FLYXION

Beyond Process

Constraint, Accessibility, and the Geometry of Becoming

An essay in constraint-first ontology

June 2026

Abstract

The history of philosophy frames the question of what is most real as a choice between two answers: persistent objects (substance ontology) or flowing events (process ontology). This essay proposes that neither is correct—that both objects and processes are derived from a deeper stratum which is neither a thing nor an event, but a geometry: the structure of constraints determining which configurations are accessible from which. We call this an *accessibility ontology* or *constraint-first ontology*.

The essay develops this claim across fourteen sections. The core mathematical apparatus is a projection framework $\pi : X \to M$ from a rich trajectory space to an observable manifold, together with a constraint field $\mathcal{A} : \mathcal{C} \to \mathcal{P}(\mathcal{C})$ and an entropy measure $H(\pi, m) = \ln \mu(\pi^{-1}(m))$ that quantifies information loss under projection. Objects emerge as basins invariant under $\mathcal{A}$; processes emerge as trajectories through it; entropy growth is the increase in fiber degeneracy when dynamics operate entirely on the projected manifold.

Applications are developed to computation (Spherepop as containment-based evaluation), cognition (gesture as the primary substrate of symbolic meaning), inference (admissibility theory and CLIO), cosmology (RSVP as an accessibility field theory), coordination (stigmergy unifying Ising models, cellular automata, and ant colonies), and artificial cognitive architecture (MEM|8 as a partial machine realization of the theory). A six-level implementation tower is derived, from ontology through inference operator, evaluation strategy, memory model, field dynamics, and wave-grid runtime, showing that each level is the unique tractable realization of the level above it.

The governing slogan replacing Heraclitus's *"everything flows"* is: *everything navigates a changing landscape of admissible possibilities.*

Keywords: constraint-first ontology, accessibility geometry, projection entropy, admissibility theory, Spherepop, RSVP, CLIO, MEM|8, process philosophy, stigmergy, cognitive runtime.

Summary of Principal Results

The essay establishes the following hierarchy of concepts and results.

Ontological primitives. The constraint field $\mathcal{A} : \mathcal{C} \to \mathcal{P}(\mathcal{C})$ is ontologically prior to both trajectories and objects. Objects are $\mathcal{A}$-invariant basins; processes are trajectories through $\mathcal{A}$; the accessibility relation is neither a thing nor an event. This creates a three-level hierarchy—substance ontology, process ontology, accessibility ontology—in which each level derives the primitives of the level above from those of the level below.

Projection and information loss. Any surjective map $\pi : X \to M$ with $|X| > |M|$ has nontrivial fibers (Map-Territory Theorem). Every map, model, symbol, memory, or institution is such a projection. The information lost in projection is the fiber coordinate, invisible to any observer operating entirely within M.

Entropy as degeneracy. Projection entropy $H(\pi, m) = \ln \mu(\pi^{-1}(m))$ generalizes Boltzmann entropy; Boltzmann's $S = k_B \ln \Omega$ is the special case where π is the thermodynamic coarse-graining map. The Second Law follows: dynamics indifferent to fiber coordinates cannot track degeneracy, so expected entropy is non-decreasing. Forgetting is fiber expansion; remembering is preservation of discriminating constraints.

Spherepop and gesture. Evaluation order in any calculus is containment depth (innermost resolves first). Syntactic punctuation is the residue of projecting containment structures onto strings. Symbols are compressed projections of gestural trajectories; understanding is fiber navigation, not lookup.

Stigmergy. Coordination through environmental traces unifies Ising synchronization, cellular automaton dynamics, and ant colony behavior under a single schema: local traces $\to$ modified constraints $\to$ altered accessibility $\to$ coherent structure. Persistent structures (gliders, magnetized phases, trails) are stigmergic attractors.

MEM | 8 as cognitive runtime. Wave-based memory architecture [10] is a partial machine realization of the constraint-first theory. The key correspondences are:

Theory primitive	Runtime implementation
Constraint field $\mathcal{A}$	Wave grid $(256^2 \times 65536)$
Overlapping constraint fields	Wave interference (unique by Thm. 12.3)
Constraint relaxation	Exponential decay (entropy-optimal by Thm. 12.6)
Admissibility stack	Reactive layer hierarchy (Thm. 12.8)
Curvature control	Custodian operator (Thm. 12.10)
Semantic projection	`.m8` format (Prop. 12.11)
Phase-locking / binding	Kuramoto synchronization (R = binding strength)

The Kuramoto coherence parameter R provides a quantitative measure of constraint binding strength: when two constraint perturbations are causally or temporally linked, they phase-lock, and their coherence $R \in [0, 1]$ represents the strength of the composite constraint. $R = 1$ indicates full constraint reinforcement; $R = 0$ indicates independent (phase-orthogonal) constraints.

The implementation tower. The six-level tower (Semantic Infrastructure $\to$ CLIO $\to$ Spherepop $\to$ Chain of Memory $\to$ RSVP $\to$ MEM | 8) is a deductive structure: each level is the unique tractable realization of the level above under the constraints of termination, continuity, and minimality.

The master theorem. For any system described by $(\mathcal{C}, \mathcal{A}, \pi)$: stable structures are basins of $\mathcal{A}$; processes are trajectories through $\mathcal{A}$; information loss is projection entropy; coordination is stigmergic modification of $\mathcal{A}$; inference is admissibility-constrained trajectory selection. The governing slogan: *everything navigates a changing landscape of admissible possibilities.*

Contents

1 Introduction: The Debate Between Being and Becoming

The history of philosophy can be read as a prolonged dispute between two intuitions about what is most real. The first intuition is that reality is composed of persistent things—objects, substances, states—and that change is a secondary phenomenon that happens to those things. The second intuition is that process, flux, and becoming are primary, and that stable objects are derivative structures precipitated from an underlying flow.

This opposition traces at least to the pre-Socratics. Parmenides argued that being is one, unchanging, and complete; motion and multiplicity are illusions. Heraclitus held the opposite: the river is never the same river twice, fire is the fundamental stuff, and *strife*—the ceaseless tension between opposites—is the basis of all order. Alfred North Whitehead, in the twentieth century, transformed Heraclitean process intuition into a systematic metaphysics. In *Process and Reality* (1929) he argued that the basic constituents of reality are not enduring substances but *occasions of experience*: momentary events that prehend their predecessors and pass into their successors. Objects are *societies* of occasions—patterns that persist by reproducing themselves through time.

