

Computation-Realization Geometry

Semantic Interface Laws, Decision-Sufficient Abstraction, and Cost Duality

Fitih M. Cinnor for semiAIfoundry Research

Research Lead and Accountable Author

semiAIfoundry Research, a research group of Semi AI Foundry LLC

Accountable correspondence: fcinnor@semiaifoundry.com

22 August 2026

Version 3.2.0

DOI: [10.5281/zenodo.22050676](https://doi.org/10.5281/zenodo.22050676) Version: 3.2.0

Executable companion: <https://github.com/semiaifoundry/computation-realization-geometry>

Collective scholarly attribution and research provenance. This work was produced by semiAIfoundry Research through a human-directed collaboration across the intelligence continuum—human, artificial, and hybrid. AI models and multi-agent systems made material contributions to mathematical exploration, formal derivation, computational implementation, verification, literature analysis, adversarial critique, and manuscript development. Fitih M. Cinnor led the research program and serves as the accountable human author, retaining responsibility for the scientific claims and corrections.

Abstract

This paper introduces the foundational object of Computation-Realization Geometry (CRG). A terminalized computation and resource contract define a legal realization family $\mathcal{P}$; typed signatures yield the exact attainable region $\mathcal{R}$, the upper obligation region Γ , and the additive-price-visible body $\mathcal{K}$. Contextual equivalence gives canonical directed interface states in the declared tree model, while circuit-protocol translations and interactive examples show why realization identity must remain distinct from a resource image. Monotone reconstruction and nonnegative support duality identify Γ and $\mathcal{K}$ as exact decision-sufficient quotients for their respective objective classes. The Safe-Commitment Law characterizes, in both directions, when a retained family preserves every declared future decision and supplies an explicit regret witness when it does not. A strict engineering-cost separation corollary strengthens the necessity direction to continuous, coordinatewise strictly increasing, coercive costs. The paper also develops hierarchy coherence, resource monotonicity, products, local data processing, and communication synergy.

*Companion to **A Communication Theory of Computation Realization** (v3.2.0), DOI [10.5281/zenodo.22048090](https://doi.org/10.5281/zenodo.22048090). This manuscript expands the semantic, geometric, and decision-sufficiency foundations of the canonical theory.*

Contents

1	Introduction	3
1.1	Contributions	3
2	Relation to prior work	3
I	Semantic contracts and exact resource geometry	4
3	Terminalized computation and resource contracts	4
3.1	How a resource contract changes the realization family	5
4	Canonical distributed interface states	6

5	General causal realizations and exact certificates	8
5.1	Finite causal protocols and circuit translations	8
5.2	Where interaction changes the signature	9
5.3	Complete interactive finite-field linear orthant	10
5.4	Carrier graphs and network computation	11
5.5	Temporal cut-state conservation and scoped projections	12
6	The CRG abstraction ladder and decision-sufficient geometry	14
6.1	Why the geometry has four levels	14
6.2	Invariance and coherent hierarchy maps	15
6.3	Monotone reconstruction and linear duality	16
6.4	Safe commitment and insufficiency of scalar summaries	17
7	Monotonicity and composition	20
8	Analytic component certificates	21
8.1	Dense matrix multiplication across memory levels	21
8.2	A certified 2.5D technology map	22
8.3	Exact attention: recomputation versus storage	23
8.4	Cross-stage state and layout laws	24
9	Worked realization geometries	25
9.1	Broadcast and parity	25
9.2	INDEX and price reversal	26
9.3	Why $\mathcal{R}$, Γ , and $\mathcal{K}$ differ	26
9.4	Technology as an observation of the geometry	26
10	Conclusion	26
	Collective attribution, accountability, and AI research provenance	30

1 Introduction

An abstract computation specifies a relation among inputs, outputs, and correctness conditions. A physical system must additionally choose an algorithm, protocol, circuit, schedule, state allocation, ownership arrangement, hierarchy, resource allocation, and technology embedding. Those choices determine which distinctions are retained locally, which are recomputed, and which cross interfaces. The central scientific object is therefore not one implementation and not one communication number. It is the family of legal realizations and the geometry of the resource obligations they jointly attain.

This paper develops the semantic and geometric foundations of *computation-realization geometry* (CRG). For a terminalized computation F , resource contract ρ , and hierarchy $\mathcal{T}$, let $\mathcal{P}_F^\rho(\mathcal{T})$ be the legal realizations and let $\sigma_{\mathcal{T}}$ map one realization to its typed obligations. The exact image $\mathcal{R}$ retains joint attainability, the upper region Γ is the complete representation for continuous monotone resource objectives, and the convex upper body $\mathcal{K}$ is the complete representation for additive nonnegative prices:

$$(F, \rho, \mathcal{T}) \implies \mathcal{P} \xrightarrow{\sigma} \mathcal{R} \longrightarrow \Gamma \longrightarrow \mathcal{K}.$$

The successive arrows form an abstraction ladder. Each step forgets information, and each is exact for a precisely characterized class of future decisions.

The semantic foundation is contextual distinguishability. On a tree, assignments on one side of an interface are equivalent when every assignment on the opposite side induces the same required opposite-side outputs. The quotient gives a canonical state. General causal protocols enlarge the legal family to interactive message events; circuits and protocols are translated under explicit accounting. Finite-field linear computation then supplies a nontrivial exact interactive fragment in which directional cut ranks are both simultaneous lower bounds and jointly attainable loads.

The observation foundation is decision sufficiency. Monotone queries reconstruct Γ ; nonnegative linear prices reconstruct $\mathcal{K}$. These reconstruction theorems yield the Safe-Commitment Law: a retained family is lossless for the declared decision class exactly when it preserves the corresponding geometry. The result identifies the information boundary between reversible abstraction and irreversible commitment.

1.1 Contributions

- F1.** A typed core resource contract with coherent hierarchy pushforward and an explicit quantifier order from computation to technology.
- F2.** Canonical directed contextual states and explicit circuit/protocol translations under declared semantic and fixed-bit accounting conventions.
- F3.** An exact interactive directional cut-rank orthant for finite-field linear computation, together with carrier/network-code equivalence and temporal cut-state conservation.
- F4.** Exact reconstruction of Γ by continuous monotone observations and of $\mathcal{K}$ by nonnegative linear prices.
- F5.** The Safe-Commitment Law, constructive commitment-regret separators, and strict insufficiency of marginal minima, scalar summaries, and universal convexification.
- F6.** Hierarchy coherence, resource monotonicity, local data processing, product additivity, and constructive communication synergy.
- F7.** Worked realization geometries and structured matrix, attention, state-sharing, and layout-conversion slices that connect the formal objects to computational design.

2 Relation to prior work

The foundations developed here connect several theories whose starting objects are usually different. Turing computation abstracts from hardware realization, while Gandy and Deutsch ask which principles connect abstract computation to physically admissible mechanisms [26, 27, 28]. Classical realization theory constructs internal state-space models from input–output behavior [29, 30]; automata theory constructs canonical states by contextual indistinguishability [31]. CRG extends the realization question from state dimension to a family of legal algorithms, protocols, circuits, schedules, mappings, ownership structures, hierarchies, and resource allocations.

Shannon theory characterizes codes and limits for a declared source/channel model [32]. Communication complexity characterizes the exchange required by distributed parties after the function, ownership split, and protocol class are fixed [33, 19]; information complexity refines the cost of interaction [3, 4]. Network coding

and network function computation show that legal local transformations change attainable load or rate regions [22, 34, 21]. These theories provide converses and constructive regions within declared operational models. CRG registers those models as realization classes and places their traffic coordinates beside state, work, causal depth, replication, and other declared obligations.

Thompson’s VLSI area–time theory and Hong–Kung I/O complexity convert communication service into lower bounds under physical or memory contracts [25, 5]. Cache-oblivious and distributed numerical linear-algebra results provide simultaneous multilevel constructions and communication-optimal schedules [7, 6, 8]. The component slices in this paper use such theorems as outer obligations and constructive witnesses inside the same resource geometry.

Contract-based design and interface theories make assumptions, guarantees, compatibility, and refinement explicit [35, 36]. Abstract interpretation studies sound abstraction through ordered semantic domains [37]. Blackwell’s comparison of experiments asks whether one information structure is at least as valuable for every decision problem in a class [38]. CRG’s decision-sufficiency results stand closest to this latter tradition: the upper obligation region and convex price body are characterized by the complete classes of decisions they answer, and the Safe-Commitment Law is equality of those decision-sufficient objects before and after pruning.

Vector optimization supplies upper images, Pareto order, support duality, and outer approximation [2, 39, 40]. The CRG regions use this mathematics, but retain a semantic realization fiber behind the image. This fiber is essential when two algorithms or schedules have the same declared signature but differ under a later embedding-sensitive objective. The novelty of the foundations is therefore the integrated computation-first object: legal realization semantics, typed joint obligations, coherent hierarchy maps, exact observation quotients, and a necessary-and-sufficient law of safe reduction.

Part I

Semantic contracts and exact resource geometry

3 Terminalized computation and resource contracts

The mathematical object to be realized is separated from the resources by which it is realized. This separation fixes the quantifiers used throughout: first a semantic task, contract, and hierarchy are fixed; only then are legal realizations and their interface signatures compared.

Definition 3.1 (Terminalized deterministic computation). *A terminalized deterministic computation is a map*

$$F : \mathcal{X} = \prod_{i \in I} X_i \longrightarrow \mathcal{Y} = \prod_{j \in O} Y_j,$$

where every alphabet is finite. A hierarchy consists of a finite tree $\mathcal{T} = (V, E)$ and placement maps $\lambda_I : I \rightarrow V$ and $\lambda_O : O \rightarrow V$. Vertex v owns the local input alphabet

$$X_v = \prod_{i: \lambda_I(i)=v} X_i$$

and must produce the local output alphabet

$$Y_v = \prod_{j: \lambda_O(j)=v} Y_j.$$

Definition 3.2 (Core resource contract). *A core deterministic resource contract is a tuple*

$$\rho = (\mathcal{B}, S, D, M, R, \varepsilon, \Lambda, \chi),$$

where $\mathcal{B}$ is the allowed local primitive basis, S is a local-work or circuit-size budget, D a causal-depth budget, M a storage budget, R a replication or recomputation budget, ε an allowed error (zero in the exact results below), Λ the admissible hierarchy and terminal-placement rules, and χ the communication-accounting type. A ρ -legal realization P is a finite hierarchical protocol or circuit satisfying these typed bounds. Contract relaxation is semantic inclusion: $\rho \preceq \rho'$ exactly when every ρ -legal realization is ρ' -legal.

To make hierarchy contraction formal, each scalar bound is represented by a finitely supported charge on hierarchy sites or interfaces together with a declared associative aggregation law. Local work and storage use

addition; retained interface traffic and replication charges use addition over surviving interfaces; causal depth and protocol rounds use max-plus path composition; terminal placement uses direct image; exact error remains zero, while an approximate extension must declare an associative error-composition operation $\oplus_\varepsilon$; and χ is preserved. A connected-fiber contraction $q : \mathcal{T} \rightarrow \mathcal{T}'$ pushes a charge forward by aggregating all charges in each fiber and on each induced surviving interface.

A resource budget is a legality constraint on the complete realization; a resource charge is the additive or max-plus object pushed through a hierarchy map. A budget need not equal the sum of local charges unless the contract declares that aggregation rule. Thus, for example, a package-wide peak or shared memory pool is not counted once per stage merely because stage-local timing expressions are written separately. Extensions may add stochastic, analog, reversible, or technology-specific coordinates, but they must specify legal-realization semantics, an associative aggregation law, and pushforward behavior before the coherence results apply.

3.1 How a resource contract changes the realization family

The contract is part of the mathematical problem, not metadata attached after optimization. Three simple changes illustrate its effect.

Time sharing.

Two stage-local state requirements m_1, m_2 aggregate as $m_1 + m_2$ when dedicated engines coexist and as $\max\{m_1, m_2\}$ when one state pool is reused across disjoint time intervals. The latter realization also carries a schedule and utilization coordinate. Replacing one aggregator by the other changes legality.

Replication and recomputation.

A replicated table or recomputed intermediate can reduce communication while increasing work, state, or replication charges. Those are different legal points in one typed geometry, not different estimates of one fixed graph.

Approximation.

Permitting error $\varepsilon > 0$ enlarges the legal family only after the correctness relation and the composition law $\oplus_\varepsilon$ have been declared. Approximate and exact geometries are different slices of the theory.

The hierarchy rule Λ is equally active. A partition can be a nested accounting cut, a physical carrier, a clock or ownership boundary, or a measurement hierarchy. The contract specifies which interpretation is in force and which resource charges survive contraction. This is the mechanism that prevents a static cut count from being silently read as bandwidth, latency, routed congestion, or energy.

Proposition 3.3 (Functoriality of contract pushforward). *For every core contract with the aggregation laws of Theorem 3.2, connected-fiber contractions satisfy*

$$\text{id}_* \rho = \rho, \quad (r \circ q)_* \rho = r_*(q_* \rho).$$

Proof. For an additive local charge a , the charge assigned to a coarse site w by the composite contraction is

$$((r \circ q)_* a)(w) = \sum_{v: (r \circ q)(v)=w} a(v) = \sum_{u: r(u)=w} \sum_{v: q(v)=u} a(v) = (r_*(q_* a))(w).$$

The same regrouping proves the claim for additive interface and replication charges. Max-plus path aggregation is associative because path concatenation and max are associative; direct image of terminal placement is functorial; exact error is fixed at zero; and every admitted approximate coordinate is functorial by the required associativity of $\oplus_\varepsilon$. The primitive basis and accounting type are unchanged. Identity maps leave every support and charge unchanged, so componentwise equality gives both identities. $\square$

Definition 3.4 (Typed interface signature). *For an oriented hierarchy interface e , a typed signature is*

$$\mu_e = (s_e, \nu_e, t_e, d_e, a_e, \pi_e),$$

where s_e is bits per state or message, ν_e is message events per operation, $t_e = s_e \nu_e$ is bits per operation, d_e is a causal-round or depth contribution, a_e is an activity factor, and π_e is an optional traffic or burst distribution. An analysis may project to any declared subset. Static cut accounting uses s_e or a fixed-bit width; throughput and link calculations use t_e ; latency uses d_e ; and dynamic energy requires t_e together with a_e and a technology model. These coordinates are not interchangeable.

