square$

Thus the exact trace is precisely the observational quotient induced by the full boundary-context language. This also explains why counterexample searches can be small: to refute a proposed trace equality one needs a single boundary assignment in the symmetric difference, whose singleton context is automatically separating.

7.2 When is a coarser summary safe?

Definition 7.4 (Contextual sufficiency). Let $R : X \rightarrow \Sigma$ be a representation of a family X of modules, $\text{plug} : C \times X \rightarrow W$ an assembly operation, and $\text{observe} : W \times \Theta \rightarrow \text{Prop}$ an observable. R is *contextually sufficient* if $R(x) = R(y)$ implies $\text{observe}(\text{plug}(c, x), \theta) \Leftrightarrow \text{observe}(\text{plug}(c, y), \theta)$ for all c, θ .

Sufficiency is upward closed under refinement (if the coarse value is a function of the fine value and the coarse representation is sufficient, so is the fine one), and the quotient by observational equivalence is the canonical sufficient representation.

Theorem 7.5 (Realized-image criterion). *Let $x \mapsto M_x$ assign a boundary model to each realized module and let σ map trace families $\Theta \rightarrow \mathcal{P}(\mathcal{B})$ to summary values. The representation $x \mapsto \sigma(\text{Tr}_{M_x})$ is contextually sufficient for every boundary context if and only if σ is injective on $\{\text{Tr}_{M_x} : x \text{ realized}\}$.*



safe coarsening $\iff \sigma$ injective on realized traces (Thm. 7.5)

collision + separating context = certified loss

Figure 1: The exact decision boundary for coarsening. The trace is canonical (Theorem 7.3); the summary is conditional on the realized model class.

Proof. If summaries agree and σ is injective on realized traces, the traces agree and Theorem 7.2 applies. Conversely, sufficiency makes equal-summary modules indistinguishable in every context, and Theorem 7.3 forces equal traces. $\square$

The restriction to realized traces is exact: global injectivity would reject safe summaries merely because they collide on traces the model class never produces, and anything weaker admits a realized collision that a singleton context separates.

Definition 7.6 (Loss witness). A *loss witness* for $(R, \text{plug}, \text{observe})$ consists of modules x, y with $R(x) = R(y)$, a context c , a parameter θ , a proof that $\text{observe}(\text{plug}(c, x), \theta)$ holds, and a proof that $\text{observe}(\text{plug}(c, y), \theta)$ fails.

A loss witness refutes contextual sufficiency immediately. Positive interfaces and negative witnesses are treated symmetrically in what follows.

7.3 Source-faithful factorization is an exact trace

Theorem 7.7 (Source trace exactness). *For an ownership factorization of Q , let the boundary model have parameter q , internal witness the family of owned residuals $(\rho_f)_f$, boundary the reaction flux x , and validity “ $x \geq 0$, $x \neq 0$, each ρ_f equals the owned residual of x , and $\sum_f \rho_f \geq 0$ ”. Then $x \in \text{Tr}(q)$ iff x is globally feasible at q . Consequently the literal source compatibility with any outside predicate on the flux boundary equals trace compatibility, and any two certified ownership factorizations of the same network are interchangeable in every flux-boundary context.*

Property specificity is essential here: a structural RAF factor, an amplification price factor, and a thermodynamic current factor need not expose the same boundary, and the calculus is generic over the relation but refuses to erase the adapter that states what the relation means. This is what makes the loss witnesses of Section 9 theorems about literal reaction families rather than about freely chosen Boolean predicates.

8 Bounded-width composition on rooted join trees

8.1 Trees, messages, and two semantics

Definition 8.1 (Rooted join tree). Over variables Var with decidable equality and values Val , a *rooted join tree* is a binary tree each of whose nodes carries a bag $\beta \subseteq \text{Var}$, a separator $\sigma \subseteq \beta$, a scope $\varsigma \subseteq \beta$, and a local relation $\ell(\theta, a)$ on assignments $a : \text{Var} \rightarrow \text{Val}$; a node is a leaf or has two children. allVars of a node is its bag united with the variables of its subtrees. Its *width* is the maximum separator cardinality over the tree.