The present essay argues that both positions share an unexamined assumption. Whether one takes objects or processes as fundamental, one is still asking *what there is*. One answer is "things"; the other is "events". But there is a third possibility: neither things nor events are fundamental. What is fundamental is the *geometry of possibility*—the structure of constraints determining which trajectories, which changes, which continuations are admissible at all. Objects emerge as basins within that geometry. Processes emerge as trajectories through it. The geometry itself—the accessibility landscape—is what underlies both.

The essay proceeds as follows. Section 2 states the ontological hierarchy precisely and derives its consequences. Section 3 develops the central mathematical apparatus: projection, fiber, and information loss. Section 4 reinterprets entropy as a measure of degeneracy under projection. Section 5 examines Spherepop as a containment-based calculus in which evaluation is geometric collapse. Section 6 treats gesture as the primary substrate of symbolic computation. Section 7 introduces admissibility theory and CLIO as a generalized inference framework. Section 8 situates the RSVP field equations within the accessibility framework. Section 9 unifies stigmergy, Ising synchronization, and cellular automata under a common constraint-modification schema. Section 10 places the resulting ontology in relation to process philosophy and identifies where it diverges. Section 13 gathers the threads into a master theorem.

Throughout, formal definitions, theorems, and proofs are given at each step. The goal is not merely to assert the framework but to derive its structure from minimal commitments.

2 The Ontological Hierarchy

We begin by making the hierarchical claim precise. Three frameworks are distinguished by where they locate the fundamental stratum of reality.

Definition 2.1 (Substance Ontology)**.** A *substance ontology* is one in which the primitive entities are persistent objects $x \in X$ equipped with a state space $S(x)$. Change is a function $\phi : X \times T \to S$, where T is a time parameter.

Definition 2.2 (Process Ontology)**.** A *process ontology* is one in which the primitive entities are events or occasions $e \in E$. An object is defined as a temporally extended pattern: a sequence $(e_1, e_2, \ldots)$ of occasions bound by causal prehension relations $e_i \prec e_j$ (e_i is causally prior to e_j).

Definition 2.3 (Accessibility Ontology)**.** An *accessibility ontology* is one in which the primitive entities are *accessibility relations*. Let $\mathcal{C}$ be a space of configurations. A *constraint field* is a function

$$\mathcal{A} : \mathcal{C} \to \mathcal{P}(\mathcal{C})$$

assigning to each configuration c the set $\mathcal{A}(c)$ of configurations that are *immediately accessible* from c. A *trajectory* is a sequence $(c_0, c_1, c_2, \ldots)$ such that $c_{i+1} \in \mathcal{A}(c_i)$ for all i. An *object* is an $\mathcal{A}$-invariant basin: a set $B \subseteq \mathcal{C}$ such that for all $c \in B$, $\mathcal{A}(c) \cap B \neq \emptyset$.

Proposition 2.4 (Ontological Depth)**.** *In any dynamical system admitting a constraint field $\mathcal{A}$:*

- *(i) Every trajectory is determined by $\mathcal{A}$ and an initial condition; but $\mathcal{A}$ is not determined by any single trajectory.*
- *(ii) Every object (basin) is determined by $\mathcal{A}$; but $\mathcal{A}$ is not determined by any collection of objects.*
- *(iii) $\mathcal{A}$ is therefore epistemically and ontologically prior to both trajectories and objects.*

Proof. (i) Given $\mathcal{A}$ and c_0, the set of admissible continuations at each step is $\mathcal{A}(c_i)$; thus every trajectory is a selection from this field. A single trajectory γ specifies only those transitions $c_i \to c_{i+1}$ that actually occurred; the full relation $\mathcal{A}$ includes all possible transitions from every $c \in \mathcal{C}$, which no single trajectory can witness.

(ii) A basin B satisfies $B = \{c : \mathcal{A}(c) \cap B \neq \emptyset\}$. Multiple distinct constraint fields can produce identical basins (adding transitions *between* basins does not alter the basins themselves), so $\mathcal{A}$ is not recoverable from the basins alone.

(iii) Both trajectories and objects are *derived* from $\mathcal{A}$; neither determines it. Hence $\mathcal{A}$ occupies the deeper stratum. □ □

Remark 2.5. The argument is deliberately minimal. It does not require metric structure, differentiability, or a Lagrangian. It holds for any relational system in which "what is possible next" is well defined.

3 Projection, Fibers, and Information Loss

The central mathematical tool is the theory of projections and their fibers. Let X be a "rich" space (full trajectory space, gesture space, field configuration space) and M a "coarser" observable space. A *map* is any function $\pi : X \to M$.

Definition 3.1 (Fiber)**.** For $m \in M$, the *fiber* of π over m is

$$F_m^\pi = \pi^{-1}(m) = \{x \in X : \pi(x) = m\}.$$

Definition 3.2 (Trivial and Nontrivial Projections)**.** π is *trivially informative* if $|F_m^\pi| = 1$ for all $m \in M$ (the projection is injective). It has *nontrivial fibers* if there exists $m \in M$ with $|F_m^\pi| > 1$.

Theorem 3.3 (Map–Territory Projection Theorem)**.** *Let $\pi : X \to M$ be any surjective map with $|X| > |M|$. Then π has nontrivial fibers: there exist distinct $x_1, x_2 \in X$ with $\pi(x_1) = \pi(x_2)$, and the information distinguishing them is invisible in M.*

Proof. Since $|X| > |M|$ and π is surjective, the pigeonhole principle forces at least one $m \in M$ to satisfy $|F_m^\pi| \geq 2$. Any two elements $x_1 \neq x_2$ in that fiber map to the same point m, so they are indistinguishable to any observer with access only to M. The information distinguishing them—their fiber coordinate—lies in $X \setminus M$ and cannot be recovered from m alone. □ □

Table 1 catalogs the recurring instances of this schema throughout the framework.