Definition 3.5 (Interface-accounting convention). *For a finite message alphabet Σ , define its semantic log-cardinality and fixed-length binary width by*

$$\ell_{\text{sem}}(\Sigma) = \log_2 |\Sigma|, \quad \ell_{\text{bit}}(\Sigma) = \lceil \log_2 |\Sigma| \rceil.$$

For a sequence of message events, capacities add event by event. The contract records $\chi \in \{\text{sem}, \text{bit}\}$, or an explicitly specified richer coding rule. Semantic capacity is exact for state-space and product identities; fixed-bit traffic is the physical convention for Boolean circuit nets. The two conventions coincide at one bit for a binary alphabet.

Unless stated otherwise, the event schedule is fixed independently of the input, conditional absence is charged as an alphabet symbol, and opposite directions are charged separately. The event sum is provisioned alphabet capacity, not Shannon entropy of a reachable-transcript distribution. Joint coding across events, expected variable length, compression, retransmission, broadcast cost, and link sharing require a richer χ stating decoder, timing, and charging rules. Bits per operation alone do not specify instantaneous bandwidth or latency; event-rate, burst, and service coordinates are also required.

4 Canonical distributed interface states

Removing an undirected edge $\{u, v\}$ creates two components. Write $V_{u|v}$ for the component containing u , $X_{u|v}$ for the product of inputs placed there, and $Y_{v|u}$ for the outputs placed on the opposite side.

Definition 4.1 (Separable directed-summary realization). *For every oriented edge $u \rightarrow v$, choose a finite message alphabet $\Sigma_{u \rightarrow v}$. A separable directed-summary realization consists of local maps*

$$\delta_{u \rightarrow v} : X_u \times \prod_{w \in N(u) \setminus \{v\}} \Sigma_{w \rightarrow u} \longrightarrow \Sigma_{u \rightarrow v}$$

and output maps

$$\eta_u : X_u \times \prod_{w \in N(u)} \Sigma_{w \rightarrow u} \longrightarrow Y_u.$$

Because the graph is a tree, the message on $u \rightarrow v$ is recursively a function only of $X_{u|v}$. Oppositely directed messages summarize their own sides and do not depend on one another. This is a noninteractive two-pass summary model, not the general interactive model introduced in Section 5.

Definition 4.2 (Directed contextual equivalence). *For $a, a' \in X_{u|v}$, define*

$$a \equiv_{u \rightarrow v} a' \iff F_{Y_{v|u}}(a, z) = F_{Y_{v|u}}(a', z) \text{ for every } z \in X_{v|u}.$$

Let

$$Q_{u \rightarrow v} = X_{u|v} / \equiv_{u \rightarrow v}, \quad \kappa_{u \rightarrow v} = \log_2 |Q_{u \rightarrow v}|.$$

The real logarithm is the semantic capacity; a fixed-length binary encoding uses $\lceil \kappa_{u \rightarrow v} \rceil$ bits.

Theorem 4.3 (Canonical directed-summary theorem). *For every finite terminalized computation F and every tree hierarchy $\mathcal{T}$:*

- (a) *the directed contextual classes are locally composable;*
- (b) *the class messages $Q_{u \rightarrow v}$ yield a separable directed-summary realization of all outputs;*
- (c) *every separable directed-summary realization satisfies $|\Sigma_{u \rightarrow v}| \geq |Q_{u \rightarrow v}|$ on every oriented edge; and*
- (d) *if equality holds on every oriented edge, the realization is the canonical quotient realization on every reachable local input-message tuple, up to independent renaming of message states; behavior on unreachable tuples is unconstrained.*

Thus

$$\kappa_{\mathcal{T}}^{\leftrightarrow}(F) = (\kappa_{u \rightarrow v})_{u \rightarrow v}$$

is an implementation-independent, encoding-optimal, jointly realizable directed-capacity vector for the hierarchy.

Proof. Fix $u \rightarrow v$. The component $V_{u|v}$ is the disjoint union of u and the components $V_{w|u}$ for $w \in N(u) \setminus \{v\}$. Suppose $a_w, a'_w \in X_{w|u}$ satisfy $a_w \equiv_{w \rightarrow u} a'_w$ for every such neighbor, and fix the same local assignment at u . For any assignment z on $V_{v|u}$, replace the a_w by a'_w one component at a time. By definition, each replacement preserves every output outside $V_{v|u}$ for every assignment to the remaining variables, hence in particular every output in $V_{v|u}$.

The combined assignments are therefore equivalent under $\equiv_{u \rightarrow v}$. The outgoing class is consequently a well-defined function of the local input and incoming classes.

For an output at vertex u , choose arbitrary representatives of the incoming classes. Replacing any representative by another in the same class preserves every output outside that neighbor component, including the outputs at u . The local output is therefore well defined and equals F .

For minimality, consider any separable realization. Its message $m_{u \rightarrow v}$ is a deterministic function of the assignment on $V_{u|v}$. If inequivalent assignments a, a' produced the same message, some opposite-side assignment z would distinguish an output in $V_{v|u}$. Every opposite-side state would be identical in the two runs, and the only dependence on the u -side would be through the identical message. Determinism would force identical outputs, a contradiction. Thus different contextual classes require different messages.

If every alphabet has minimum cardinality, its message map can neither merge two classes nor use spare states to split one class. It induces a bijection $Q_{u \rightarrow v} \rightarrow \Sigma_{u \rightarrow v}$, and the local maps are conjugate to the quotient maps under these independent bijections on every reachable local input–message tuple. Their values on unreachable tuples may be assigned arbitrarily without changing correctness, minimum message cardinality, or canonical equivalence on the reachable domain. $\square$

Interpretation. The theorem is simultaneously constructive, minimal, and joint. It does not obtain one lower bound per edge and then splice incompatible edgewise optima. The same quotient realization attains every directed minimum at once. The message state on $u \rightarrow v$ is the complete set of distinctions on the u -side that can still matter to an output on the v -side under some legal opposite-side context. Any smaller alphabet merges two distinguishable contexts; any larger alphabet contains representation redundancy.

The result also explains why terminal placement is load bearing. Moving an output to another vertex changes which contexts can distinguish two assignments and therefore changes the quotient itself. A “communication requirement of the function” is consequently incomplete until input and output ownership have been declared.

Example 4.4 (Three-node worked quotient). *Let a path $u - v - w$ carry input bits x at u and y at w , with output $x \oplus y$ required at v . Across $u \rightarrow v$, the two values of x are contextually inequivalent because fixing $y = 0$ distinguishes the required output; hence $Q_{u \rightarrow v} \cong \{0, 1\}$. Symmetrically, $Q_{w \rightarrow v} \cong \{0, 1\}$. In the reverse directions no output lies beyond the edge, so the quotient has one state. The canonical realization sends one bit inward from each leaf and computes parity locally at v . The directed capacity vector follows from semantics alone and is attained without choosing a binary encoding convention in advance.*

Corollary 4.5 (Fixed-bit canonical widths). *Under fixed-length binary accounting,*

$$\hat{\kappa}_{u \rightarrow v} = \lceil \kappa_{u \rightarrow v} \rceil$$

is jointly realizable and componentwise minimal among all fixed-length encodings in the separable directed-summary model. Each ceiling adds less than one bit beyond semantic capacity.

Corollary 4.6 (Single-root quotient). *If all outputs are placed at a root r , every edge directed toward r has the contextual-equivalence state space of its subtree, while every reverse edge has one state. The canonical realization is bottom-up and is analogous to a Myhill–Nerode quotient for the selected variable tree.*

Corollary 4.7 (Independent-bit routing). *Let source bit x_i be placed at one vertex and copied to an arbitrary set of output terminals. For oriented edge $u \rightarrow v$, let $k_{u \rightarrow v}$ be the number of source bits located in $V_{u|v}$ with at least one requested copy in $V_{v|u}$. Then*

$$|Q_{u \rightarrow v}| = 2^{k_{u \rightarrow v}}, \quad \kappa_{u \rightarrow v} = k_{u \rightarrow v}.$$

Consequently, total directed semantic capacity across an undirected edge equals the number of independent routed nets separated by that cut.

Proof. The opposite-side outputs depend on exactly those $k_{u \rightarrow v}$ source bits. Two assignments on the u -side are contextually equivalent precisely when they agree on all of them, giving $2^{k_{u \rightarrow v}}$ classes. $\square$

Example 4.8 (Broadcast and all-reduce). *A single source bit broadcast to many consumers has one bit of semantic capacity on every directed edge separating the source from at least one consumer and zero in the reverse direction. One bit of information can thus accumulate arbitrarily large multiscale incidence. By contrast, if every node must output the parity of all source bits, each side of every nontrivial edge has two summary states (its local parity), so one bit is required in each direction. Semantic information and physical distribution incidence are therefore distinct quantities.*

5 General causal realizations and exact certificates

The quotient theorem is exact within the separable two-pass model. A general circuit or protocol may cross the same interface in several causal stages, so the realization class must be enlarged before making interactive claims. We use four explicit conventions: circuits are finite, combinational, and acyclic; each Boolean signal is one driver-to-sinks hyperedge net; gates and terminals are placed at hierarchy vertices; and each hierarchy edge is either a nested cut-accounting boundary or a physical tree hop, as declared by ρ . Net incidence counts one bit per crossed tree edge, not one bit per sink pin. Sequential reuse, registers, variable schedules, routed congestion, and nonfree branching require additional time, storage, or carrier coordinates.

The section makes three successive enlargements of the realization model. First, a circuit is represented as a causal protocol so that static net incidence and message traffic can be compared under one accounting contract. Second, interaction is admitted, showing precisely how nonlinear computations can alter the lower bounds derived from one-way contextual states. Third, a temporal service window is introduced, allowing retained state and communication to enter one conservation law. The three steps separate representation, interaction, and time instead of folding them into one overloaded “communication cost.”

For each enlargement the proof pattern is the same:

semantic range or quotient $\implies$ outer obligation and explicit protocol or circuit $\implies$ inner witness.

When the two meet, the corresponding CRG slice is exact.

5.1 Finite causal protocols and circuit translations

Definition 5.1 (Finite acyclic tree protocol). *A finite acyclic protocol P on a tree hierarchy consists of a finite event DAG. Each message event a is labeled by an oriented tree edge, outputs a symbol in a finite alphabet Σ_a , and is a deterministic function of the local input at its source vertex and earlier messages received there. A conditional or absent message is an ordinary alphabet symbol, so the event schedule is static. Each output event is similarly a deterministic function of local input and earlier received messages. The protocol computes F when all terminal outputs equal $F(x)$ for every x .*

For tree edge e , define semantic alphabet traffic and fixed-length binary traffic by

$$a_e(P) = \sum_{a: a \text{ labeled by } e} \log_2 |\Sigma_a|, \quad c_e(P) = \sum_{a: a \text{ labeled by } e} \lceil \log_2 |\Sigma_a| \rceil.$$

Write $t_e^x(P) = a_e(P)$ under semantic accounting and $t_e^x(P) = c_e(P)$ under fixed-bit accounting. A round-resolved signature may retain every message event separately; $t^x(P)$ is its declared static projection.

Definition 5.2 (Circuit cut vector). *Let C be a finite Boolean circuit with one hyperedge net per signal, and let $\mathcal{T}$ place every gate and terminal at a tree vertex. For tree edge e , let M_e denote either side of its cut and define*

$$b_e(C, \mathcal{T}) = |\{n : \text{net } n \text{ has pins on both sides of } e\}|.$$

The vector $b(C, \mathcal{T})$ is the multiscale cut-incidence vector of the placed implementation.

The translation treats a gate or terminal at v as a local deterministic event, a one-bit net as one-bit messages on its minimal connecting subtree, and fanout as replication at a tree branch. Replication is free only in the cut-only contract; otherwise it contributes local work, storage, and possibly depth. Topological gate order becomes event-DAG order. A k -symbol message uses a chosen injective $\lceil \log_2 k \rceil$ -bit code whose encoder and decoder logic is included in the local event circuits. Reuse through time is either unrolled into distinct events or represented by stored sequential state and therefore cannot be inferred from a static cut vector.

Theorem 5.3 (Circuit-to-protocol factorization). *For every finite acyclic Boolean circuit C and tree placement $\mathcal{T}$ of its gates and terminals, there is a finite acyclic tree protocol computing the same terminalized function with*

$$c_e(P) = b_e(C, \mathcal{T})$$

for every tree edge e , under a communication-only contract in which forwarding, replication, and multicast at tree branch vertices are free. The equality is an undirected cut-total identity; the net driver determines the direction on every crossed edge. Moreover,

$$\sum_{e \in E(\mathcal{T})} b_e(C, \mathcal{T}) = \sum_{n \in E(C)} |E(\text{Steiner}_{\mathcal{T}}(\text{pins}(n)))|,$$

where the right side counts edges in the minimal subtree connecting the pins of each circuit net.