Definition 8.2 (Running intersection and local scoping). A tree satisfies *running intersection* if, recursively at every branch with bag β and children L, R ,

$$\text{allVars}(L) \cap (\beta \cup \text{allVars}(R)) \subseteq \sigma_L, \quad \text{allVars}(R) \cap (\beta \cup \text{allVars}(L)) \subseteq \sigma_R.$$




Figure 2: The central repair. Independent existential child witnesses do not glue; running intersection routes each cross-subtree shared variable through separators, and local-scope extensionality makes the splice semantically invisible.

It is *locally scoped* if every local relation depends extensionally only on its declared scope: assignments agreeing on ς receive the same verdict.

Definition 8.3 (Computed and global messages). The *computed message* of a node at θ to an outside assignment o is defined recursively: there is an assignment h agreeing with o on the separator, satisfying the local relation, and (at a branch) such that both children have computed messages to h . The *global semantic message* asserts one global assignment g agreeing with o on the node's separator and satisfying every local relation in the subtree.

The computed message is what a local algorithm evaluates; the global message is what it is meant to mean. They are not automatically equal.

Example 8.4 (Failed gluing). Let a parent P have children L and R that share a variable x not routed through their separators, with L 's relation forcing $x = 0$ and R 's forcing $x = 1$. Each child has a computed message to any outside assignment, so the parent does. But no single global assignment satisfies both, so the global message is false. The recursive semantics is too weak precisely when a relation can inspect an out-of-scope variable or when a shared variable bypasses a separator (Figure 2).

8.2 Exact messages, substitutability, and width

Theorem 8.5 (Exact message semantics). *If a rooted join tree satisfies running intersection and is locally scoped, then at every node, parameter, and outside assignment the computed message holds iff the global semantic message holds.*

Proof sketch. Induct on the tree. At a branch with local witness h and global child witnesses g_L, g_R , define the splice

$$g(v) = \begin{cases} g_L(v) & v \in \text{allVars}(L), \\ g_R(v) & v \in \text{allVars}(R) \setminus \text{allVars}(L), \\ h(v) & \text{otherwise.} \end{cases}$$

Running intersection proves agreement on all overlaps: a variable of L that also occurs in the bag or in R lies in σ_L , where g_L agrees with h , and symmetrically. Scope extensionality transports each local predicate from its own witness to g . The reverse direction restricts a single global witness to each recursive obligation. $\square$

Removing either hypothesis reintroduces Example 8.4. The weaker dependent witness-tree semantics is provably equal to the computed message without hypotheses, but it is not the intended meaning; identifying this gap and isolating the locality condition was the central correction of the formal development.

Corollary 8.6 (Separator contextual sufficiency). *Under the hypotheses of Theorem 8.5, the boundary model “computed message” and the boundary model “one global assignment agreeing on the separator” have equal traces, so the recursive message may replace the literal subtree in every outside separator context.*

Theorem 8.7 (Bounded-width checker soundness). *A message certificate for a node is a relation description together with proofs that its denotation is sound and complete for the computed message. Under running intersection, local scoping, and width $\leq w$, every certificate denotes exactly the global semantic message, and the node’s separator has cardinality at most w .*

The theorem is representation-neutral: a relation of low arity may still have an expensive description or an infinite domain. What is controlled is the number of exposed coordinates, which is exactly what contextual substitution requires. A rational-linear description language (finite conjunctions of atoms $\sum_i c_i x_i \leq \rho$ with an explicit strictness flag, so that projection cannot silently close an open face) is provided as one instance.

Corollary 8.8 (Decomposition invariance). *Two certified rooted join trees whose global semantic messages agree pointwise compute extensionally identical messages, regardless of shapes, bags, or elimination orders.*

Theorem 8.9 (Simultaneous exact replacement). *Let $(L_i)_{i \in \mathcal{I}}$ be an arbitrary indexed family of exact local interfaces: each carries its own code type, a code, a description relation, a literal relation, and a proof that the description of its code equals its literal relation pointwise. Then the conjunction of described relations equals the conjunction of literal relations, and the described and literal assemblies are indistinguishable by every boundary context.*

Heterogeneity is the point: members may use genuinely different description languages, and pointwise exactness converts the conjunction of encodings into the conjunction of literal relations, after which Theorem 7.3 lifts equality to all contexts.