Rich space X	**Projection**	**Observable** M
Gesture trajectory	compression	Symbol
Expression tree	evaluation	Value
Territory	map-making	Map
Accessibility landscape	stabilization	Object / state
Full utterance history	canonicalization	Semantic field
Civilizational trajectory	institutionalization	Institution
Field configuration (RSVP)	coarse-graining	Cosmological observable
Colony dynamics	pheromone deposition	Stigmergic trace

Table 1: Instances of the projection schema $X \xrightarrow{\pi} M$.

Corollary 3.4. *Full reconstruction of $x \in X$ from $m = \pi(x)$ requires additional information specifying the fiber coordinate of x within F_m^π. In general, this coordinate is unavailable to an observer operating entirely within M.*

Proof. From m alone, an observer can identify F_m^π but cannot distinguish among its elements. The fiber coordinate parameterizes the kernel of π; since this kernel lies outside M, no map $M \to X$ can reconstruct it without auxiliary information. □ □

4 Entropy as Degeneracy Under Projection

Standard thermodynamic entropy is defined by Boltzmann as $S = k_B \ln \Omega$, where Ω is the number of microstates consistent with a given macrostate. We show this is a special case of a more general notion: the logarithmic size of a fiber.

Definition 4.1 (Projection Entropy)**.** Let $\pi : X \to M$ be a projection and $m \in M$ an observable. The *projection entropy* of m is

$$H(\pi, m) = \ln |F_m^\pi|.$$

When X carries a measure μ and fibers need not be finite, set $H(\pi, m) = \ln \mu(F_m^\pi)$.

Proposition 4.2 (Boltzmann Entropy as Special Case)**.** *Thermodynamic entropy $S = k_B \ln \Omega$ is projection entropy with π the coarse-graining map from microstates to macrostates, X the microstate space, and k_B a unit conversion factor.*

Proof. Set $X =$ (microstate space), $M =$ (macrostate space), and $\pi(x) =$ the macrostate of microstate x. Then $F_m^\pi = \{x : \pi(x) = m\}$ is the set of microstates compatible with macrostate m, which is Ω. Hence $H(\pi, m) = \ln |F_m^\pi| = \ln \Omega = S/k_B$. □ □

Proposition 4.3 (Second Law as Increasing Degeneracy)**.** *In a system whose dynamics operates entirely on M (i.e. is indifferent to fiber coordinates), the expected projection entropy $\langle H \rangle$ is non-decreasing over time.*

Proof. If the dynamics $\Phi_t : M \to M$ does not track fiber coordinates, then an initial microstate x_0 with $\pi(x_0) = m_0$ evolves under any trajectory consistent with the macrostate path $m_t = \Phi_t(m_0)$. At each step, the set of microstates consistent with m_t can only grow or stay constant, because the dynamics cannot eliminate fiber elements it does not track. Formally, $|F_{m_t}^\pi| \geq |F_{m_0}^\pi|$ on average, so $\langle H(\pi, m_t) \rangle \geq H(\pi, m_0)$. □ □

Remark 4.4. The Second Law is not a brute fact about disorder but a structural consequence of operating with projections. Any system that compresses information—maps, memories, concepts, institutions—loses the ability to reverse its own coarse-graining. Irreversibility is in the projection, not in any intrinsic asymmetry of the underlying dynamics.

Example 4.5 (Semantic Entropy)**.** Let X be the space of utterance histories leading to a conversational state, and M the space of semantic representations. Projection entropy measures how many distinct histories could have produced a given meaning m. High semantic entropy means a meaning is accessible via many routes; low entropy means the path that produced it is nearly recoverable from the meaning alone.

Example 4.6 (Institutional Entropy)**.** Let X be the space of civilizational trajectories and M the space of institutional configurations. As institutions become more self-reproducing, their fiber grows: the same institutional configuration becomes reachable from more and more distinct historical pasts.

5 Spherepop: Evaluation as Geometric Collapse

Spherepop is a containment-based calculus in which expressions are represented as nested regions and evaluation is the collapse of those regions. Syntactic structure (scope, precedence, binding) is encoded geometrically as containment depth.

Definition 5.1 (Containment Structure)**.** A *containment structure* is a pair $(R, \subseteq)$ where R is a finite set of *regions* and $\subseteq$ is a partial order representing containment, satisfying: if $r_1 \cap r_2 \neq \emptyset$ then either $r_1 \subseteq r_2$ or $r_2 \subseteq r_1$ (regions are strictly nested).

Definition 5.2 (Spherepop Expression)**.** A *Spherepop expression* is a labeled containment structure $(R, \subseteq, \ell)$ where $\ell : R \to \Sigma$ assigns each region a label from an alphabet Σ. The *depth* of region r is

$$d(r) = |\{r' \in R : r' \supsetneq r\}|.$$

Definition 5.3 (Pop, Refuse, Collapse)**.** The three primitive operations are:

(i) **Pop**: Remove the innermost region r^* (one with no proper sub-regions) and return its evaluated value to its parent.
(ii) **Refuse**: Mark a region as irreducible until its context admits it.
(iii) **Collapse**: Apply Pop exhaustively to all innermost regions, yielding the fully evaluated normal form.

Theorem 5.4 (Containment Encodes Precedence)**.** *In a Spherepop expression representing an arithmetic formula, the evaluation order induced by iterated Pop (innermost first) agrees with the standard evaluation order defined by operator precedence.*

Proof. Standard precedence rules encode binding tightness: sub-expressions that bind more tightly are placed in regions at greater depth. Since Pop always collapses the deepest region first, it evaluates the most tightly bound sub-expression first. Formally: let e contain operators $\circ_1$ (higher precedence) and $\circ_2$ (lower precedence). Encoding e as a Spherepop structure places the $\circ_1$ sub-expression at depth $d + 1$ and the $\circ_2$ sub-expression at depth d. Iterated Pop reduces the depth-$(d + 1)$ region before the depth-d region, matching the precedence-respecting evaluation order. □ □

Corollary 5.5. *Scope, variable binding, and recursive call structure are similarly encoded by containment depth. Parentheses, keywords, and indentation in conventional string-based syntax are the residue of applying the projection π : containment structures $\to$ strings: they are the additional notation required to partially recover the fiber coordinate—i.e. to reconstruct the containment structure from its linearized image.*

Proof. By Theorem 3.3, the linearization map π has nontrivial fibers: multiple distinct containment structures map to similar strings. Syntactic punctuation (parentheses, keywords) is exactly the notation that a writer adds to reduce the fiber, allowing a reader to reconstruct the intended containment structure from the string. □ □

6 Gesture as Primary Substrate

Symbols are compressed projections of gestural trajectories. Gesture is not a supplement to symbolic computation but its original substrate.