Proof. Process circuit nets in topological order of their drivers. Once the driver of net n is known, route its one-bit value along the minimal subtree of $\mathcal{T}$ connecting the driver and all sinks, replicating at branch vertices. A tree edge is used by n exactly when deleting that edge separates at least two pins of n , equivalently when n crosses the associated circuit cut. Exactly one bit for n traverses each such edge, proving the pointwise and summed identities. A net is routed only after its driver is computed, so the event schedule is acyclic. $\square$

Remark 5.4 (Resource cost of materialized forwarding). *If forwarding is materialized as logic and every physical tree hop consumes a causal stage, the construction adds $O(\sum_e b_e)$ forwarding events and at most a factor $\text{diam}(\mathcal{T})$ to path depth. If hierarchy edges are nested accounting boundaries, the theorem is a cut representation and adds no logical depth. The contract must choose the interpretation.*

Theorem 5.5 (Protocol-to-circuit compilation). *Let P be a finite acyclic tree protocol. Choose an injective fixed-length binary code for every message alphabet, and suppose each encoded message or output event, including required encoder and decoder logic, is supplied by a Boolean circuit over $\mathcal{B}$ with size s_a and depth d_a as a function of its locally available bits. Then P compiles into a Boolean circuit computing the same terminalized function with*

$$S(C) \leq \sum_a s_a,$$

and depth at most the maximum weighted length of a path in the event DAG, where event a has weight d_a . Every message bit becomes a circuit net crossing the corresponding tree interface, hence

$$b_e(C, \mathcal{T}) \leq c_e(P).$$

Combinational unrolling represents a message reused in several causal rounds by separate nets; a sequential realization may instead trade storage and time for wires under an enriched contract.

Proof. Topologically order the event DAG. Instantiate the stated circuit for each event and connect its input pins to local input terminals and the earlier message outputs on which it depends. Acyclicity gives a legal circuit. Summing local event sizes proves the size bound, and the longest weighted event path proves the depth bound. Each encoded message bit is one net crossing its labeled interface, which proves the cut bound. $\square$

What the two translations establish. The circuit-to-protocol direction preserves the exact multiscale incidence of one placed combinational realization under the cut-only contract. The protocol-to-circuit direction records the local logic and causal depth needed to implement arbitrary finite message alphabets. Together they show that “circuit” and “protocol” are two representations of realizations inside CRG, not competing definitions of the design space. Their signatures agree only after the contract specifies whether forwarding, fanout, storage, and time are free or charged.

As a simple example, a one-bit signal driven at one leaf and consumed at two other leaves uses each edge of the minimal connecting subtree once. Its summed circuit-cut incidence equals the number of edges in that subtree; the protocol witness routes the bit along exactly those edges and replicates only at their branch point. If branch replication is charged, the same semantic realization acquires an additional local-work or activity coordinate without changing the cut identity.

Corollary 5.6 (Representation equivalence with explicit overhead). *Circuits and finite acyclic hierarchical protocols are mutually translatable. The forward translation exposes multiscale cut incidence exactly under the communication-only contract; the reverse compilation preserves semantics with explicit local-computation and causal-depth overhead. Arbitrary interface encoding and interaction therefore belong inside the feasible realization family rather than being excluded by definition.*

5.2 Where interaction changes the signature

Example 5.7 (INDEX interaction boundary). *Alice holds $x \in \{0, 1\}^n$, Bob holds $i \in [n]$, and Bob must output x_i . A one-way Alice-to-Bob summary must distinguish all 2^n strings, so its semantic capacity is n bits. A two-round protocol lets Bob send i using $\lceil \log_2 n \rceil$ bits and Alice return x_i using one bit. Thus*

$$D_{\text{one-way}}^{A \rightarrow B}(\text{INDEX}_n) = n, \quad D_{\text{interactive}}(\text{INDEX}_n) \leq \lceil \log_2 n \rceil + 1.$$

The legal directional signatures $(n, 0)$ and $(1, \lceil \log_2 n \rceil)$ can be selected by different price vectors. With uniform bit prices and $n \geq 4$, the interactive realization wins; if Bob-to-Alice traffic is sufficiently expensive, the one-way realization wins. A one-way contextual summary, total communication, and a technology-weighted interface objective are different observations.

In the noninteractive model, contextual quotients are both constructive points and matching lower bounds. In a general model, a semantic outer bound requires a problem-specific converse from communication complexity, information complexity, algebra, or another appropriate method. The operational discipline is therefore to retain certified brackets

$$\mathcal{R}_{F,\text{in}}^\rho(\mathcal{T}) \subseteq \mathcal{R}_F^\rho(\mathcal{T}) \subseteq \mathcal{R}_{F,\text{out}}^\rho(\mathcal{T})$$

and declare every gap.

5.3 Complete interactive finite-field linear orthant

Theorem 5.8 (Linear contextual states are cut ranks). *Let $F(x) = Mx$ over $\text{GF}(2)$, split inputs as $x = (x_A, x_B)$, and place a set of outputs Y_B on side B . If $M_{B,A}$ is the submatrix mapping x_A to outputs required on B , then $a, a' \in \text{GF}(2)^A$ are contextually equivalent for $A \rightarrow B$ exactly when*

$$M_{B,A}(a - a') = 0.$$

The quotient has $2^{\text{rank } M_{B,A}}$ classes, and

$$\boxed{\kappa_{A \rightarrow B} = \text{rank}_{\text{GF}(2)} M_{B,A}.$$

Proof. For any outside assignment b , the outputs on B differ by $M_{B,A}(a - a')$. They agree for every b exactly when this vector is zero. The number of cosets of $\ker M_{B,A}$ is $2^{\text{rank } M_{B,A}}$. $\square$

Lemma 5.9 (Interactive directional range bound). *Delete one tree edge and call the sides A and B . In any exact deterministic finite causal protocol with a static charged event schedule,*

$$t_{A \rightarrow B} \geq \max_b \log_2 |\{F_B(a, b) : a \in X_A\}|$$

under semantic alphabet accounting; fixed binary width rounds the right side up.

Proof. Fix b and retain the ordered tuple of every A -to- B symbol. Given that tuple and the receiver input, every receiver-local event and reverse message is determined recursively by induction over the event DAG. Equal forward tuples must therefore give equal required receiver outputs. The tuple has at least the displayed number of values and at most the product of its event alphabet cardinalities. $\square$

Lemma 5.10 (Local composition of linear quotient states). *For a component C exposed by deleting a tree edge, let*

$$L_C = M_{O_{\bar{C}}, I_C}, \quad Q_C = \mathbb{F}_{\mathfrak{q}}^{I_C} / \ker L_C,$$

where $\mathfrak{q}$ is the field cardinality. If $O_B \subseteq O_{\bar{C}}$, then

$$\boxed{\ker M_{O_{\bar{C}}, I_C} \subseteq \ker M_{O_B, I_C}.$$

Consequently $M_{O_B, I_C} x_C$ is a well-defined linear function of Q_C . If disjoint incident components $C_1, \dots, C_s$ meet at v and O_B lies outside every C_i , their aggregate contribution to O_B is a linear function of x_{I_v} and the states Q_{C_i} .

Proof. Deleting output rows from $M_{O_{\bar{C}}, I_C}$ can only enlarge its kernel. Thus

$$L_C x_C \mapsto M_{O_B, I_C} x_C$$

is well defined on $\text{im } L_C \cong Q_C$. Apply this factor map to each incident component and add the resulting vectors together with the local term $M_{O_B, I_v} x_{I_v}$. $\square$

Definition 5.11 (Interactive linear traffic contract). *Let $F(x) = Mx$ over a fixed finite field $\mathbb{F}_{\mathfrak{q}}$, with input coordinates and required output rows placed on $\mathcal{T}$. For an oriented edge $e = A \rightarrow B$, put*

$$M_e = M_{O_B, I_A}, \quad r_e = \text{rank}_{\mathbb{F}_{\mathfrak{q}}} M_e.$$

The contract admits exact deterministic finite causal tree protocols, arbitrary deterministic local maps (not necessarily linear), and a static charged event schedule. Each vertex begins with only its placed input and fixed input-independent advice; no input-dependent state is prepositioned across a cut. Traffic is charged by [Theorem 3.5](#); local work, storage, rounds, and depth are unrestricted in this slice. Conditional absence remains an ordinary charged symbol.

Theorem 5.12 (Exact interactive linear lower corner and orthant). *Under Theorem 5.11, arbitrary interaction and nonlinear local encoding cannot reduce the directional load below $r_e \log_2 \mathfrak{q}$. One quotient realization attains all edge bounds simultaneously. Hence*

$$\Gamma_F^\rho(\mathcal{T}) = \mathcal{K}_F^\rho(\mathcal{T}) = \prod_{e \in \vec{E}(\mathcal{T})} [r_e \log_2 \mathfrak{q}, \infty)$$

under semantic traffic, with lower corner $(\lceil r_e \log_2 \mathfrak{q} \rceil)_e$ under fixed binary widths. For every nonnegative directional price vector w ,

$$\inf_P \langle w, t(P) \rangle = \sum_{e \in \vec{E}(\mathcal{T})} w_e r_e \log_2 \mathfrak{q}.$$

Proof. Fix the complete B -side input. The required B -side output ranges over an affine translate of $\text{im } M_{OB,IA}$, which has $\mathfrak{q}^{r_e}$ points. The directional range bound gives the converse without assuming linear encoders. Contextual equivalence identifies the sender quotient

$$\mathbb{F}_{\mathfrak{q}}^{IA} / \ker M_{OB,IA},$$

of dimension r_e .

For simultaneous achievability, root the tree arbitrarily. In the inward pass, suppose child components C_i meet their parent at v and $C = \{v\} \cup_i C_i$. Since $O_{\bar{C}} \subseteq O_{\bar{C}_i}$, Theorem 5.10 lets v compute Q_C from its local input and the child states. In the outward pass toward child C_j , the parent combines its local input, the inward states of all sibling components, and its incoming outward state. The target rows O_{C_j} lie outside every contributing component, so the same lemma computes the complementary state sent into C_j . Every vertex then combines its local input and incident component states to produce its placed outputs. Identifying each quotient with $\text{im } L_C$ and choosing a basis sends exactly r_e field symbols on every oriented edge in one realization. The attainable lower corner and componentwise converse give the upper orthant, already convex. $\square$

A worked three-site linear realization. Let a path 1–2–3 carry input bits x_1, x_2, x_3 at the three sites, and require outputs

$$y_1 = x_1 \oplus x_2, \quad y_2 = x_1 \oplus x_2 \oplus x_3, \quad y_3 = x_2 \oplus x_3$$

at sites 1, 2, 3, respectively. Across the left edge, the contribution of x_1 to outputs on the right has rank one, and the contribution of (x_2, x_3) to the left output also has rank one. The same is true across the right edge. The exact directional lower corner is therefore one bit in each direction on each edge. An inward/outward quotient protocol attains all four bounds simultaneously by exchanging the appropriate local parities. The example illustrates the content of the orthant theorem: interaction may change event order, but it cannot reduce the dimension of the cross-cut image, and one realization achieves all cut-rank obligations at once.

The orthant theorem determines the complete static charged-traffic geometry for this finite-field linear contract. Rounds, depth, local work, storage, approximation, randomization, variable-length expected cost, and graph-carrier coordinates define additional CRG projections and can be added without altering the exact traffic result. INDEX lies outside this linear class because its opposite input selects which distinction matters; for a linear map the cross-side image is context independent. On a graph with parallel paths, the rank converse applies to aggregate cut traffic, not independently to each arc. The one-cut transcript and rank facts are standard; the joint conclusion is the simultaneous, arbitrary-placement directed orthant even when reverse interaction is free and local encoders may be nonlinear [19, 20, 21].

5.4 Carrier graphs and network computation

A tree records nested accounting boundaries; a directed carrier graph records which local channels exist. Let $\mathcal{N} = (\mathcal{G}_{\text{car}}, \alpha_I, \alpha_O)$ be a finite directed acyclic multigraph with source-access and output-demand relations. A carrier realization assigns a finite alphabet to every arc, a local encoder from the tail's accessible sources and incoming symbols, and exact decoders at demand vertices. Its signature is the vector $\ell(P)$ of log alphabet sizes. Define the equal-instance normalized zero-error link-load region

$$\mathcal{L}_F(\mathcal{N}) = \text{cl up } \bigcup_{k \geq 1} \left\{ \frac{\ell(P)}{k} : P \text{ realizes } F^{\otimes k} \text{ exactly on } \mathcal{N} \right\}.$$

Block concatenation makes this closed upper region convex.

Theorem 5.13 (Carrier/network-code equivalence). *Under deterministic zero-error semantics on a finite acyclic carrier, single-shot carrier realizations and network function codes are identical up to internal-symbol renaming, with identical arc loads. Applying the bijection to $F^{\otimes k}$ for every k , normalizing by k , and taking closed upper block closure gives exactly the zero-error network-computation link-load region. Identity or multicast output demands specialize this to ordinary multi-source network codes under equal replication of the source instance fixed by F . This does not identify the full independent-source rate-vector region.*

Proof. Both objects assign precisely the same local encoder to each directed arc and the same decoder to each demand vertex. A topological order makes every carrier encoder a legal network-code encoder; the reverse translation changes nothing. The identity translation commutes with independent block products, normalization, and closure. $\square$

For arc capacities c_a , the asymptotic computation rate is

$$\sup (\{0\} \cup \{R > 0 : \exists \ell \in \mathcal{L}_F(\mathcal{N}), R\ell_a \leq c_a \forall a\}).$$

The adjoined zero fixes the gauge: if the carrier is zero-error unsolvable and $\mathcal{L}_F(\mathcal{N}) = \emptyset$, its computation capacity is zero. Placing carrier vertices in $\mathcal{T}$ linearly projects arc loads to tree-cut loads, but parallel arcs show the projection is not injective. Matched network-code solvability and nonlinear-code counterexamples therefore transfer exactly. Linear-versus-nonlinear comparisons require the explicit subcontract in which source and arc alphabets are vector spaces over one fixed field and all local maps are linear.