9 Certified instances: safe compression and certified loss

9.1 The amplification price profile

Theorem 9.1 (MAF profile interface). *For networks on a common species set, the representation $N \mapsto (q \mapsto D_N(q))$ is contextually sufficient for every price-boundary context, and it composes under shared-species coupling by pointwise intersection, $D_{N_1 \parallel N_2}(q) = D_{N_1}(q) \cap D_{N_2}(q)$.*

Theorem 9.2 (Scalar MAF loss). *The constant summary assigning the scalar 2 to both N_1 and N_2 of Theorem 3.8 admits a loss witness: with common partner N_1 and parameter $q = 4$, $N_2 \parallel N_1$ is feasible while $N_1 \parallel N_1$ is not, because $4 \leq 2$ is false. Hence the scalar summary is not contextually sufficient.*

In the formalization the parallel composition with the fixed partner is encoded in the observable and the context type is a singleton; the instance is a faithful but schematic loss witness, and the compositional content is carried by Theorem 9.1. The construction is not used to weaken the positive theorem; it identifies the exact place where a popular compression loses information and directs the interface back to the price cone.

9.2 A thermodynamic interval interface and two layered losses

The thermodynamic material comes from a companion study [20] of the Strong Compatibility hypothesis of Kosc et al. [11] for families of interacting primitive autocatalytic cores in reversible mass-action networks with barrier-weighted currents $j_r = b_r(\mathbf{z}^{\text{react}(r)} - \mathbf{z}^{\text{prod}(r)})$. Four verdicts form a hierarchy of increasing strength for a family F of oriented cores:

- $\text{MultiPAC}(F)$: a flow v with positive oriented components making every core productive;
- $\text{Dir}(F)$: a species potential μ with strictly positive oriented affinity on every reaction;
- $\text{LinCx}(F)$: independent positive activities per *complex* realising oriented currents that make every core productive;
- $\text{MultiCAC}(F)$: one positive species-activity vector whose monomials realise such currents (the common species-activity lift).

$\text{MultiCAC} \Rightarrow \text{LinCx}$, and each step of the hierarchy is strict.

Theorem 9.3 (Exact interval interface). *For the triangle $A \rightarrow B \rightarrow C \rightarrow mA$ with gain $m > 1$ and barriers $b_0, b_1, b_2 > 0$, let the literal model have internal currents (j_0, j_1, j_2) and boundary q , valid when $j_0 > 0$, $mj_2 - j_0 > 0$, $j_1 - j_2 > 0$, $j_0 - j_1 > 0$ (positive net production of A, C, B) and $q = b_0(j_1/(b_1j_0) + j_2/(b_2j_0))$. Then the trace of the literal model is exactly the open interval*

$$\frac{b_0(1/b_1 + 1/b_2)}{m} < q < b_0\left(\frac{1}{b_1} + \frac{1}{b_2}\right),$$

so the literal current model and the interval model are contextually substitutable.

Two such triangles sharing $A \rightarrow B$ with left and right coefficients R_L, R_R admit a common response iff $1/m_L < R_R/R_L < m_R$; for gain 2 on both sides this is the window $1/2 < k < 2$ of the barrier ratio k , and the companion study proves the complete phase diagram $\text{MultiPAC} \wedge \text{Dir} \wedge (\text{MultiCAC} \iff 1/2 < k < 2)$ for that family.

Theorem 9.4 (Two strict thermodynamic losses). *Let the control family be the two-triangle family at the interior point $k = 1$ (species A, B, C, D ; reactions $A \rightleftharpoons B$, $B \rightleftharpoons C$, $C \rightleftharpoons 2A$, $B \rightleftharpoons D$, $D \rightleftharpoons 2A$, all barriers 1), which satisfies all four verdicts. Then:*

- (i) (magnitude loss) *the same chemistry with right barriers 10 satisfies MultiPAC and Dir but not LinCx ; the constant two-verdict summary (true, true) admits a loss witness at the observable LinCx ;*
- (ii) (toric loss) *the two-species family $A \rightleftharpoons B$ (barrier 1), $B \rightleftharpoons 2A$ (barrier 1), $B \rightleftharpoons 3A$ (barrier 2), with cores $\{A \rightleftharpoons B, B \rightleftharpoons 2A\}$ and $\{A \rightleftharpoons B, B \rightleftharpoons 3A\}$, satisfies MultiPAC , Dir , and LinCx but not MultiCAC ; the constant three-verdict summary admits a loss witness at the observable MultiCAC .*

The toric obstruction is elementary once found: left productivity forces $A > B > A^2$, hence $0 < A < 1$ and $B - A^3 > B - A^2$, which contradicts the right core's requirement $2(B - A^3) < A - B$ against the left's $B - A^2 > (A - B)/2$. Both witnesses pair a positive source theorem with a negative source theorem in the generic loss-witness structure; the summary maps are constant by construction, and the mathematical content is the paired verdicts (Figure 3).