Definition 6.1 (Gesture Space)**.** A *gesture space* is a metric space $(\mathcal{G}, d_{\mathcal{G}})$ of continuous trajectories $\gamma : [0,1] \to \mathbb{R}^n$ (for embodied motor gestures, n is the dimension of the effector space). The group of reparameterizations $\mathrm{Diff}^+([0,1])$ acts on $\mathcal{G}$ by precomposition; the *shape* of a gesture is its equivalence class under this action.

Definition 6.2 (Symbol as Gesture Projection)**.** A *symbol* $s \in \Sigma$ is an equivalence class of gestures under a projection $\pi_{\mathrm{sym}} : \mathcal{G} \to \Sigma$ that identifies all gestures a community conventionally interprets as s. The fiber $F_s^{\pi_{\mathrm{sym}}}$ is the set of gestures conventionally producing s.

Proposition 6.3 (Semantic Reconstruction is Fiber Navigation)**.** *Understanding a symbol s requires navigating from $s \in \Sigma$ back into a representative gesture in $F_s^{\pi_{\mathrm{sym}}}$. This navigation is not uniquely determined by s alone; it requires contextual constraints to select a specific fiber element.*

Proof. By Corollary 3.4, $s \in \Sigma$ corresponds to a fiber with potentially many elements. A receiver interpreting s must select one $\hat{\gamma} \in F_s^{\pi_{\mathrm{sym}}}$ as the "intended" gesture. This selection is underdetermined by s alone; it depends on prior context, shared conventions, and the receiver's constraint field $\mathcal{A}$. Understanding is therefore constrained trajectory reconstruction, not symbol lookup. □ □

Remark 6.4. This formalizes the underdetermination of meaning by linguistic form as a structural consequence of the projection schema. Meaning is more stable in high-constraint contexts (where $\mathcal{A}$ strongly selects a single fiber element) and more ambiguous in low-constraint contexts.

7 Admissibility Theory and CLIO

Admissibility theory generalizes accessibility from dynamical systems to inference problems. CLIO (Constrained Landscape of Inferential Openings) is the associated formal framework.

Definition 7.1 (Admissibility Space)**.** An *admissibility space* is a triple $(\mathcal{Q}, \mathcal{D}, \mathrm{Adm})$ where $\mathcal{Q}$ is a set of queries, $\mathcal{D}$ is a document (a structured set of propositions), and $\mathrm{Adm} : \mathcal{D} \to \mathcal{P}(\mathcal{Q})$ assigns to each document state the set of queries meaningfully constrained by it.

Definition 7.2 (CLIO Projection Engine)**.** A *CLIO projection engine* is a map $\Pi : \mathcal{D} \times \mathcal{Q} \to \mathcal{A} \cup \{\perp\}$ where $\mathcal{A}$ is an answer space, with $\Pi(\mathcal{D}, q)$ defined for $q \in \mathrm{Adm}(\mathcal{D})$ and $\Pi(\mathcal{D}, q) = \perp$ for inadmissible queries.

Theorem 7.3 (Admissibility Determines Semantic Coherence)**.** *The document $\mathcal{D}$ is* semantically coherent *(consistently supports its admissible queries) if and only if* Adm *satisfies:*

(i) Closure under entailment*: if $q_1 \in \mathrm{Adm}(\mathcal{D})$ and $q_1 \vdash q_2$, then $q_2 \in \mathrm{Adm}(\mathcal{D})$.*

(ii) Projection consistency*: for every $q \in \mathrm{Adm}(\mathcal{D})$, re-querying a document constructed from the answer yields the same answer:*

$$\Pi(\mathcal{D}, q) = \Pi\big(\sigma(\Pi(\mathcal{D}, q)), q\big)$$

for some section $\sigma : \mathcal{A} \to \mathcal{D}$.

Proof. (Necessity.) If closure under entailment fails, there exists $q_1 \in \mathrm{Adm}(\mathcal{D})$, $q_1 \vdash q_2$, with $q_2 \notin \mathrm{Adm}(\mathcal{D})$. Answering q_1 generates a context in which q_2 is natural but inadmissible, creating a gap in the document's inferential structure—incoherence. If projection consistency fails, an answer regenerates a document from which the same question yields a different answer—inconsistency.

(Sufficiency.) Given both conditions, every answer lies within the admissible region, and every answer is stable under re-projection. The admissible queries form a closed inferential network over $\mathcal{D}$, which is semantic coherence. □ □

Remark 7.4. CLIO generalizes both deductive closure (formal logic) and consistency conditions (databases) into a single framework parameterized by an accessibility structure over queries. The key shift is treating "what can be meaningfully asked" as the primary object, rather than "what is true."

8 RSVP: Cosmology as Accessibility Field Theory

The Relativistic Scalar-Vector Plenum (RSVP) framework situates cosmological dynamics within the accessibility ontology. Rather than beginning with a metric $g_{\mu\nu}$ and deriving dynamics from a Lagrangian, RSVP begins with an accessibility field and derives the metric as a derived quantity.

Definition 8.1 (RSVP Field Configuration)**.** An *RSVP field configuration* on a spacetime manifold $\mathcal{M}$ is a pair (ϕ, v_μ) where:

- (i) $\phi : \mathcal{M} \to \mathbb{R}$ is a scalar *accessibility potential*, measuring the local density of admissible transitions.
- (ii) $v_\mu : \mathcal{M} \to T^*\mathcal{M}$ is a co-vector field encoding the preferred direction of admissible flow (the plenum current).

Proposition 8.2 (RSVP Equations of Motion)**.** *The stationary-action principle for the Lagrangian density*

$$\mathcal{L}_{\mathrm{RSVP}} = \tfrac{1}{2}(\partial_\mu \phi)(\partial^\mu \phi) - \tfrac{1}{4} F_{\mu\nu} F^{\mu\nu} - V(\phi), \qquad F_{\mu\nu} = \partial_\mu v_\nu - \partial_\nu v_\mu,$$

yields the equations of motion

$$\Box \phi = V'(\phi), \tag{1}$$

$$\partial_\mu F^{\mu\nu} = J^\nu, \tag{2}$$

where J^ν is a source current derived from ϕ.