The inherited undecidability statement is specifically the variable-alphabet matched-solvability problem: given finite source, carrier, and demand incidence, decide whether some common alphabet size at least two and deterministic local encoders satisfy every demand exactly. It does not make fixed-alphabet feasibility or every network-code decision problem undecidable; a finite instance with every alphabet supplied is decidable by finite enumeration [22, 23, 24].

5.5 Temporal cut-state conservation and scoped projections

Definition 5.14 (Schedule-reachable temporal cut-state certificate). *Fix a deterministic realization with a finite acyclic event DAG, a static charged schedule, a complete site cut $A \mid B$, and a service window I . Conditional absence is an ordinary charged symbol. With receiver-local exogenous context fixed, let S_I be the receiver's retained state at the start of the window, with $\log_2 |S_I| \leq m_I$; let $\mathcal{A}_I$ contain all forward A -to- B events in the window, with $\ell_a = \log_2 |\Sigma_a|$; and let O_I be a receiver state or output required at the end, with*

$$\kappa_I = \log_2 |\text{im } O_I|$$

as compatible sender histories vary. The certificate is valid when: (i) every retained-state coordinate is available at the beginning of I ; (ii) every charged incoming event occurs before the certified output and lies on a directed path from a source-side input or state event to O_I ; (iii) every directed dependency path by which a varied source-side history can alter O_I intersects either the retained-state set or the charged incoming-event set; and (iv) each charged event is counted once under the registered accounting rule. These conditions are decidable by reachability and cut tests on the finite event DAG.

Theorem 5.15 (Cut-state conservation). *Every temporal cut-state certificate satisfies*

$$\boxed{\kappa_I \leq m_I + \sum_{a \in \mathcal{A}_I} \ell_a}, \quad \boxed{\sum_{a \in \mathcal{A}_I} \ell_a \geq (\kappa_I - m_I)_+}.$$

Interaction and nonlinear local encoding do not improve the inequality.

Proof. The initial receiver state and ordered forward transcript determine every receiver-local event, reverse message, and required output by induction over the event DAG. Their joint alphabet has at least 2^{κ_I} values and at most $2^{m_I} \prod_{a \in \mathcal{A}_I} 2^{\ell_a}$. Taking logarithms proves both forms. $\square$

Theorem 5.16 (Fractional certificate packing). *For any family of valid certificates j , with forward-event sets $\mathcal{A}_j$ and fractional weights $\alpha_j \geq 0$ satisfying*

$$\sum_{j: a \in \mathcal{A}_j} \alpha_j \leq 1 \quad \text{for every event } a,$$

total charged traffic obeys

$$Q \geq \sum_j \alpha_j (\kappa_j - m_j)_+.$$

Proof. This is the standard fractional-packing/linear-programming composition [51]. Multiply each instance of Theorem 5.15 by α_j , sum, exchange the order of summation, and use the per-event charging constraint. $\square$

Conservation as a realization identity. The inequality is an information balance across a temporal cut. At the beginning of the service window, the receiver possesses at most m_I bits of relevant state. During the window, the ordered forward transcript contributes at most $\sum_a \ell_a$ additional bits of log-cardinality. Every receiver state attainable at the end must be indexed by the pair consisting of initial state and forward transcript. The number of required distinctions therefore cannot exceed the product of those two state spaces. Reverse interaction may determine which forward symbol is useful, but once the receiver input is fixed it cannot create a sender-side distinction that never appears in the initial state or forward transcript.

This form makes the storage–communication tradeoff explicit: prepositioning one additional relevant bit can reduce the lower bound on forward traffic by at most one bit in the same service window. Replication, caching, and checkpointing are realizations that move along a joint state–traffic geometry.

Three classical projections follow after adding distinct service contracts.

- (i) *Bandwidth–time.* If a complete cut supplies at most B bits per unit time and no relevant state is prepositioned, cut-state conservation gives

$$BT \geq \kappa.$$

- (ii) *Atomic I/O partition.* Under the standard atomic-word red–blue pebble contract, let $H_D(S)$ be the minimum number of parts in a classical S -partition of the noninput vertices of DAG D . The parts are disjoint and exhaustive, their quotient is acyclic, each has a dominator set of at most S vertices intercepting every input-to-part path, and

$$\text{Min}(V_i) = \{v \in V_i : \text{no immediate successor of } v \text{ lies in } V_i\}$$

has cardinality at most S . Segmenting an execution after every M I/O moves creates a Hong–Kung $2M$ -partition, hence

$$Q \geq M(H_D(2M) - 1).$$

This inequality follows from the atomic-DAG partition spectrum and pebble contract; it is not obtained by silently identifying a memory move with a static tree-cut symbol.

- (iii) *Two-dimensional area–time.* Suppose every admissible balanced bisection has an orientation $A_0 \rightarrow B_0$ with forward service $B_{\text{bis}}^{A_0 \rightarrow B_0} \leq \gamma\sqrt{A}$, a valid oriented cut-state certificate requiring κ_{bis} distinctions, and at most m_{bis} relevant prepositioned bits on B_0 . With

$$\kappa_{\text{bis}}^{\text{net}} = (\kappa_{\text{bis}} - m_{\text{bis}})_+,$$

conservation yields

$$AT^2 \geq \frac{(\kappa_{\text{bis}}^{\text{net}})^2}{\gamma^2}.$$

This projection uses cut-state conservation plus the two-dimensional service and separator law.

Table 1: Temporal cut-state conservation and its separately registered projections. Bandwidth–time follows from the service-window theorem; Hong–Kung I/O and Thompson-style AT^2 require, respectively, atomic-DAG memory and when-and-where-oblivious planar-separator contracts.

Projection	Additional service contract	Resulting bound
Bandwidth–time	Complete cut, service at most B bits per unit time, no relevant prepositioned state	$BT \geq \kappa$
Atomic I/O	Red–blue pebble execution, M -move segmentation, valid $2M$ -partition	$Q \geq M(H_D(2M) - 1)$
Two-dimensional VLSI	Balanced cut, forward service at most $\gamma\sqrt{A}$, net distinction $\kappa_{\text{bis}}^{\text{net}}$	$AT^2 \geq (\kappa_{\text{bis}}^{\text{net}})^2 / \gamma^2$

The causal state-and-transcript count becomes a reusable typed certificate by registering cut, initial state, service window, and accounting rule. Fractional packing composes overlapping certificates, while the three service laws above translate the same semantic conservation statement into bandwidth–time, I/O, and area–time consequences. No genealogy may be inferred from static κ alone, and the Hong–Kung and area–time projections remain model scoped [5, 25].

6 The CRG abstraction ladder and decision-sufficient geometry

Typed units and normalization. CRG signatures are typed rather than dimensionless: state bytes, traffic bytes per operation, event counts, causal depth, work, energy, area, temperature, and reliability coordinates retain their declared units. Coordinates are not assumed commensurate. Convex combinations, norms, or scalar prices are interpreted only after the resource contract or technology ledger supplies the required normalization and unit conversion.

Definition 6.1 (Admissible decomposition system). *The decomposition system $\mathbf{Dec}_\rho(F)$ has as objects the tree hierarchies, terminal placements, and local site rules admitted by Λ . Its structure maps are semantic hierarchy isomorphisms and connected-fiber contractions $q : \mathcal{T} \rightarrow \mathcal{T}'$ equipped with the typed pushforward $q_*\rho$ of Theorem 3.2. Identity and composition are inherited from the tree maps. A contract outside the core calculus is admissible only when its legal-realization semantics and pushforward coherence are supplied explicitly.*

Definition 6.2 (Realization bundle, signature regions, and CRG). *Fix F , ρ , and $\mathcal{T} \in \mathbf{Dec}_\rho(F)$. Let*

$$\mathcal{P}_F^\rho(\mathcal{T})$$

be the registry of ρ -legal hierarchical realizations P , modulo semantic-preserving renaming of internal states. Let

$$\sigma_{\mathcal{T}} : \mathcal{P}_F^\rho(\mathcal{T}) \longrightarrow \mathbb{R}_+^{d(\mathcal{T})}$$

extract the declared typed interface signature, or a declared projection of it. Define the exact attainable region

$$\mathcal{R}_F^\rho(\mathcal{T}) = \sigma_{\mathcal{T}}(\mathcal{P}_F^\rho(\mathcal{T})),$$

the closed upper obligation region

$$\Gamma_F^\rho(\mathcal{T}) = \text{cl}(\text{up } \mathcal{R}_F^\rho(\mathcal{T})),$$

and the convex linear-cost body

$$\mathcal{K}_F^\rho(\mathcal{T}) = \text{cl conv } \Gamma_F^\rho(\mathcal{T}).$$

The computation-realization geometry is the decomposition-indexed system

$$\mathfrak{G}_F^\rho = \left(\mathbf{Dec}_\rho(F), \left\{ \mathcal{P}_F^\rho(\mathcal{T}) \xrightarrow{\sigma_{\mathcal{T}}} \mathcal{R}_F^\rho(\mathcal{T}), \Gamma_F^\rho(\mathcal{T}), \mathcal{K}_F^\rho(\mathcal{T}) \right\}_{\mathcal{T}}, \{\pi_q\}_q \right).$$

The levels retain different information. The registry $\mathcal{P}$ distinguishes architectures, event structures, and embeddings. The exact image $\mathcal{R}$ retains jointly attainable, possibly nonconvex signatures but forgets architecture identity inside a signature fiber. The upper region Γ is the exact quotient for continuous monotone interface-obligation observations. The convex upper body $\mathcal{K}$ is the further quotient seen by additive nonnegative prices and explicitly admitted amortized mixtures. The term *geometry* denotes this resource-indexed feasible system, not a Euclidean placement, signal constellation, or smooth manifold. A selected algorithm may have an intermodule communication geometry in the classical sense [1]; here topology and technology remain unfixed until the contract and observation are chosen.

6.1 Why the geometry has four levels

The ladder is derived from the information required by the decision, not from an aesthetic preference for several sets. Three distinct losses occur.

From $\mathcal{P}$ to $\mathcal{R}$: forgetting realization identity. The signature map partitions $\mathcal{P}$ into fibers

$$\sigma_{\mathcal{T}}^{-1}(r) = \{P \in \mathcal{P}_F^{\rho}(\mathcal{T}) : \sigma_{\mathcal{T}}(P) = r\}.$$

Two realizations in one fiber consume the same retained coordinates, yet can have different event DAGs, schedules, layouts, reliability structure, or legal embeddings. Any objective that inspects those distinctions remains a function on $\mathcal{P}$, not merely on $\mathcal{R}$.

From $\mathcal{R}$ to Γ : forgetting exact slack. For $r \in \mathcal{R}$ and $s \geq r$ coordinatewise, s represents a provisioned budget that can serve the realization with signature r . A continuous monotone cost can never prefer s to r , so adjoining every such s changes neither the optimum nor the complete response map. Closure adds limits of attainable or provisioned sequences, exactly the information needed for continuous queries. Hence $\Gamma = \text{clup } \mathcal{R}$ is the canonical quotient for that observation class.

From Γ to $\mathcal{K}$: forgetting nonconvex distinctions. For an additive nonnegative price w , the value of a signature is $\langle w, r \rangle$. Linear functionals have the same infimum on a set, its convex hull, and its closure. Upper closure is likewise invisible because $w \geq 0$. Therefore every such price sees precisely $\mathcal{K} = \text{clconv } \Gamma$. A nonlinear monotone objective can still distinguish Γ from $\mathcal{K}$, and a realization-sensitive objective can distinguish $\mathcal{P}$ from all three signature images.

Definition 6.3 (Decision equivalence and sufficient representation). *Let $\mathfrak{D}$ be a class of real-valued decision objectives on a realization family. Two retained families $\mathcal{A}, \mathcal{B} \subseteq \mathcal{P}$ are $\mathfrak{D}$ -equivalent when*

$$\inf_{P \in \mathcal{A}} D(P) = \inf_{P \in \mathcal{B}} D(P) \quad \text{for every } D \in \mathfrak{D}.$$

A representation Q is sufficient for $\mathfrak{D}$ when equality of Q for two families implies their $\mathfrak{D}$ -equivalence, and complete when the converse also holds.

Theorems 6.9 and 6.11 show that Γ and $\mathcal{K}$ are complete representations for their respective decision classes. In the terminology of statistical decision theory, they play a role analogous to sufficient information structures; in the terminology of abstract interpretation, they are exact abstractions for the named query classes. CRG adds realization semantics and technology-conditioned selection to that sufficiency statement.

The diagram is also a statement about information flow through a design process. The map $\sigma_{\mathcal{T}}$ forgets which realization produced a given signature. Upper closure forgets whether a larger provisioned point is itself attained, because every monotone resource objective treats it as no better than a dominated attainable point. Convexification forgets nonconvex distinctions that no additive price can observe. No further reduction is licensed without specifying a smaller decision class.