9.3 Finite Type-II passive tails are exact two-ports

Type-II cores in the classification of [5, 9] may carry passive reversible tails. In the companion unistationarity study [21] a *unit chain* is built inductively: a direct edge with forward and backward rates $c, \beta > 0$, or an extension of a chain by an edge (c, β) that turns the previous right endpoint into an intermediate with degradation $d > 0$. Its *summary* $(c, \beta, \ell_L, \ell_R)$ is defined by the recursion, with prefix summary $(c_q, \beta_q, \ell_q, r_q)$ and $L = \beta_q + r_q + c + d > 0$,

$$c' = \frac{c_q c}{L}, \quad \beta' = \frac{\beta_q \beta}{L}, \quad \ell'_L = \ell_q + \frac{c_q(r_q + d)}{L}, \quad \ell'_R = \frac{\beta(r_q + d)}{L},$$




Figure 3: Two source-valid abstraction-loss cuts. Green nodes are shared positive layers; each red arrow marks the first verdict that separates the paired sources of Theorem 9.4.

starting from $(c, \beta, 0, 0)$. The chain-flux relation states that the internal concentrations are steady, j_L leaves the left endpoint with concentration X , and j_R enters the right endpoint with concentration Y .

Theorem 9.5 (Exact passive-tail compression and reconstruction). *For every finite unit chain and all boundary data (X, Y, j_L, j_R) : an internal steady state exists iff*

$$j_L = (c + \ell_L)X - \beta Y \quad \text{and} \quad j_R = cX - (\beta + \ell_R)Y.$$

Whenever a steady state exists it is determined by (X, Y) alone. Consequently the literal chain and its four-coefficient affine two-port have identical boundary traces and are indistinguishable by every boundary context.

Proof sketch. Forward elimination is the known Schur-complement step. The converse is recursive: from the two compressed equations, set $Z = (c_q X + \beta Y)/L$ and $j_{\text{mid}} = c_q X - (\beta_q + r_q)Z$; these satisfy the two uncompressed component equations, the node balance $j_{\text{mid}} - cZ + \beta Y - dZ = 0$, and the output equation, and the prefix is reconstructed by induction. Positivity of L is the strictly positive total outflow of the appended intermediate. $\square$

Internal degradation of a reversible chain is therefore boundary-observable *only* as two endpoint leaks, which are absorbed additively into the endpoint species' degradation rates when the tail is inserted into a core. This is an exact compression in the sense of Section 7, in contrast with the scalar MAF and the thermodynamic verdict vectors.

10 Formal verification

10.1 Environment and discipline

All theorems are formalised in Lean 4 [22] against Mathlib [23]. The environment is frozen: toolchain `leanprover/lean4:v4.30.0`, Mathlib at commit `c5ea0035` (full: `c5ea00351c28e24afc9f0f84379aa41082b1188f`), and a Lake manifest with SHA-256 `a8df0b606ec258561c226df89856884be433c19b49fb0a27335d69bb2b6e9fac`. The amplification-price receipt of Table 2 was re-issued on the date of this manuscript after a universe generalisation of three of its modules; all other receipts match the current sources. Every module is compiled with `-DwarningAsError=true`; a verification receipt records the exact command, exit code 0, empty standard output and error, the source hash, the transitive local dependency list, the compiled artifact hash, and a fingerprint of the elaborated type of each exported declaration. No declaration uses `sorry`.

10.2 Trusted base and provenance

The degradation, boundary-trace, and thermodynamic packages report exactly the standard axioms `propext`, `Classical.choice`, and `Quot.sound`. The factor-interaction bundle additionally inherits three `native_decide` axioms from the Example 4.3 regression it imports;

Table 2: Terminal Lean declarations and receipts. Hashes are SHA-256 of the proof source file; “deps” counts transitive local modules materialised and strictly compiled in the receipt.