Proof. Euler–Lagrange variation of the scalar sector: $\partial_{\mathcal{L}}/\partial\phi = -V'(\phi)$ and $\partial_{\mathcal{L}}/\partial(\partial_\mu\phi) = \partial^\mu\phi$, giving $\partial_\mu\partial^\mu\phi + V'(\phi) = 0$, i.e. equation (1). For the vector sector, the $F_{\mu\nu}$ term is the Lagrangian of an Abelian gauge field; variation with respect to v_ν yields equation (2) by the standard Maxwell derivation, with J^ν arising from any ϕ-dependent interaction term. □ □

Theorem 8.3 (Cosmological Redshift as Accessibility Gradient)**.** *In an RSVP cosmology with decreasing accessibility potential $\partial_t\phi < 0$ on large scales, photons propagating along accessibility gradients undergo a frequency shift*

$$\frac{\delta\omega}{\omega} \propto \frac{\partial_t\phi}{\phi},$$

which reproduces the Hubble redshift as a consequence of the evolving constraint field rather than spatial expansion.

Proof sketch. The dispersion relation for a photon wavepacket in a background field ϕ is modified: $\omega^2 = |\mathbf{k}|^2 + m_{\text{eff}}^2(\phi)$ for a ϕ-dependent effective mass. As ϕ decreases, m_{eff} changes, and the frequency observed by a distant detector differs from that at emission. A first-order expansion in $\partial_t\phi/\phi$ yields the stated proportionality. The derivation parallels that of the gravitational redshift in standard GR, with the accessibility gradient playing the role of the gravitational potential gradient. □ □

Remark 8.4. RSVP is not proposed here as empirically established cosmology. Its function in this essay is structural: to show that the accessibility-field formalism, developed for purely foundational reasons, admits a cosmological sector as an application. The speculative content of RSVP is acknowledged; the formal embedding within the accessibility framework is not.

9 Stigmergy, Ising Models, and Cellular Automata

Stigmergy is coordination through environmental traces. An agent does not transmit its state to other agents; it modifies a shared medium, and that modification alters the constraints experienced by future agents. We show that Ising synchronization and cellular automata are special cases of this general mechanism.

Definition 9.1 (Stigmergic System)**.** A *stigmergic system* is a quadruple $(\mathcal{S}, \mathcal{E}, \tau, \mathcal{A}_{\mathcal{E}})$ where: $\mathcal{S}$ is a set of agents; $\mathcal{E}$ is the shared environment; $\tau : \mathcal{S} \times \mathcal{E} \to \mathcal{E}$ is the *trace function* (agent i in state s_i modifies e to $\tau(s_i, e)$); and $\mathcal{A}_{\mathcal{E}} : \mathcal{E} \to \mathcal{P}(\mathcal{S})$ is the *environmental constraint function* (the environment state e determines which agent transitions are admissible). The update cycle is: agents leave traces $\tau \to e$ updates $\to \mathcal{A}_{\mathcal{E}}(e)$ changes $\to$ future agent transitions are drawn from $\mathcal{A}_{\mathcal{E}}(e)$.

Theorem 9.2 (Ising Model as Stigmergic System)**.** *The Ising model on a lattice Λ is a stigmergic system in which the local magnetic field is the shared environment.*

Proof. Set agents $\mathcal{S} = \{\sigma_i : i \in \Lambda\}$ (spins with values in $\{-1, +1\}$), environment $\mathcal{E} =$ (local field configurations $\{h_i\}$). Define the trace function:

$$\tau(\sigma_i, h)_j = h_j + J_{ij}\sigma_i \quad \text{for neighbors } j \text{ of } i,$$

so spin σ_i modifies the field experienced by its neighbors. Define the constraint function:

$$\mathcal{A}_{\mathcal{E}}(h) \ni \sigma_i \text{ with probability } P(\sigma_i \to -\sigma_i \mid h_i) = \frac{1}{1 + e^{2\beta\sigma_i h_i}},$$

so the local field determines the admissibility of spin flips. The Ising dynamics—spins update according to local fields; fields are updated by spin states—is exactly the stigmergic update cycle. Coherent magnetized phases are stigmergic fixed points of this cycle. □ □

Theorem 9.3 (Cellular Automata as Stigmergic Systems)**.** *Every deterministic cellular automaton is a stigmergic system in which the cell grid is the shared environment.*

Proof. Let Λ be a lattice, Q a finite state set, and $f : Q^{N(i)} \to Q$ the local transition rule. Set agents = cells, environment = grid state $e \in Q^\Lambda$. Define $\tau(c_i, e) = e$ with $e_i \mapsto f(e_{N(i)})$: each cell writes its new state to the grid based on its neighborhood. Define $\mathcal{A}_{\mathcal{E}}(e) = f$ (the rule is uniform): the current grid state determines exactly which next states are admissible. Every cell state is a trace in the shared medium; those traces determine the admissible next steps for neighboring cells. □ □

Corollary 9.4 (Persistent Structures as Stigmergic Attractors)**.** *A persistent structure in a cellular automaton (e.g. a Life glider) is a* stigmergic attractor*: a finite region of cell states that, through its traces, continuously recreates the constraint field required for its own persistence.*

Proof. A glider is a pattern $P \subset Q^\Lambda$ that recurs (shifted by $\mathbf{d}$) after k steps. Each cell in P leaves a trace that, via $\mathcal{A}_{\mathcal{E}}$, determines the admissible next states of its neighbors, which in turn reproduce P shifted by $\mathbf{d}$. Thus P is a fixed point of the stigmergic dynamics modulo translation: a trace-constraint loop that sustains itself indefinitely. □ □

Proposition 9.5 (Unification Principle)**.** *Ant colony foraging, Ising criticality, cellular automaton gliders, semantic memory traces, and institutional path-dependence are all instances of the stigmergic schema:*

local traces $\longrightarrow$ *modified constraints* $\longrightarrow$ *altered accessibility* $\longrightarrow$ *coherent structure.*

Proof. Each system fits Definition of stigmergic system by the correspondences:

Ant colony: pheromone field = environment; deposition = trace function τ; evaporation and reinforcement threshold = $\mathcal{A}_{\mathcal{E}}$.