6.2 Invariance and coherent hierarchy maps

Proposition 6.4 (Semantic isomorphism invariance). *Suppose F' is obtained from F by bijective relabelings of input alphabets, output alphabets, and terminal identities. Relabel every hierarchy and conjugate every local realization map by those bijections. This gives coordinate-preserving bijections*

$$\mathcal{R}_F^{\rho}(\mathcal{T}) \cong \mathcal{R}_{F'}^{\rho}(\mathcal{T}'), \quad \Gamma_F^{\rho}(\mathcal{T}) \cong \Gamma_{F'}^{\rho}(\mathcal{T}'), \quad \mathcal{K}_F^{\rho}(\mathcal{T}) \cong \mathcal{K}_{F'}^{\rho}(\mathcal{T}').$$

Thus the geometry depends on semantic structure and the resource contract, not on symbol names or terminal labels.

Proof. Conjugating a legal protocol or circuit by terminal bijections preserves its event structure, resource use, accounting convention, and interface traffic; the inverse relabeling recovers the original realization. $\square$

Definition 6.5 (Hierarchy contraction and signature projection). *A contraction $q : \mathcal{T} \rightarrow \mathcal{T}'$ partitions $V(\mathcal{T})$ into connected fibers and replaces each fiber by one supervertex; terminal placements are pushed forward by q . Since $\mathcal{T}$ is a tree, every edge e' of $\mathcal{T}'$ corresponds to a unique edge e of $\mathcal{T}$ joining the endpoint fibers. For an edge-indexed signature c , define*

$$(\pi_q c)_{e'} = c_e.$$

Other retained coordinates are pushed forward by their declared aggregation laws in $q_* \rho$.

Theorem 6.6 (Hierarchy-contraction consistency). *For every admissible contraction $q : \mathcal{T} \rightarrow \mathcal{T}'$:*

- (a) the canonical directed contextual state on every surviving oriented cut is unchanged up to canonical isomorphism;
- (b) every legal causal realization on $\mathcal{T}$ pushes forward to a legal realization on $\mathcal{T}'$ with retained traffic vector $\pi_q c$ under $q_*\rho$; and
- (c) consequently,

$$\pi_q \mathcal{R}_F^\rho(\mathcal{T}) \subseteq \mathcal{R}_F^{q_*\rho}(\mathcal{T}'), \quad \pi_q \Gamma_F^\rho(\mathcal{T}) \subseteq \Gamma_F^{q_*\rho}(\mathcal{T}'), \quad \pi_q \mathcal{K}_F^\rho(\mathcal{T}) \subseteq \mathcal{K}_F^{q_*\rho}(\mathcal{T}').$$

Proof. For a surviving edge, contraction does not change the two terminal sets separated by its cut, so the contextual equivalence relation is identical. For a general protocol, treat each connected fiber as one local superprocessor and execute all events internal to it locally. Only messages crossing the unique fine edge corresponding to a coarse edge remain, producing $\pi_q c$. Upper closure, topological closure, convexification, and coordinate projection give the stated inclusions. $\square$

Theorem 6.7 (Decomposition-system coherence). *The structure maps satisfy*

$$\pi_{\text{id}_\mathcal{T}} = \text{id}, \quad \pi_{r \circ q} = \pi_r \circ \pi_q,$$

and their contracts satisfy $(r \circ q)_*\rho = r_*(q_*\rho)$. Semantic hierarchy isomorphisms act by coordinate bijections and commute with the projections. Thus $\mathfrak{G}_F^\rho$ is a coherent decomposition-indexed system, not a bare collection of feasible regions.

Proof. Every surviving coarse edge corresponds to the same unique original tree edge whether contractions are applied at once or successively, so coordinate restriction composes exactly. Connected-fiber localization and aggregation of retained resource charges are associative by Theorem 3.3. Relabelings only rename coordinates and local symbols and hence commute with projection. $\square$

6.3 Monotone reconstruction and linear duality

Definition 6.8 (Monotone interface-cost response). *For a continuous coordinatewise nondecreasing $\Phi : \mathbb{R}_+^{d(\mathcal{T})} \rightarrow \mathbb{R}_+$, define*

$$V_{F,\rho,\mathcal{T}}(\Phi) = \inf_{r \in \mathcal{R}_F^\rho(\mathcal{T})} \Phi(r),$$

with value $+\infty$ if the hierarchy is infeasible.

Theorem 6.9 (Reconstruction by monotone interface observations). *For every feasible hierarchy $\mathcal{T}$:*

- (a) every continuous coordinatewise nondecreasing cost satisfies

$$V_{F,\rho,\mathcal{T}}(\Phi) = \inf_{x \in \Gamma_F^\rho(\mathcal{T})} \Phi(x);$$

- (b) the full response $\Phi \mapsto V_{F,\rho,\mathcal{T}}(\Phi)$ reconstructs $\Gamma_F^\rho(\mathcal{T})$ exactly. In particular, for $a \in \mathbb{R}_+^{d(\mathcal{T})}$, let

$$\Phi_a(x) = \|(x - a)_+\|_\infty.$$

Then

$$a \in \Gamma_F^\rho(\mathcal{T}) \iff V_{F,\rho,\mathcal{T}}(\Phi_a) = 0;$$

- (c) the decomposition-indexed response functional $(\mathcal{T}, \Phi) \mapsto V_{F,\rho,\mathcal{T}}(\Phi)$ and the family of upper obligation regions $\{\Gamma_F^\rho(\mathcal{T})\}_\mathcal{T}$ determine one another, up to coordinate relabelings declared by ρ .

Conversely, those regions determine every query in this observation class. The completeness claim does not extend automatically to objectives depending on a realization fiber, causal timing, switching activity, or physical embedding.

Proof. Monotonicity gives $\inf_{\text{up}} \mathcal{R} \Phi = \inf_{\mathcal{R}} \Phi$: increasing any provisioned coordinate cannot improve the cost. Continuity preserves the infimum under closure, proving (a).

If $a \in \Gamma$, then $\Phi_a(a) = 0$. Conversely, if the infimum is zero, choose $x_n \in \Gamma$ with $\|(x_n - a)_+\|_\infty \rightarrow 0$ and set $y_n = x_n \vee a$ coordinatewise. Upper closure gives $y_n \in \Gamma$, while $y_n \rightarrow a$; closedness gives $a \in \Gamma$. Applying this membership test on every admissible hierarchy proves (c), and the converse follows from (a). $\square$

Definition 6.10 (Lower linear support cost). *For $w \in \mathbb{R}_+^{d(\mathcal{T})}$, define*

$$h_{F,\rho,\mathcal{T}}(w) = \inf_{c \in \mathcal{K}_F^\rho(\mathcal{T})} \langle w, c \rangle.$$

Theorem 6.11 (Linear-cost-body duality). *Let $\mathcal{K} \subseteq \mathbb{R}_+^m$ be nonempty, closed, convex, and upper closed, and put $h_{\mathcal{K}}(w) = \inf_{x \in \mathcal{K}} \langle w, x \rangle$ for $w \geq 0$. Then*

$$\mathcal{K} = \{x \in \mathbb{R}_+^m : \langle w, x \rangle \geq h_{\mathcal{K}}(w) \text{ for every } w \in \mathbb{R}_+^m\}.$$

Thus the optimum under every additive nonnegative interface-price system uniquely determines the linear-cost body [2].

Proof. Every $x \in \mathcal{K}$ satisfies the inequalities by definition. Conversely, let $x \notin \mathcal{K}$. Strong separation gives a and α with

$$\langle a, x \rangle < \alpha \leq \langle a, y \rangle \quad \text{for every } y \in \mathcal{K}.$$

The separator must satisfy $a \geq 0$: if $a_j < 0$, upper closure would put $y + te_j$ in $\mathcal{K}$ for every $t \geq 0$, driving $\langle a, y + te_j \rangle$ to $-\infty$. Hence a is an allowed price vector and $\langle a, x \rangle < h_{\mathcal{K}}(a)$. $\square$

Corollary 6.12 (Linear-price completeness). *For a fixed hierarchy, $\mathcal{K}_F^p(\mathcal{T})$ and the response $h(w)$ determine one another, and distinct bodies are separated by a nonnegative price vector. Convexification is exact for additive linear interface prices; general monotone costs require Γ , while objectives distinguishing realizations with one signature require $\mathcal{P}$.*

Definition 6.13 (Decision-response equivalence). *Let $\mathfrak{C}$ be a class of objectives on a common signature space. Two nonempty realization families $\mathcal{A}$ and $\mathcal{B}$ are $\mathfrak{C}$ -equivalent when*

$$\inf_{P \in \mathcal{A}} C(\sigma_{\mathcal{T}}(P)) = \inf_{P \in \mathcal{B}} C(\sigma_{\mathcal{T}}(P)) \quad \text{for every } C \in \mathfrak{C}.$$

A statistic of a realization family is decision sufficient for $\mathfrak{C}$ when equality of that statistic is equivalent to $\mathfrak{C}$ -equivalence.

Theorem 6.14 (Canonical decision quotients). *Fix F , ρ , and $\mathcal{T}$.*

- (a) *For the class $\mathfrak{C}_{\uparrow}$ of continuous coordinatewise nondecreasing signature objectives, the map $\mathcal{A} \mapsto \Gamma_{\mathcal{A}}$ is a complete decision-sufficient statistic.*
- (b) *For the class $\mathfrak{C}_{+}$ of additive objectives $c \mapsto \langle w, c \rangle$ with $w \geq 0$, the map $\mathcal{A} \mapsto \mathcal{K}_{\mathcal{A}}$ is a complete decision-sufficient statistic.*
- (c) *Neither quotient is complete for objectives that distinguish realizations in the same signature fiber; those objectives require the registry $\mathcal{P}$ or an enriched signature.*

Consequently, the ladder $\mathcal{P} \rightarrow \mathcal{R} \rightarrow \Gamma \rightarrow \mathcal{K}$ is an ordered family of exact observation quotients rather than a sequence of heuristic summaries.

Proof. Part (a) follows from Theorem 6.9: the complete monotone response reconstructs the upper region, and the upper region computes every response. Part (b) follows from Theorems 6.11 and 6.12. For (c), choose two distinct realizations P_1, P_2 with $\sigma_{\mathcal{T}}(P_1) = \sigma_{\mathcal{T}}(P_2)$ and an objective that assigns different costs to their schedules, causal structures, or embeddings. Every statistic factoring only through the common signature gives them the same value, whereas the realization-sensitive objective separates them. $\square$

The theorem is a deterministic design analogue of statistical decision sufficiency: the sufficient object is defined by the family of questions that must remain answerable. It also supplies the precision order for abstraction. When only additive prices remain, $\mathcal{K}$ is sufficient. When nonlinear monotone constraints or bottleneck objectives remain, Γ is required. When exact discrete alternatives matter, $\mathcal{R}$ is required. When schedule, architecture, or physical embedding matters, the realization fiber must be retained.

6.4 Safe commitment and insufficiency of scalar summaries

For a retained family $\emptyset \neq \mathcal{A} \subseteq \mathcal{P}_F^p(\mathcal{T})$ and an admissible signature objective C , define the commitment regret

$$\text{Regret}_C(\mathcal{A}) = \inf_{P \in \mathcal{A}} C(\sigma_{\mathcal{T}}(P)) - \inf_{P \in \mathcal{P}_F^p(\mathcal{T})} C(\sigma_{\mathcal{T}}(P)) \geq 0. \quad (6.1)$$

A reduction is safe for a declared decision class exactly when it leaves every optimum in that class unchanged. The next theorem identifies the corresponding geometric invariants.

Theorem 6.15 (Safe-Commitment Law: universal lossless retention). *Let $\emptyset \neq \mathcal{A} \subseteq \mathcal{P}_F^\rho(\mathcal{T})$ be a retained realization family and define*

$$\Gamma_{\mathcal{A}} = \text{clup } \sigma_{\mathcal{T}}(\mathcal{A}), \quad \mathcal{K}_{\mathcal{A}} = \text{clconv } \Gamma_{\mathcal{A}}.$$

Then:

- (i) $\mathcal{A}$ preserves the optimum of every continuous coordinatewise nondecreasing signature objective if and only if $\Gamma_{\mathcal{A}} = \Gamma_F^\rho(\mathcal{T})$;
- (ii) it preserves the optimum under every additive nonnegative price if and only if $\mathcal{K}_{\mathcal{A}} = \mathcal{K}_F^\rho(\mathcal{T})$; and
- (iii) if either equality fails, a member of the corresponding observation class has strictly positive commitment regret in the sense of Equation (6.1).

Proof. Suppose first that $\Gamma_{\mathcal{A}} = \Gamma_F^\rho(\mathcal{T})$. Every continuous monotone objective has the same infimum on a realization family and on the upper closure of its signatures by Theorem 6.9; hence the retained and full optima agree. The same argument with linear-cost duality proves sufficiency of $\mathcal{K}_{\mathcal{A}} = \mathcal{K}_F^\rho(\mathcal{T})$ for additive prices.

For necessity in the monotone class, assume every monotone optimum agrees. For any $a \in \mathbb{R}_+^{d(\mathcal{T})}$, apply the membership probe $\Phi_a(x) = \|(x - a)_+\|_\infty$. The retained response is zero exactly when $a \in \Gamma_{\mathcal{A}}$, and the full response is zero exactly when $a \in \Gamma_F^\rho(\mathcal{T})$. Equality of all responses therefore gives pointwise equality of the two regions.

For necessity in the linear-price class, equality of every nonnegative linear response gives the same support function on the nonnegative orthant. Theorem 6.11 reconstructs the two closed convex upper bodies from those responses, so they are equal.