Result	Declaration	deps	source hash (prefix)
Section 3 (MAF calculus)	MAFComposition. mafCompositionCalculusKernel and the nine public interfaces of proofs/MAFComposition/Main.lean	11	7fc801a6... d3c2
Section 4 (Thm. 4.7)	AutocatalyticCS.sourceDirectCSCoreEnum_ exact	15	0a54d083... bee8d2
Section 5 (Thms. 5.4–5.5)	DegradationControl.sourceNetwork_ arbitraryDegradationRegions, ...SpectralState, ...partialActuation_ stabilizable_iff_unactuatedCertificate, ...partialActuation_spectralStabilizable_ iff_unactuatedHurwitz, ...threeCycle_waterfill_global, ...arbitraryDegradationFiniteAtlas	15	f773c6e9... 54ed
Section 6 (Thms. 6.3–6.7)	CoreInteraction.coreInteractionCalculus; CoreInteraction.Degradation. SourceStarData.schurLoad_ spectralTrichotomy	35; 7	817b43c2... fee3; e3f82de9... 8449
Sections 7–8, 9 (package)	ContextualCoreInteraction. contextualSufficiency_boundedWidth_ coreInteraction	38	971c2657... 9661
Sections 7, 9 (post-proof closure)	ContextualCoreInteraction.publication_ sourceLossWitnesses; imports trace_eq_iff_boundaryIndistinguishable, summary_contextuallySufficient_iff_ injectiveOnRealizedTraces, computedMessage_decompositionInvariant, simultaneous_exactLocalReplacement_ contextual, TypeIIPassiveTail.literal_ passive_trace_eq	45	c64ff546... 81f9

these certify finite computations by the compiled evaluator rather than by the kernel and are confined to that regression. The terminal core-extraction theorem does not import that file.

The finite-dimensional Schauder fixed point used in Theorem 5.2 is a current-toolchain port of a public Lean development (a cube-parity Poincaré–Miranda argument followed by a simplex retraction and a compactness step); the port changed only compatibility details and the exported statement is the standard one. A broader Scarf–Brouwer development was tried first and abandoned because the port exposed roughly twenty API breaks without reducing the mathematical barrier. The gist from which the artifact was taken displayed no explicit license at retrieval, which does not affect the logical result but must be resolved before redistribution of the copied file.

Two junction-tree theorems coexist in the sources: a hypothesis-free equivalence with the dependent witness-tree semantics, and Theorem 8.5. Only the latter enters the public package. Likewise two checker theorems coexist and only the sound one is exported. The width-one and width-two regressions are single-leaf trees with trivial relations; they check the arithmetic, not a substantive instance.

10.3 Experiments and their role

Every numerical experiment in this work was a falsifier for a specific implication and was retired once it stopped discriminating hypotheses: the two-species MAF searches over coefficients in $\{0, 1, 2, 4\}$; the nested-matching and source-enumerator regressions of Section 4; the Type-I sweep, certificate adversaries, partial-actuation rays, Type-II₂ branch selection over 200 random log-uniform parameter sets, and 500 interval-corner samples of Section 5; the exact-rational degradation-star preflight; and a three-element exhaustive check of Theorems 7.3 and 7.5 (all 8 traces against all 8 contexts; a label summary is sufficient and a cardinality summary is not). The preflights selected proof routes; none is treated as evidence for a universal theorem.

11 Related work and novelty boundary

Amplification. The MAF is a von Neumann max–min model [12, 13] transported to reaction networks [10]; its ancestry in production/price duality and theorems of the alternative [14, 15] is classical and we claim none of it. Optimisation of growth bounds over subnetworks and the polyhedral geometry of autocatalytic ecosystems [24] are adjacent; recognising autocatalysis and maximising output are NP-complete in general [25]. The defensible contribution is the exact operation calculus on price profiles and the two impossibility boundaries.

Cores. Stoichiometric autocatalytic cores [5], their explicit-catalysis refinement [6], the bridge to RAF sets [7], and unstable cores as sources of instability [26] supply the semantics. The exchange-path theorem, the support-versus-edge correction, and the exact executable enumerator are new to our knowledge; the enumeration question itself is posed in [6].

Degradation. Perron–Frobenius theory for Metzler matrices, Collatz–Wielandt characterisations, M-matrix equivalences, and projective diagonal control [17–19] are baseline. The diluted onset model and small-degradation results are due to [8]; the classification and the open question to [9]. Our contribution is the source-faithful composition from a reaction list to the spectral partition for every degradation vector, the partial-actuation theorem, and the branch diagnostics.