Semantic memory: prior conversational context = environment; utterance = trace; admissibility landscape of continuations = $\mathcal{A}_{\mathcal{E}}$.

Institution: legal precedent = trace; social and legal admissibility = $\mathcal{A}_{\mathcal{E}}$; institutional stability = stigmergic fixed point.

In every case the update cycle holds. □ □

10 Beyond Process Philosophy

The framework developed above stands in a determinate relationship to the process philosophy tradition. It agrees with process philosophy in rejecting substance primacy; it departs from process philosophy in its identification of the fundamental stratum.

Proposition 10.1 (Comparison with Whitehead). *Let W denote Whitehead's process ontology and A the accessibility ontology of Definition 2.3. Then:*

(i) A is strictly more general than W: every Whiteheadian system can be embedded in an accessibility system, but not conversely.
(ii) A is constraint-first *where W is* event-first*: in W, the occasion of experience is primitive and constraints emerge from causal relations; in A, the constraint field is primitive and occasions emerge as trajectories selected from it.*

Proof. (i) In Whitehead, each occasion e prehends a set $\mathrm{Pr}(e)$ of prior occasions. Define $\mathcal{A}(e) = \{e' : e \in \mathrm{Pr}(e')\}$ (the occasions causally influenced by e). This embeds W into A. Conversely, A admits constraint fields over continuous manifolds with no discrete occasion structure, which W does not represent. The inclusion is strict.

(ii) In W, causal relations between occasions define what is possible. In A, the accessibility relation $\mathcal{A}$ is defined independently of any occasions; occasions are derived as trajectories through $\mathcal{A}$. The order of derivation is reversed. □ □

The three frameworks are positioned as follows:

Framework	**Primitive**	**Derived**
Substance ontology	Objects (substances)	Change, process
Process ontology (Whitehead)	Events (occasions)	Objects, substances
Accessibility ontology	Constraint field $\mathcal{A}$	Trajectories, events, objects

The process philosopher asks: *what is happening?* The accessibility theorist asks: *what can happen from here?* The former studies motion; the latter studies the geometry of possibility that gives rise to motion. For the process philosopher, the river is a flow. For the accessibility theorist, the river is an accessibility landscape whose admissible trajectories manifest as a flow.

Remark 10.2 (Relation to Structural Realism). One might ask whether the regress terminates. Is there something more primitive than the constraint field? The answer implicit in the framework is that the constraint field is formal: it is a relation, not a substance or an event. Relations can be specified without positing any underlying "stuff." In this sense the accessibility ontology is closer to structural realism (Ladyman, French) than to either substance or process realism.

11 MEM|8: Cognitive Runtime for Constraint-First Theory

The previous sections established a constraint-first ontology, a projection framework, an entropy theory, a containment calculus, a gesture semantics, an inference operator, a field theory, and a stigmergic account of coordination. This section addresses the question: what does this theory look like when *compiled into an executable memory architecture*?

The answer proposed here is that MEM|8 [10]—a wave-based cognitive architecture by Chenoweth and Chenoweth—is not merely *compatible with* the constraint-first framework. It is a *partial machine realization of it*. The relationship is analogous to that between a programming language and a processor architecture: not identity, but a concrete execution model for the language's primitives. The wave grid executes constraint modifications. Interference executes constraint composition. Temporal blankets execute entropy management. Reactive layers execute the admissibility stack. The Custodian executes curvature control. The `.m8` format executes file-level semantic projection.

Each claim requires a proof stronger than "waves can encode constraints." The following subsections establish: (i) that wave interference is the *natural* representation of overlapping constraint fields, not merely a possible one; (ii) that temporal decay implements entropy increase under weakening reconstruction constraints in a precise variational sense; (iii) that the reactive layer hierarchy is a stratified admissibility stack derivable from urgency-ordered constraint evaluation; (iv) that the Custodian is a curvature-control operator on the accessibility manifold; and (v) that the `.m8` format is a file-level implementation of semantic projection under the fiber theorem.

Definition 11.1 (Chain of Memory)**.** A *chain of memory* is a sequence $X_0 \xrightarrow{\pi_0} M_0 \xrightarrow{R_0} X_1 \xrightarrow{\pi_1} M_1 \xrightarrow{R_1} \cdots$ of projections and reconstructions such that each X_{n+1} preserves enough constraints from X_n to maintain identity: $|\mathcal{C}(X_n) \cap \mathcal{C}(X_{n+1})| \geq \theta$ for a coherence threshold θ.

Proposition 11.2 (Memory Effectiveness Criterion)**.** *A memory $\mathbf{m} : \mathcal{C} \to \mathcal{C}$ is cognitively inert if and only if for every constraint c and every trajectory γ, $\gamma \in \mathcal{A}(c) \iff \gamma \in \mathcal{A}(\mathbf{m}(c))$.*

Proof. If no admissible region changes, no future behavior is altered; the memory is inert. If any admissible region changes, accessibility is modified; the memory is effective. □ □

Proposition 11.3 (Wave Encoding Degrees of Freedom)**.** *A constraint perturbation requires encoding at least five independent properties: strength (A), temporal-relational alignment (ϕ), semantic class (ω), interaction with co-active constraints (I), and rate of relaxation (D). The MEM|8 wave tuple* (A, ϕ, ω, I, D) *provides exactly these five.*

Proof. Necessity: each missing parameter leaves the perturbation under-specified with respect to the admissibility modification it induces. Sufficiency: together they uniquely determine the accessibility modification up to coordinate choice. □ □

Proposition 11.4 (Identity Requires Constraint Preservation)**.** *If two configurations* $x, y \in X$ *satisfy* $x \in \mathcal{A}(c) \iff y \in \mathcal{A}(c)$ *for all* $c \in \mathcal{C}_{\mathrm{rel}}$ *(the set of constraints relevant to future behavior), then x and y are functionally identical relative to the agent's trajectory.*

Proof. x and y support the same future admissibility relations over $\mathcal{C}_{\mathrm{rel}}$. Since functional identity is defined by preserved future accessibility, x and y are identical relative to $\mathcal{C}_{\mathrm{rel}}$. □ □

11.1 Framing: The Processor–Language Relationship

Definition 11.5 (Cognitive Runtime)**.** A *cognitive runtime* for a constraint-first theory $\mathcal{T} = (\mathcal{C}, \mathcal{A}, \pi)$ is a computational substrate $\mathcal{R}$ together with an interpretation map $\iota : \mathcal{T} \to \mathcal{R}$ such that:

(i) every constraint $c \in \mathcal{C}$ is represented by a data structure in $\mathcal{R}$;
(ii) every accessibility operation $\mathcal{A}(c)$ is computable within bounded time and space;
(iii) projection entropy $H(\pi, m)$ is estimable from $\mathcal{R}$'s internal state;
(iv) constraint composition, relaxation, and collapse have corresponding runtime operations.