Finally, if $\Gamma_{\mathcal{A}} \neq \Gamma_F^\rho(\mathcal{T})$, choose a in their set difference. The probe Φ_a has full optimum zero and retained optimum strictly positive, producing commitment regret. If the convex bodies differ, strong separation by a nonnegative vector w gives an additive price whose retained optimum is strictly larger than the full optimum. These objectives are explicit witnesses to the information lost by the reduction. $\square$

Corollary 6.16 (Strict engineering-cost separation). *Let $\Gamma_{\mathcal{A}} \subsetneq \Gamma$ be closed upper sets generated by a retained family $\mathcal{A}$ and by the full legal realization family, respectively. Then there exists a continuous, coordinatewise strictly increasing, coercive cost $\Phi : \mathbb{R}_+^d \rightarrow \mathbb{R}_+$ such that*

$$\inf_{x \in \Gamma} \Phi(x) < \inf_{x \in \Gamma_{\mathcal{A}}} \Phi(x).$$

Consequently, loss of the monotone-decision geometry is witnessed by an engineering-style cost that penalizes every declared resource, not only by a membership probe.

Proof. Choose $a \in \Gamma \setminus \Gamma_{\mathcal{A}}$ and set

$$\delta := \inf_{x \in \Gamma_{\mathcal{A}}} \|(x - a)_+\|_\infty.$$

Closed upperness implies $\delta > 0$: otherwise a sequence in $\Gamma_{\mathcal{A}}$ would approach a after taking coordinatewise maxima with a , and closedness would place a in $\Gamma_{\mathcal{A}}$. Choose

$$0 < \varepsilon < \frac{\delta}{1 + \sum_i a_i}$$

and define

$$\Phi(x) := \|(x - a)_+\|_\infty + \varepsilon \sum_i x_i.$$

The positive linear term makes Φ coordinatewise strictly increasing and coercive. Every $x \in \Gamma_{\mathcal{A}}$ satisfies $\Phi(x) \geq \delta$, whereas $\Phi(a) = \varepsilon \sum_i a_i < \delta$. Since $a \in \Gamma$, strict separation follows. $\square$

Corollary 6.17 (Finite-registry retention tests). *Suppose $\mathcal{P}_F^\rho(\mathcal{T})$ is finite.*

- (a) *A retained family is complete for all continuous monotone signature objectives exactly when every discarded signature is coordinatewise dominated by a retained signature.*
- (b) *It is complete for all additive nonnegative prices exactly when every full signature lies in $\text{conv}(\sigma_{\mathcal{T}}(\mathcal{A})) + \mathbb{R}_+^{d(\mathcal{T})}$.*
- (c) *If a named finite technology ledger Θ has attained minima, the smallest requirement is argmin hitting: $\mathcal{A} \cap \text{argmin}_P C_\theta(P) \neq \emptyset$ for every $\theta \in \Theta$.*

Proof. For a finite set, the minimal points of its upper closure are precisely the nondominated signatures, proving (a). Convexification of a finite upper set is $\text{conv}(\sigma(\mathcal{P})) + \mathbb{R}_+^d$, proving (b). For (c), the retained and full minima agree at technology θ exactly when the retained family contains at least one member of the attained full argmin set; quantifying over Θ proves the statement. $\square$

Safe-commitment principle. Within the declared observation class, a realization may be discarded only when the relevant geometry is unchanged. Equivalently, discard a realization only when every future decision in that class is unable to detect that it is gone.

Derivation of the law

The equivalence can be read constructively in both directions.

Sufficiency. If $\Gamma_{\mathcal{A}} = \Gamma$, then every continuous monotone objective has the same infimum on the retained and full families because both optimize over the same complete upper region. If $\mathcal{K}_{\mathcal{A}} = \mathcal{K}$, the same statement holds for every additive nonnegative price by linear-cost duality.

Necessity. Suppose the retained upper region omits a point $a \in \Gamma$. The probe

$$\Phi_a(x) = \|(x - a)_+\|_\infty$$

has full-family optimum zero but retained-family optimum strictly positive; it is therefore an explicit witness of commitment regret. If instead the convex bodies differ, a nonnegative separating price w gives a strict value gap. The law is thus not a heuristic preference for diversity. It is a complete equivalence between a geometric equality and preservation of an entire class of future decisions.

Three worked retention regimes

Universal monotone regime.

Every resource objective that is continuous and worsens when any coordinate increases remains admissible. A discarded finite-registry point is safe exactly when a retained point dominates it.

Additive-price regime.

Future decisions are restricted to $w \cdot r$ with $w \geq 0$. A discarded point can remain safe even when it is not dominated, provided it lies in the convex upper hull generated by retained points.

Named technology ledger.

Only a finite set of technology/objective pairs remains open. It is sufficient to retain a family that intersects the argmin set of every ledger entry, as formalized later by the argmin-hitting theorem. This quotient can be much smaller than either universal geometry.

These regimes separate a universal theorem from a practical retention policy. The decision class is declared first; the smallest sufficient realization record follows from that declaration.

For a finite registry, universal monotone retention means that every discarded signature is dominated by a retained one. Linear-price retention permits the weaker test

$$\mathcal{R}_F^\rho(\mathcal{T}) \subseteq \text{conv}(\sigma_{\mathcal{T}}(\mathcal{A})) + \mathbb{R}_+^{d(\mathcal{T})}.$$

These statements are universal only relative to their named observation classes; a fixed finite technology ledger can require less data.

Theorem 6.18 (Strict insufficiency of marginal and scalar summaries). *There are two closed upper feasible regions with identical coordinate floors, minimum total traffic, and minimum bottleneck width, but different technology-price responses. Let*

$$\Gamma_A = \text{up}\{(1, 2), (2, 1)\}, \quad \Gamma_B = \text{up}\{(1, 3), (2, 1)\}.$$

Both satisfy

$$(\inf x_1, \inf x_2) = (1, 1), \quad \inf(x_1 + x_2) = 3, \quad \inf \max\{x_1, x_2\} = 2.$$

Yet for $w = (10, 1)$,

$$\inf_{x \in \Gamma_A} \langle w, x \rangle = 12, \quad \inf_{x \in \Gamma_B} \langle w, x \rangle = 13.$$

The jointly realizable region therefore retains operational information lost by coordinatewise minima, total width, and bottleneck width.

Proof. Evaluate the two Pareto-minimal points of each upper region. Under $w = (10, 1)$, $(1, 2)$ minimizes Γ_A at 12, while $(1, 3)$ minimizes Γ_B at 13; $(2, 1)$ costs 21 in both. $\square$

Example 6.19 (Why the exact region cannot be convexified universally). Suppose the only exact signatures are $r = (0, 2)$ and $s = (2, 0)$. Their convex linear-cost body contains $(1, 1)$. Every linear price has the same infimum on the two raw points and on the body, but the monotone convex objective $\Phi(c) = \max\{c_1, c_2\}$ has exact optimum 2 and convex-relaxation optimum 1. Thus $\mathcal{K}$ is exact for linear prices, $\mathcal{R}$ preserves discrete joint alternatives, and $\mathcal{P}$ is required to preserve architecture identity within each signature fiber.

Proposition 6.20 (Price reversal for incomparable signatures). If $r, s \in \mathbb{R}_+^m$ are incomparable in coordinate order, there are nonnegative prices w, w' such that

$$\langle w, r \rangle < \langle w, s \rangle, \quad \langle w', s \rangle < \langle w', r \rangle.$$

No price-independent monotone scalar ranking can therefore agree with every nonnegative linear price system.

Proof. Incomparability gives coordinates i, j with $r_i < s_i$ and $r_j > s_j$. Take $w = e_i$ and $w' = e_j$. $\square$

7 Monotonicity and composition

Proposition 7.1 (Resource monotonicity). If $\rho \preceq \rho'$, then on every hierarchy admissible under both contracts,

$$\mathcal{R}_F^\rho(\mathcal{T}) \subseteq \mathcal{R}_F^{\rho'}(\mathcal{T}), \quad \Gamma_F^\rho(\mathcal{T}) \subseteq \Gamma_F^{\rho'}(\mathcal{T}), \quad \mathcal{K}_F^\rho(\mathcal{T}) \subseteq \mathcal{K}_F^{\rho'}(\mathcal{T}),$$

and

$$h_{F, \rho', \mathcal{T}}(w) \leq h_{F, \rho, \mathcal{T}}(w) \quad \text{for all } w \geq 0.$$

If Λ' also admits more hierarchies than Λ , a global hierarchy optimum can only improve further.

Proof. Relaxation is inclusion of legal realizations by definition. Applying $\sigma_{\mathcal{T}}$, upper closure, closure, and convexification preserves inclusion. Taking an infimum over the larger body proves the price inequality. $\square$

Proposition 7.2 (Local data processing). Suppose $F' = p_{\text{out}} \circ F \circ p_{\text{in}}$, where p_{in} and p_{out} are implementable locally at the input and output terminals for free under ρ . Every realization of F then induces a realization of F' with the same interface signature. Hence

$$\mathcal{K}_{F'}^\rho(\mathcal{T}) \supseteq \mathcal{K}_F^\rho(\mathcal{T}), \quad h_{F', \rho, \mathcal{T}}(w) \leq h_{F, \rho, \mathcal{T}}(w).$$

Local processing can encode or discard communication demand, but cannot make the derived task more expensive under the same interface prices.

Proof. Precompose each input terminal with p_{in} , run the realization of F , and postcompose its output terminals with p_{out} . The added maps are local and free, so no interface coordinate changes. Region inclusion and the support-cost inequality follow. $\square$

Theorem 7.3 (Exact product additivity for canonical summaries). Let F_1 and F_2 have disjoint input and output alphabets placed on the same tree, and consider their product computation $F_1 \times F_2$. In the separable directed-summary model,

$$Q_{u \rightarrow v}^{F_1 \times F_2} \cong Q_{u \rightarrow v}^{F_1} \times Q_{u \rightarrow v}^{F_2}$$

on every oriented edge, and therefore

$$\kappa_{u \rightarrow v}(F_1 \times F_2) = \kappa_{u \rightarrow v}(F_1) + \kappa_{u \rightarrow v}(F_2).$$

If $\hat{\kappa} = \lceil \kappa \rceil$ is the minimum fixed-bit width of one canonical summary, then

$$\hat{\kappa}(F_1) + \hat{\kappa}(F_2) - 1 \leq \hat{\kappa}(F_1 \times F_2) \leq \hat{\kappa}(F_1) + \hat{\kappa}(F_2).$$

Semantic capacities add exactly; joint fixed-bit packing can save at most one bit per directed interface relative to separate rounding.

Proof. If two product-side assignments are equivalent for both component outputs on the opposite side, they are equivalent for the ordered output pair. Conversely, if they are inequivalent for either component, fix all variables of the other component and use a distinguishing opposite-side context for the first; the ordered output differs. The equivalence relation is therefore the Cartesian product relation, class cardinalities multiply, and logarithms add. The fixed-bit inequalities are the elementary bounds for the ceiling of a sum. $\square$

Proposition 7.4 (Constructive subadditivity and communication synergy). *Let F_1, F_2 be governed by ρ_1, ρ_2 , and suppose the product contract $\rho_1 \oplus \rho_2$ permits side-by-side execution with additive local resources. In any interactive realization model closed under that parallel execution,*

$$h_{F_1 \times F_2, \rho_1 \oplus \rho_2, \mathcal{T}}(w) \leq h_{F_1, \rho_1, \mathcal{T}}(w) + h_{F_2, \rho_2, \mathcal{T}}(w).$$

Define

$$\text{Syn}_w(F_1, F_2) = h_{F_1, \rho_1, \mathcal{T}}(w) + h_{F_2, \rho_2, \mathcal{T}}(w) - h_{F_1 \times F_2, \rho_1 \oplus \rho_2, \mathcal{T}}(w) \geq 0.$$

This is the price-weighted communication saved by shared encodings, intermediate computation, or control relative to separate realization. Exact direct-sum behavior is the case $\text{Syn}_w = 0$; information-complexity direct-sum results are neighboring tools [3, 4].

Proof. Choose realizations of the two tasks arbitrarily close to their separate infima and execute them side by side. The product contract makes their signatures additive, so the resulting price is the sum. Taking the infimum over all joint realizations gives the inequality and nonnegativity of Syn_w . $\square$

Serial composition has the analogous constructive bound: concatenate realizations and eliminate the intermediate terminal. Joint optimization may improve it by fusing the intermediate representation; that serial fusion gain is the sequential analogue of Syn_w . It is a quantity defined by the full geometry, not an automatic equality.

8 Analytic component certificates

The design calculus becomes operational when a computation has both a scoped converse and constructive realizations. Dense matrix multiplication and exact attention provide complementary cases. Matrix multiplication exposes a memory–traffic–collective frontier; attention exposes a traffic–storage–recomputation frontier. Neither has a technology-independent scalar winner.

8.1 Dense matrix multiplication across memory levels

Consider classical $n \times n$ dense matrix multiplication restricted to the conventional $\Theta(n^3)$ scalar-product DAG. For a two-level memory hierarchy with M fast words, established pebble and geometric bounds give $Q(M) = \Omega(n^3/\sqrt{M} + n^2)$, while blocked algorithms attain the same order under the usual cache assumptions [5, 6].

Theorem 8.1 (Multilevel matrix-multiplication bracket). *Let $M_1 < \dots < M_L$ be the cumulative fast-side capacities exposed by an inclusive multilevel contract, and let q_i be words moved across interface i . Inputs begin beyond each modeled interface and the final output is returned beyond it. There are constants $c_0, C_0 > 0$ such that every legal classical realization satisfies*

$$q_i \geq c_0 \left(\frac{n^3}{\sqrt{M_i}} + n^2 \right) \quad (1 \leq i \leq L),$$

while one cache-oblivious recursively blocked realization simultaneously satisfies

$$q_i \leq C_0 \left(\frac{n^3}{\sqrt{M_i}} + n^2 \right) \quad (1 \leq i \leq L)$$

at every level obeying the declared ideal-cache assumptions [7]. Thus one outer orthant and one jointly realizable inner point certify the multilevel traffic projection within a constant factor.