Composition and interfaces. Exact local computation on decomposable graphical structures [27], bounded-width constraint satisfaction [28], acyclic database schemes [29], and factor-graph message passing [30] establish that tree-structured factorization supports exact local elimination. Behavioural systems [31] and categorical open reaction networks [32] establish relational black-boxing and composition by interconnection. RAF theory [3, 4] supplies the structural language. We claim priority for none of these individually. The novelty claim is the conjunction: a source-faithful factorization adapter; exact all-context trace semantics; the realized-image criterion; the isolated locality condition that repairs existential witness gluing; certified bounded separator-message arity; source-level amplification, thermodynamic and Type-II instances; and a warning-free formal closure. The contribution is not a new name for belief propagation: the semantic carrier is a property-specific feasibility relation over literal source data, and the theorems establish both its maximal contextual meaning and certified cases where tempting coarsenings fail.

12 Conclusion and open directions

Exact interfaces unify the positive and negative sides of modular proof. For the MAF, the interface is the price profile and the scalar is a certified loss. For catalyst-aware cores, the

interface between an ordinary core and its extensions is a pack of exchange paths. For degradation, the interface is a positive composition, and for a star of branches it is a scalar Schur load that adds. In general, equality of exact boundary traces licenses replacement in every environment, inequality automatically supplies a separating context, and certified tree decompositions compute the same relation while exposing only bounded-arity separators.

Several directions are open and none is required by the results above. The interaction layer covers forests and one-level stars; simple-cycle gain formulas, cactus-shaped incidence graphs, and degradation trees of greater depth are the natural next targets, as is a proof-carrying rational-linear message compiler for Theorem 8.7. For core extraction, output-sensitive or parameterised delay bounds, canonical anchoring to avoid duplicate work, and an end-to-end benchmark against a faithful implementation of the published complete enumeration would sharpen the practical claim. For degradation, the reducible Metzler equivalence “Hurwitz iff strict positive supersolution” would remove the irreducibility scope from the spectral partial-actuation corollary, and nonlinear continuation from the critical surface is a separate campaign. Finally, the MAF calculus, the degradation regions, and the trace theory together suggest a two-stage certified analysis of any proposed network edit: does the price profile prove a structural increase, and does the degradation region enlarge?

Data and code availability

The Lean sources, verification receipts, and the deterministic experiment scripts are part of the authors’ repository; the terminal declaration names and source hashes in Table 2 identify the verified artifacts. A public release tag will accompany submission.

References

- [1] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi:10.1016/S0022-5193(86)80047-9.
- [2] Wim Hordijk and Mike Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004. doi:10.1016/j.jtbi.2003.11.020.
- [3] Wim Hordijk, Peter R. Wills, and Mike Steel. Autocatalytic sets and biological specificity. *Bulletin of Mathematical Biology*, 76:201–224, 2014. doi:10.1007/s11538-013-9916-4.
- [4] Mike Steel, Wim Hordijk, and Joana C. Xavier. Autocatalytic networks in biology: structural theory and algorithms. *Journal of the Royal Society Interface*, 16:20180808, 2019. doi:10.1098/rsif.2018.0808.
- [5] Alex Blokhuis, David Lacoste, and Philippe Nghe. Universal motifs and the diversity of autocatalytic systems. *Proceedings of the National Academy of Sciences*, 117(41):25230–25236, 2020. doi:10.1073/pnas.2013527117.
- [6] Richard Golnik, Thomas Gatter, Peter F. Stadler, and Nicola Vassena. Autocatalytic cores in reaction networks with explicit catalysis. arXiv preprint arXiv:2603.02770 [math.CO], 2026. doi:10.48550/arXiv.2603.02770.
- [7] Richard Golnik, Thomas Gatter, Wim Hordijk, Peter F. Stadler, and Nicola Vassena. Bridging two theoretical frameworks of autocatalysis: RAF sets and stoichiometric autocatalysis. arXiv preprint arXiv:2605.25523, 2026. doi:10.48550/arXiv.2605.25523.