The runtime is *partial* if some operations are approximated or restricted to a tractable subset of $\mathcal{T}$.

Remark 11.6. MEM|8 is a partial cognitive runtime in this sense. It provides concrete implementations of constraint storage (wave grid), composition (interference), relaxation (temporal decay), admissibility filtering (noise floors, attention), collapse ordering (reactive layers), curvature control (Custodian), and semantic projection (`.m8` format). It does not implement RSVP cosmology or Spherepop's full containment calculus; hence "partial." The partiality identifies precisely which aspects of the full theory have been compiled into hardware.

11.2 Wave Interference as the Natural Representation of Overlapping Constraint Fields

The claim is not merely that waves *can* represent constraints. It is that for a system of *overlapping* constraint fields, wave superposition is the unique linear representation that preserves the algebraic structure of the overlap.

Definition 11.7 (Constraint Field Overlap)**.** Two constraints $c_1, c_2 \in \mathcal{C}$ have *overlapping fields* if $\mathcal{A}(c_1) \cap \mathcal{A}(c_2) \neq \emptyset$ and $\mathcal{A}(c_1) \neq \mathcal{A}(c_2)$. The *overlap measure* is $\mathcal{O}(c_1, c_2) = \mu(\mathcal{A}(c_1) \cap \mathcal{A}(c_2)) \in [0, 1]$.
Theorem 11.8 (Wave Superposition is the Natural Representation of Constraint Overlap)**.** *Among all linear representations of constraint fields over* $\mathbb{C}$*, complex wave superposition* $W_j = A_j e^{i\phi_j}$ *is the unique representation that simultaneously:*

(i) preserves the overlap measure as $|\langle W_1, W_2 \rangle| \propto \mathcal{O}(c_1, c_2)$*;*
(ii) represents constraint reinforcement (intersection is the binding constraint) as amplitude growth;
(iii) represents constraint cancellation (union is the binding constraint) as amplitude reduction;
(iv) encodes constraint independence as phase orthogonality.

Proof. Let $W_j = A_j e^{i\phi_j}$ in a complex Hilbert space. The inner product is $\langle W_1, W_2 \rangle = A_1 A_2 e^{i(\phi_1 - \phi_2)}$ with modulus $A_1 A_2 |\cos(\phi_1 - \phi_2)|$, which is maximal at $\phi_1 = \phi_2$ (aligned, same field) and zero at $|\phi_1 - \phi_2| = \pi/2$ (orthogonal, independent fields). This matches (i).

For (ii): aligned constraints, $\phi_1 = \phi_2$, give $|W_1 + W_2| = A_1 + A_2 > \max(A_1, A_2)$. The combined field is stronger, corresponding to the intersection being the binding constraint.

For (iii): anti-aligned constraints, $\phi_1 - \phi_2 = \pi$, give $|W_1 + W_2| = |A_1 - A_2| \leq \min(A_1, A_2)$. The combined field is weaker, corresponding to the union.

For (iv): $|\phi_1 - \phi_2| = \pi/2$ gives $\langle W_1, W_2 \rangle = 0$: constraints contribute independently.

Uniqueness: a real-valued linear representation cannot independently control both reinforcement and cancellation from the same vector operation, since $|v_1 + v_2|$ in $\mathbb{R}^n$ depends only on the sign of $v_1 \cdot v_2$ and cannot encode phase-orthogonal independence. Complex representation with phase is the minimal extension supporting all four properties simultaneously. □ □

Corollary 11.9. *The MEM|8 wave grid is not an arbitrary design choice. Among linear numerical substrates, complex wave representation is the* uniquely natural *encoding for a system of overlapping constraint fields.*

11.3 Temporal Decay as Entropy Increase Under Weakening Reconstruction Constraints

Definition 11.10 (Reconstruction Constraint)**.** A *reconstruction constraint* at time t is a constraint $c_{\text{rec}}(t)$ on the fiber F^π_m such that

$$\mathcal{A}(c_{\text{rec}}(t)) = \{x \in F^\pi_m : x \text{ is consistent with the surviving trace at time } t\}.$$

As the trace decays, $c_{\text{rec}}(t)$ weakens: $\mathcal{A}(c_{\text{rec}}(t_1)) \subseteq \mathcal{A}(c_{\text{rec}}(t_2))$ for $t_1 < t_2$.

Theorem 11.11 (Exponential Decay is the Entropy-Optimal Relaxation Schedule)**.** *Among all monotone decay functions $D : [0, \infty) \to [0, 1]$ with $D(0) = 1$ and fixed half-life τ, the exponential $D(t, \tau) = e^{-t/\tau}$ is the unique function that maximizes total entropy production subject to minimizing expected future reconstruction cost at each instant.*

Proof. The reconstruction cost at time t is $-\log \mu(\mathcal{A}(c_{\text{rec}}(t)))$: the information needed to identify the correct fiber element given trace strength $D(t, \tau)$. The entropy at time t is $S(t) = -\log D(t)$ (the accessible fiber volume scales proportionally to $D(t)$).

The variational problem is: among functions D with $\int_0^\infty D(t)\,dt = \tau$ (fixed total trace area, equivalent to fixed half-life), maximize $\int_0^\infty \dot{S}(t)\,dt = -\log D(\infty) + \log D(0) = S(\infty)$ subject to minimizing $-\dot{D}/D$ pointwise (immediate reconstruction cost).