Proof. At level i , contract every faster level into the declared cumulative M_i -word fast side and apply the two-level lower bound to that complete interface. For achievability, the same cache-oblivious recursion induces a legal blocked schedule at every ideal-cache level simultaneously. The claim therefore does not combine incompatible per-level optima or count one transfer at several levels without accounting. A noninclusive hierarchy must replace M_i by the actual capacity available on the contracted fast side. $\square$

The same structured computation also yields a scale exponent. For a set $S \subseteq [n]^3$ of $g = |S|$ multiplication triples, the discrete Loomis–Whitney inequality gives

$$|S|^2 \leq |\pi_A(S)| |\pi_B(S)| |\pi_C(S)|.$$

Table 2: Declared public technology-coordinate library. Payload rates are ideal one-direction cluster projections, not protocol or contention measurements.

Technology row	Payload	Link energy	Status
UCIe-S, 32 GT/s	64 GB/s	0.50 pJ/bit	public target anchor
UCIe-A, 32 GT/s	256 GB/s	0.25 pJ/bit	public target anchor
UCIe-A, 64 GT/s	512 GB/s	0.25 pJ/bit	rate anchor; energy proxy
Local fast state	32 or 96 MiB	–	public cache-capacity anchors

Table 3: Modeled $n = 8192$, $K_{\text{chip}} = 64$ projection with two-byte words. Time uses one ideal UCIe-A32 cluster on the critical interface.

c	Memory/chiplet	Traffic/chiplet	Stages	Transfer time	Link energy
1	6.0 MiB	32.00 MiB	8	131.1 μs	2.147 mJ
2	12.0 MiB	22.63 MiB	4	92.7 μs	1.519 mJ
4	24.0 MiB	16.00 MiB	3	65.5 μs	1.074 mJ

Hence at least one operand/output projection has size $\Omega(g^{2/3})$; cubic blocks attain $O(g^{2/3})$. The natural classical blocked realization therefore has the interior demand law

$$R(g) = \Theta(g^{2/3}).$$

This is a result for the registered structured computation, not for every Boolean circuit computing matrix multiplication.

8.2 A certified 2.5D technology map

On K_{chip} processor-chiplets with per-chiplet fast memory M , classical distributed lower bounds include a per-chiplet communication floor of order $\Omega(n^3/(K_{\text{chip}}\sqrt{M}))$ under the declared distribution assumptions [8, 6]. The 2.5D family introduces a replication factor $1 \leq c \leq K_{\text{chip}}^{1/3}$ and, up to constants and logarithmic latency terms, obeys

$$M(c) = \Theta\left(\frac{3cn^2}{K_{\text{chip}}}\right), \quad Q(c) = \Theta\left(\frac{n^2}{\sqrt{cK}}\right), \quad R(c) = O\left(\sqrt{\frac{K_{\text{chip}}}{c^3}} + \log c\right)$$

[9]. The numerical certificate declares the constructive leading model

$$q(c) = \frac{2n^2}{\sqrt{cK}}, \quad m(c) = \frac{3cn^2}{K_{\text{chip}}}, \quad r(c) = \left\lceil \sqrt{K_{\text{chip}}/c^3} \right\rceil + \lceil \log_2 c \rceil.$$

The constants are model coordinates; the exchange of memory for traffic and collective stages is inherited from the cited algorithms and converses.

For $K_{\text{chip}} = 64$, two-byte words, and $c \in \{1, 2, 4\}$, feasibility under local memory budget M_{max} is exactly

$$\frac{3cn^2}{K_{\text{chip}}}(2 \text{ bytes}) \leq M_{\text{max}}.$$

Both (qc) and (rc) decrease with c , so the largest feasible replication factor is Pareto dominant for every monotone traffic-round objective in this declared family.

The UCIe rows use public rate classes and the declared energy proxies [10, 11]; the local-state anchors follow public processor documentation [12]. Together they define a public-anchor technology library whose rows are explicit technology coordinates for the design calculus.

The resulting memory-conditioned selection is

32 MiB/chiplet	c^*
$n \leq 9459$	4
$9459 < n \leq 13377$	2
$13377 < n \leq 18918$	1
$n > 18918$	infeasible in the declared local-memory model.




Figure 1: The 2.5D realization family is retained until memory and link technology are declared.

With 96 MiB per chiplet, the $c = 4$ region extends to $n = 16384$. The public link rows quantify the value of a selected realization but do not change the largest-feasible- c ordering in this idealized family.

Collective capability can reverse a different pair of realizations. At $K_{\text{chip}} = 64$, a one-dimensional replicated- B point with effective collective stage count r has typed signature

$$z_{1D}(r) = (63, r, 1.03125),$$

while the declared two-dimensional SUMMA point is $z_{2D} = (14, 8, 0.046875)$ [13, 14]. Under price $C = w_Q q + w_R r + w_M m$, the one-dimensional realization can win only for $r < 8$, and then exactly when

$$w_R > \frac{49w_Q + 0.984375w_M}{8 - r}.$$

The reversal is attached to an explicit collective capability rather than an unstated communication primitive.

8.3 Exact attention: recomputation versus storage

Let $Q, K, V \in \mathbb{R}^{N \times d}$ and consider exact attention $O = \text{softmax}(QK^T)V$ with fast memory M scalar elements. The conventional realization materializes $N \times N$ score and probability matrices; the tiled realization retains rowwise normalization state and recomputes blocks when required.

Theorem 8.2 (Imported exact-attention I/O bracket). *For $d \leq M \leq Nd$, standard exact attention has*

$$Q_{\text{std}} = \Theta(Nd + N^2), \quad S_{\text{std}} = \Theta(N^2),$$

whereas *FlashAttention* achieves

$$Q_{\text{flash}} = \Theta\left(\frac{N^2 d^2}{M}\right), \quad S_{\text{flash}} = O(N),$$

while spending additional local work through recomputation. The cited analysis proves I/O optimality over the stated SRAM range [15].

The theorem supplies incomparable typed points, not one scalar winner. For $d = 64$ and 128, the asymptotic crossover $M > d^2$ corresponds to only 8 and 32 KiB for two-byte elements, but implementation constants and additional work still require an explicit objective.

For the published A100 forward-plus-backward GPT-2-medium measurement, the reported points are

$$z_{\text{std}} = (Q, F, T) = (40.3 \text{ GB}, 66.6 \text{ GFLOP}, 41.7 \text{ ms}),$$

$$z_{\text{flash}} = (4.4 \text{ GB}, 75.2 \text{ GFLOP}, 7.3 \text{ ms})$$




Figure 2: Traffic, work, and storage jointly determine the attention realization boundary.

[15]. Under typed additive cost $C = w_Q Q + w_F F + w_S S$, the tiled realization is preferred exactly when

$$35.9w_Q + \Delta S w_S > 8.6w_F, \quad \Delta S = S_{\text{std}} - S_{\text{flash}} > 0.$$

With zero storage price, the break-even is approximately 0.240 FLOP/byte. Public A100 peak/bandwidth ratios are approximately $838\times$ larger on the BF16 tensor-peak basis and about $52\times$ larger on the FP32 basis [16]; this establishes the arithmetic-intensity separation; measured kernel time additionally depends on utilization, scheduling, and the remaining typed coordinates.

8.4 Cross-stage state and layout laws

The end-to-end application joins attention and an MLP. Two general facts prevent isolated stage optima from being multiplied together without a joint resource check.

Theorem 8.3 (Static provisioning and time-shared state). *Let two isolated modes require usable state m_A and m_F on a chiplet with budget M .*

- (i) *A static fused realization with dedicated simultaneously instantiated engines is infeasible when $m_A + m_F > M$.*
- (ii) *A sequential universal realization that reuses one state pool is state-feasible when $\max(m_A, m_F) \leq M$.*

The two statements belong to different realization contracts; time sharing must expose any efficiency cost rather than treating reconfiguration as free.

Proof. Dedicated state occupies disjoint capacity and adds. State used in disjoint time intervals occupies the maximum. An embedding cannot transform one resource aggregator into the other without changing the realization. $\square$

For layout conversion, let $K_{\text{chip}} = CT = C'T'$ chiplets implement aligned two-dimensional layouts over a normalized activation tensor. Let $S_{i=1}^{K_{\text{chip}}}$ and $T_{j=1}^{K_{\text{chip}}}$ be the equal-area source and target rectangles. Each tensor element has one source and one target owner; chiplet relabeling is free and bijective. Compression, recomputation, replication, coding, cached duplicates, unequal-capacity placement, routed congestion, latency, and energy are outside this conversion contract.

Theorem 8.4 (Exact aligned-layout assignment). *After optimal chiplet relabeling, the exact retained and minimum moved fractions are*

$$r^*(C, C') = \max_{\pi \in \mathfrak{S}_{K_{\text{chip}}}} \sum_{i=1}^{K_{\text{chip}}} |S_i \cap T_{\pi(i)}|, \quad \mu^*(C, C') = 1 - r^*(C, C').$$

If $C \mid C'$ or $C' \mid C$, then

$$r^*(C, C') = \frac{\min(C, C')}{\max(C, C')}, \quad \mu^*(C, C') = 1 - \frac{\min(C, C')}{\max(C, C')}.$$




Figure 3: The $K_{\text{chip}} = 12$ nonnested example separates the general assignment law from the nested closed form.

For nonnested factors the closed form need not hold.

Proof. Under a relabeling π , chiplet i retains exactly $|S_i \cap T_{\pi(i)}|$, so maximizing retention is a maximum-weight perfect assignment. Retain every selected intersection and move only each target remainder from its unique source owner; this realizes movement equal to one minus retained measure. Conversely, every element outside the retained intersections must change owner.

For the nested case $C' = kC$, each nonempty intersection has area at most $1/(kK)$, so total retention is at most $1/k$. The intersection graph is k -regular and bipartite, hence has a perfect matching attaining $1/k$. The case $C > C'$ is symmetric. $\square$

At $K_{\text{chip}} = 12, C = 4, C' = 6$, the exact retained fraction is $1/2$, not $2/3$, which shows why nesting is load bearing. Auto-sharding systems already price resharding; the result here is the exact optimal-relabeling law under its declared ownership contract [17, 18].

These component certificates now supply the end-to-end case with scoped converses, constructive stage modes, state aggregators, and an exact ownership conversion law. The remaining question is which joint realization wins once all resources and package obligations are conserved simultaneously.

9 Worked realization geometries

The following examples make the abstraction ladder explicit without relying on asymptotics or physical calibration.

9.1 Broadcast and parity

Place one binary source at a vertex of a tree and require copies at a set of terminals. Across an oriented edge that separates the source from at least one consumer, the two assignments 0 and 1 are contextually inequivalent, so the canonical forward quotient has two states and capacity one bit. The reverse quotient is trivial. The same one-bit

distinction can cross arbitrarily many interfaces as the distribution tree grows; semantic information and physical multiscale incidence are therefore distinct coordinates.

If every vertex instead holds a source bit and must output the parity of all source bits, each side of every nontrivial cut has one local parity summary. One bit is required in each direction. Broadcast and parity thus have the same one-bit local state size but different directional obligation geometry.

9.2 INDEX and price reversal

In the INDEX problem, Alice holds $x \in \{0, 1\}^n$, Bob holds $i \in [n]$, and Bob outputs x_i . A one-way Alice-to-Bob realization has signature $r = (n, 0)$ in the two directional traffic coordinates. An interactive realization has signature $s = (1, \lceil \log_2 n \rceil)$. For price $w = (w_{AB}, w_{BA})$,

$$C_w(r) = nw_{AB}, \quad C_w(s) = w_{AB} + \lceil \log_2 n \rceil w_{BA}.$$

The interactive protocol is selected when $(n-1)w_{AB} > \lceil \log_2 n \rceil w_{BA}$; the one-way protocol is selected in the reverse regime. The computation has no price-independent preferred protocol.

9.3 Why $\mathcal{R}$, Γ , and $\mathcal{K}$ differ

Let the exact signatures be $r = (0, 2)$ and $s = (2, 0)$. Then

$$\mathcal{R} = \{r, s\}, \quad \Gamma = (r + \mathbb{R}_+^2) \cup (s + \mathbb{R}_+^2), \quad \mathcal{K} = \text{conv}\{r, s\} + \mathbb{R}_+^2.$$

The midpoint $(1, 1)$ belongs to $\mathcal{K}$ but not to $\mathcal{R}$ or Γ . Therefore

$$\inf_{x \in \mathcal{R}} \max\{x_1, x_2\} = 2, \quad \inf_{x \in \mathcal{K}} \max\{x_1, x_2\} = 1.$$

Convexification preserves every additive price but changes this monotone nonlinear decision. Coordinatewise minima are also fictitious: each minimum is zero, but no legal realization attains $(0, 0)$.

Retaining only r incurs regret under $w = (1, 3)$:

$$\text{Regret}_w(\{r\}) = \langle w, r \rangle - \min\{\langle w, r \rangle, \langle w, s \rangle\} = 6 - 2 = 4.$$

This explicit separator is the elementary form of the Safe-Commitment Law.