- [8] Jérémie Unterberger and Philippe Nghe. Stoechiometric and dynamical autocatalysis for diluted chemical reaction networks. *Journal of Mathematical Biology*, 85:26, 2022. doi:10.1007/s00285-022-01791-x.
- [9] Praneet Nandan, Philippe Nghe, and Jérémie Unterberger. Autocatalytic cores in the diluted regime: classification and properties. *Journal of Mathematical Biology*, 92:36, 2026. doi:10.1007/s00285-026-02357-7; preprint arXiv:2507.15546.
- [10] Praful Gagrani, Jiwei Wang, Yannick De Decker, and David Lacoste. Topological bounds on the dynamical growth rate of chemical reaction networks. arXiv preprint arXiv:2603.02627, 2026. doi:10.48550/arXiv.2603.02627.
- [11] Thomas Kosc, Denis Kuperberg, Etienne Rajon, and Sylvain Charlat. Thermodynamic consistency of autocatalytic cycles. *Proceedings of the National Academy of Sciences*, 122(18):e2421274122, 2025. doi:10.1073/pnas.2421274122.
- [12] John von Neumann. A model of general economic equilibrium. *The Review of Economic Studies*, 13(1):1–9, 1945. doi:10.2307/2296111.
- [13] John G. Kemeny, Oskar Morgenstern, and Gerald L. Thompson. A generalization of the von Neumann model of an expanding economy. *Econometrica*, 24(2):115–135, 1956. doi:10.2307/1905746.
- [14] Paul Gordan. Ueber die Auflösung linearer Gleichungen mit reellen Coefficienten. *Mathematische Annalen*, 6:23–28, 1873. doi:10.1007/BF01442864.
- [15] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley, Chichester, 1986.
- [16] Juliusz Schauder. Der Fixpunktsatz in Funktionalräumen. *Studia Mathematica*, 2:171–180, 1930. <https://eudml.org/doc/217247>.
- [17] Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*, volume 9 of *Classics in Applied Mathematics*. SIAM, 1994. doi:10.1137/1.9781611971262.
- [18] Lorenzo Farina and Sergio Rinaldi. *Positive Linear Systems: Theory and Applications*. Wiley, 2000. doi:10.1002/9781118033029.
- [19] Ron Ofir, Francesco Bullo, and Michael Margaliot. Minimum effort decentralized control design for contracting network systems. *IEEE Control Systems Letters*, 6:2731–2736, 2022. doi:10.1109/LCSYS.2022.3176196; preprint arXiv:2203.10392.
- [20] Thermodynamic compatibility of interacting autocatalytic cores: an exact phase diagram and a layered obstruction theory. Companion manuscript, same authors, 2026.
- [21] Unistationarity of Type II_ℓ autocatalytic cores with positive degradation, and exact multistationarity beyond minimality. Companion manuscript, same authors, 2026.
- [22] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi:10.1007/978-3-030-79876-5_37.
- [23] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020. doi:10.1145/3372885.3373824.

- [24] Praful Gagrani, Victor Blanco, Eric Smith, and David Baum. Polyhedral geometry and combinatorics of an autocatalytic ecosystem. *Journal of Mathematical Chemistry*, 62: 1012–1078, 2024. doi:10.1007/s10910-024-01576-x.
- [25] Jakob L. Andersen, Christoph Flamm, Daniel Merkle, and Peter F. Stadler. Maximizing output and recognizing autocatalysis in chemical reaction networks is NP-complete. *Journal of Systems Chemistry*, 3:1, 2012. doi:10.1186/1759-2208-3-1.
- [26] Nicola Vassena and Peter F. Stadler. Unstable cores are the source of instability in chemical reaction networks. *Proceedings of the Royal Society A*, 480:20230694, 2024. doi:10.1098/rspa.2023.0694.
- [27] Steffen L. Lauritzen and David J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society: Series B*, 50(2):157–194, 1988. doi:10.1111/j.2517-6161.1988.tb01721.x.
- [28] Eugene C. Freuder. Complexity of K-tree structured constraint satisfaction problems. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 4–9, 1990. <https://www.aaai.org/Papers/AAAI/1990/AAAI90-001.pdf>.
- [29] Catriel Beeri, Ronald Fagin, David Maier, and Mihalik Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30(3):479–513, 1983. doi:10.1145/2402.322389.
- [30] Frank R. Kschischang, Brendan J. Frey, and Hans-Andrea Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, 47(2):498–519, 2001. doi:10.1109/18.910572.
- [31] Jan C. Willems. The behavioral approach to open and interconnected systems. *IEEE Control Systems Magazine*, 27(6):46–99, 2007. doi:10.1109/MCS.2007.906923.
- [32] John C. Baez and Blake S. Pollard. A compositional framework for reaction networks. *Reviews in Mathematical Physics*, 29(9):1750028, 2017. doi:10.1142/S0129055X17500283.