9.4 Technology as an observation of the geometry

A technology point does not alter a semantic obligation merely by pricing it. If the legal family and embedding capabilities are unchanged, a new price vector exposes another point of the same $\mathcal{K}$. If the technology introduces a new primitive or hierarchy, the resource contract changes and a new legal realization may appear outside the old geometry. The former operation is repricing; the latter is branch reopening.

10 Conclusion

Computation-realization geometry separates three questions that physical computing must answer in order: what the computation requires semantically, which realizations are legal under a resource contract, and which observation of the realization family is sufficient for a later decision. The registry $\mathcal{P}$ preserves realization identity and causal structure; $\mathcal{R}$ preserves exact joint attainability; Γ preserves every continuous monotone resource decision; and $\mathcal{K}$ preserves every additive nonnegative-price decision.

The Safe-Commitment Law turns this hierarchy into an exact design principle. A realization may be discarded precisely when the object required by the future decision class remains unchanged. The contextual-state, interactive linear, carrier, conservation, reconstruction, and composition theorems show that semantic obligation, attainable realization, and decision sufficiency can be expressed in one formal object.

CRG therefore supplies a foundational intermediate layer between abstract computation and physical technology. Communication-complexity converses, network regions, circuit constructions, memory-traffic laws, compiler transformations, and technology prices become compatible slices of a common realization geometry while retaining their distinct assumptions and units.

References

- [1] H. T. Kung, “The structure of parallel algorithms,” *Advances in Computers*, vol. 19, pp. 65–112, 1980.
- [2] R. T. Rockafellar, *Convex Analysis*. Princeton University Press, 1970.
- [3] M. Braverman, “Interactive information complexity,” *SIAM Journal on Computing*, vol. 44, no. 6, pp. 1698–1739, 2015.
- [4] B. Barak, M. Braverman, X. Chen, and A. Rao, “How to compress interactive communication,” *SIAM Journal on Computing*, vol. 42, no. 3, pp. 1327–1363, 2013.
- [5] J.-W. Hong and H. T. Kung, “I/O complexity: The red-blue pebble game,” in *Proceedings of the 13th ACM Symposium on Theory of Computing*, pp. 326–333, 1981.
- [6] G. Ballard, J. Demmel, O. Holtz, and O. Schwartz, “Minimizing communication in numerical linear algebra,” *SIAM Journal on Matrix Analysis and Applications*, vol. 32, no. 3, pp. 866–901, 2011.
- [7] M. Frigo, C. E. Leiserson, H. Prokop, and S. Ramachandran, “Cache-oblivious algorithms,” in *Proceedings of the 40th IEEE Symposium on Foundations of Computer Science*, pp. 285–297, 1999.
- [8] D. Irony, S. Toledo, and A. Tiskin, “Communication lower bounds for distributed-memory matrix multiplication,” *Journal of Parallel and Distributed Computing*, vol. 64, no. 9, pp. 1017–1026, 2004.
- [9] E. Solomonik and J. Demmel, “Communication-optimal parallel 2.5D matrix multiplication and LU factorization algorithms,” in *Euro-Par 2011*, pp. 90–109, 2011.
- [10] UCIE Consortium, “UCIE 3.0 Specification,” 2025. <https://www.uciexpress.org/specifications>.
- [11] UCIE Consortium, “UCIE 2.0 White Paper,” 2024, accessed through the official specification portal. <https://www.uciexpress.org/specifications>.
- [12] AMD, “4th Gen AMD EPYC Processor Architecture,” white paper, 2023. <https://www.amd.com/content/dam/amd/en/documents/products/epyc/4th-gen-epyc-processor-architecture-white-paper.pdf>.
- [13] R. A. van de Geijn and J. Watts, “SUMMA: Scalable universal matrix multiplication algorithm,” *Concurrency: Practice and Experience*, vol. 9, no. 4, pp. 255–274, 1997.
- [14] R. W. Hockney, “The communication challenge for MPP: Intel Paragon and Meiko CS-2,” *Parallel Computing*, vol. 20, no. 3, pp. 389–398, 1994.
- [15] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Ré, “FlashAttention: Fast and memory-efficient exact attention with IO-awareness,” in *Advances in Neural Information Processing Systems*, vol. 35, 2022, pp. 16344–16359. Available: <https://arxiv.org/abs/2205.14135>.
- [16] NVIDIA, “NVIDIA A100 Tensor Core GPU,” product datasheet and architecture white paper, 2020–2021. <https://www.nvidia.com/en-us/data-center/a100/>.
- [17] L. Zheng et al., “Alpa: Automating inter- and intra-operator parallelism for distributed deep learning,” in *OSDI*, pp. 559–578, 2022.
- [18] D. Lepikhin et al., “GSPMD: General and scalable parallelization for ML computation graphs,” arXiv:2105.04663, 2021. <https://arxiv.org/abs/2105.04663>.
- [19] E. Kushilevitz and N. Nisan, *Communication Complexity*. Cambridge University Press, 1997.
- [20] R. Ahlswede and N. Cai, “On communication complexity of vector-valued functions,” *IEEE Transactions on Information Theory*, vol. 40, no. 6, pp. 2062–2067, 1994.
- [21] H. Kowshik and P. R. Kumar, “Optimal function computation in directed and undirected graphs,” *IEEE Transactions on Information Theory*, vol. 58, no. 6, pp. 3407–3418, 2012.
- [22] R. Ahlswede, N. Cai, S.-Y. R. Li, and R. W. Yeung, “Network information flow,” *IEEE Transactions on Information Theory*, vol. 46, no. 4, pp. 1204–1216, 2000.

- [23] R. Dougherty, C. Freiling, and K. Zeger, “Networks, matroids, and non-Shannon information inequalities,” *IEEE Transactions on Information Theory*, vol. 53, no. 6, pp. 1949–1969, 2007.
- [24] C. T. Li, “Undecidability of network coding, conditional information inequalities, and conditional independence implication,” *IEEE Transactions on Information Theory*, vol. 69, no. 6, pp. 3493–3510, 2023.
- [25] C. D. Thompson, “Area-time complexity for VLSI,” in *Proceedings of the 11th ACM Symposium on Theory of Computing*, pp. 81–88, 1979.
- [26] A. M. Turing, “On computable numbers, with an application to the Entscheidungsproblem,” *Proceedings of the London Mathematical Society*, series 2, vol. 42, pp. 230–265, 1936.
- [27] R. Gandy, “Church’s thesis and principles for mechanisms,” in *The Kleene Symposium*, J. Barwise, H. J. Keisler, and K. Kunen, Eds. North-Holland, 1980, pp. 123–148, doi:10.1016/S0049-237X(08)71257-6.
- [28] D. Deutsch, “Quantum theory, the Church–Turing principle and the universal quantum computer,” *Proceedings of the Royal Society of London A*, vol. 400, no. 1818, pp. 97–117, 1985.
- [29] R. E. Kalman, “Mathematical description of linear dynamical systems,” *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, vol. 1, no. 2, pp. 152–192, 1963.
- [30] B. L. Ho and R. E. Kalman, “Effective construction of linear state-variable models from input/output functions,” *Regelungstechnik*, vol. 14, pp. 545–548, 1966.
- [31] A. Nerode, “Linear automaton transformations,” *Proceedings of the American Mathematical Society*, vol. 9, no. 4, pp. 541–544, 1958.
- [32] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, pp. 379–423 and 623–656, 1948.
- [33] A. C.-C. Yao, “Some complexity questions related to distributive computing,” in *Proceedings of the 11th ACM Symposium on Theory of Computing*, pp. 209–213, 1979.
- [34] R. Appuswamy, M. Franceschetti, N. Karamchandani, and K. Zeger, “Network coding for computing: Cut-set bounds,” *IEEE Transactions on Information Theory*, vol. 57, no. 2, pp. 1015–1030, 2011.
- [35] L. de Alfaro and T. A. Henzinger, “Interface automata,” in *Proceedings of the 8th European Software Engineering Conference*, pp. 109–120, 2001, doi:10.1145/503209.503226.
- [36] A. Benveniste, B. Caillaud, D. Nickovic, R. Passerone, J.-B. Raclet, P. Reinkemeier, A. Sangiovanni-Vincentelli, W. Damm, T. A. Henzinger, and K. G. Larsen, *Contracts for System Design*. Foundations and Trends in Electronic Design Automation, vol. 12, no. 2–3, pp. 124–400, 2018, doi:10.1561/10000000053.
- [37] P. Cousot and R. Cousot, “Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints,” in *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, 1977, pp. 238–252.
- [38] D. Blackwell, “Equivalent comparisons of experiments,” *The Annals of Mathematical Statistics*, vol. 24, no. 2, pp. 265–272, 1953.
- [39] H. P. Benson, “An outer approximation algorithm for generating all efficient extreme points in the outcome set of a multiple objective linear programming problem,” *Journal of Global Optimization*, vol. 13, pp. 1–24, 1998.
- [40] A. Löhne, *Vector Optimization with Infimum and Supremum*. Springer, 2011.
- [41] R. Bojanić and E. Seneta, “A unified theory of regularly varying sequences,” *Mathematische Zeitschrift*, vol. 134, pp. 91–106, 1973.
- [42] J. Galambos and E. Seneta, “Regularly varying sequences,” *Proceedings of the American Mathematical Society*, vol. 41, no. 1, pp. 110–116, 1973.
- [43] J. Hajnal, “Weak ergodicity in non-homogeneous Markov chains,” *Proceedings of the Cambridge Philosophical Society*, vol. 54, no. 2, pp. 233–246, 1958.

- [44] J. E. Cohen, “Ergodic theorems in demography,” *Bulletin of the American Mathematical Society*, vol. 1, no. 2, pp. 275–295, 1979.
- [45] E. Seneta, *Non-negative Matrices and Markov Chains*, revised ed. Springer, 2006.
- [46] D. J. Hartfiel, *Nonhomogeneous Matrix Products*. World Scientific, 2002.
- [47] S. Williams, A. Waterman, and D. Patterson, “Roofline: An insightful visual performance model for multicore architectures,” *Communications of the ACM*, vol. 52, no. 4, pp. 65–76, 2009.
- [48] J. Symons et al., “Stream: Layer-fused multi-core dataflow exploration for deep neural networks,” 2023.
- [49] Y. S. Shao et al., “Simba: Scaling deep-learning inference with multi-chip-module-based architecture,” in *Proceedings of MICRO*, 2019.
- [50] E. T. Kalns and L. M. Ni, “Processor mapping techniques toward efficient data redistribution,” *IEEE Transactions on Parallel and Distributed Systems*, 1995.
- [51] A. Schrijver, *Theory of Linear and Integer Programming*. Wiley, 1986.
- [52] R. M. Karp, “Reducibility among combinatorial problems,” in *Complexity of Computer Computations*, pp. 85–103, 1972.
- [53] J. H. Davenport and J. Heintz, “Real quantifier elimination is doubly exponential,” *Journal of Symbolic Computation*, vol. 5, pp. 29–35, 1988.

Collective attribution, accountability, and AI research provenance

Fitih M. Cinnor initiated, directed, and led the research program; established its scientific objectives and evidentiary standards; selected and adjudicated the claims advanced here; and reviewed the supporting mathematics and computation. He is the named and accountable human author, responsible for the accuracy, integrity, scope, and correction of the work. Accountable correspondence: fcinnor@semiaifoundry.com.

semiAIfoundry Research, a research group of Semi AI Foundry LLC, is the collective scholarly attribution for the coordinated research effort. It recognizes integrated intellectual work across human, artificial, and hybrid forms of intelligence rather than masking those contributions behind an exclusively individual account. Collective attribution does not displace human accountability or create AI legal authorship, ownership, signatory authority, representation authority, or accountable-person status.

The research used a human-directed, orchestrated multi-agent process. AI systems made material contributions to candidate generation, mathematical derivation, theorem formalization, counterexample construction, numerical implementation, software testing, high-precision and interval-oriented verification, literature mapping, prior-art comparison, research-record synthesis, adversarial critique, visualization, drafting, editing, and publication engineering. OpenAI GPT-5.6 Pro receives explicit credit for cross-corpus synthesis, mathematical exploration, derivation checking, theorem formalization, counterexample search, software and verifier development, exposition, consistency review, and manuscript construction. Anthropic Claude systems, including Fable 5, receive explicit credit for mathematical exploration, numerical-instrument development, adversarial review, gap identification, and editorial critique. The Fable 5 review was an internal AI-assisted audit, not independent external peer review or third-party validation.

These credits record provenance, not scientific status. AI-generated material acquires no standing merely by being generated. Scientific standing comes from proofs, computations, verification controls, independent reproductions, source checks, and correction records. No AI system is an author, legal owner, signatory, authorized representative, or accountable person. No identification or credit implies that OpenAI, Anthropic, EPFL, an artifact source, or any other provider or institution owns, sponsors, approves, warrants, or endorses this work, or accepts responsibility for its claims. Third-party rights remain governed by the applicable notices. The research program received no external funding. The accountable human author declares no competing interests or conflicts of interest.

Copyright, licensing, and artifact availability

Copyright © 2026 Semi AI Foundry LLC. Except where otherwise noted, the manuscript and project-created figures in the exact versioned public payload are licensed under CC BY 4.0 (<https://creativecommons.org/licenses/by/4.0/>), with requested attribution to “Fitih M. Cinnor for semiAIfoundry Research.” Reusers must give appropriate attribution, provide the license link, and indicate changes; attribution does not imply endorsement. Project-created code is licensed under Apache-2.0. Third-party materials remain governed by their recorded upstream terms. The licensed payload is identified by its versioned source, results, artifacts, version number, public file index, root SHA-256 manifest, and archival deposit. These grants apply only to the files and version identified in that payload and do not extend to other releases or evolving research-program content.