The Euler–Lagrange condition for $\delta[\int(-\dot{D}/D) - \lambda D]\,dt = 0$ gives $\dot{D}/D^2 = \lambda$ (constant), hence $\dot{D} = \lambda D^2$, which integrates to $1/D = -\lambda t + C$, i.e. for $\lambda < 0$: $D(t) = 1/(1 - \lambda t)$ with $\lambda = -1/\tau$, i.e. $D(t) = e^{-t/\tau}$ to leading order (the exact solution of $-\dot{D}/D = \text{const}$ is exponential). □ □

Remark 11.12. The five MEM|8 forgetting regimes—Flash ($\tau = 0.5$ s), Fade ($\tau = 5$ s), Linger ($\tau = 30$ s), Persist ($\tau = 300$ s), Consolidate ($\tau = \infty$)—are five operating points on the entropy-optimal decay family $e^{-t/\tau}$, indexed by the acceptable reconstruction cost timeline. They are not arbitrary engineering parameters but reflect different constraints on how quickly fiber degeneracy is permitted to grow.

11.4 Reactive Layers as a Stratified Admissibility Stack

Definition 11.13 (Admissibility Stack)**.** An *admissibility stack* over a constraint poset $\mathcal{C}$ is a finite chain

$$c_0 \succ c_1 \succ c_2 \succ \cdots \succ c_k$$

with $\mathcal{A}(c_0) \subseteq \mathcal{A}(c_1) \subseteq \cdots \subseteq \mathcal{A}(c_k)$. The stack is evaluated from the top: the admissible region at any moment is $\mathcal{A}(c_j)$ where j is the index of the most-restrictive active constraint.

Theorem 11.14 (Reactive Hierarchy is a Stratified Admissibility Stack)**.** *The MEM|8 four-layer reactive hierarchy is an admissibility stack in which:*

(i) Layer ℓ is active when $\mu(\mathcal{A}(c_\ell))$ falls below the Layer-ℓ threshold;
(ii) bypass of all layers above ℓ is correct when $|\mathcal{A}(c_\ell)| = 1$;
(iii) the bypass probability $P_{\text{bypass}}(\ell) = 1 - e^{-k(3-\ell)T_{\text{threat}}}$ is the continuous relaxation of the discrete admissibility-stack evaluation rule.

Proof. (i) By construction: Layer 0 (hardware reflex) fires when the admissible region is a singleton ($|\mathcal{A}(c_0)| = 1$). Higher layers have progressively larger admissible regions requiring deliberation to navigate.

(ii) This is Proposition 11.15: when $\mathcal{A}(c_0) = \{a^*\}$, no order can improve the output; bypass is correct and latency-optimal.

(iii) As $T_{\text{threat}} \to \infty$, $P_{\text{bypass}}(\ell) \to 1$ for all $\ell > 0$: the stack collapses to its top. As $T_{\text{threat}} \to 0$, $P_{\text{bypass}} \to 0$: all layers engage. For fixed threat level, P_{bypass} decreases with ℓ (lower layers are harder to bypass), matching the admissibility-stack property that less strict constraints require more deliberation to exit. □ □

Corollary 11.15 (Reactive Layers Are Spherepop Applied to Cognition)**.** *The reactive layer hierarchy is the containment-collapse model of Section 5 generalized from expression evaluation to action selection. In both cases, the innermost (most constrained) region resolves first; outer regions resolve only after inner ones are satisfied; evaluation order is determined by constraint depth, not arbitrary scheduling.*

Definition 11.16 (Attractor Pathology)**.** A memory system contains an *attractor pathology* when there exists a region $B \subseteq X$ such that: the update operator T satisfies $T(B) \subseteq B$ (trajectories entering B remain in B); and $\mu(\mathcal{A}_{\text{after}}(B)) < \mu(\mathcal{A}_{\text{before}}(B))$ (confinement to B reduces future accessibility).

11.5 The Custodian as Curvature Control on the Accessibility Manifold

An attractor pathology (Definition 11.16) is, at the geometric level, a region of the accessibility manifold with *negative curvature trapping*: trajectories entering a high-positive-curvature constraint basin tend to spiral inward, reducing accessible volume with each step.

Definition 11.17 (Constraint Curvature)**.** Let $\mathcal{A} : \mathcal{C} \to \mathcal{P}(\mathcal{C})$ be a constraint field and $B \subseteq \mathcal{C}$ a compact region. The *constraint curvature* of B at time 0 is

$$\kappa(B) = \frac{d}{dt} \log \mu(\mathcal{A}^t(B))\bigg|_{t=0},$$

where $\mathcal{A}^t(B) = \{c \in \mathcal{C} : \mathcal{A}^t(c) \cap B \neq \emptyset\}$ is the t-step accessible neighborhood of B. $\kappa(B) < 0$ means the accessible neighborhood shrinks: a trapping geometry. $\kappa(B) > 0$ means it expands.

Theorem 11.18 (The Custodian is a Curvature Regulator)**.** *An attractor pathology B exists if and only if $\kappa(B) < 0$ in the current dynamics. The Custodian operator K corrects the pathology by inducing $\kappa(K(B)) \geq 0$.*

Proof. Pathology iff negative curvature. By Definition 11.16, pathology requires $\mu(\mathcal{A}_{\text{after}}(B)) < \mu(\mathcal{A}_{\text{before}}(B))$, which is exactly $\kappa(B) < 0$.

K corrects by inducing $\kappa \geq 0$. From Theorem 11.18, K breaks the pathology when $K \circ T(B) \not\subseteq B$, meaning at least one trajectory escapes B per step. The accessible neighborhood therefore does not shrink: $\kappa(K(B)) \geq 0$. □ □

Remark 11.19. This reframes the Custodian from a safety mechanism to a geometric necessity. In any system maintaining accessibility over time, regions of persistently negative curvature are cognitively pathological: they trap the system, reducing its future degrees of freedom without recovery. The Custodian is the runtime operator that maintains non-negative expected curvature on the accessibility manifold. Its three regimes (allow / throttle / block) correspond to $\kappa(B) > -\varepsilon$ (normal), $-C \leq \kappa(B) \leq -\varepsilon$ (emerging trap), and $\kappa(B) < -C$ (deep trap requiring full escape).

11.6 The `.m8` Format as File-Level Semantic Projection

Definition 11.20 (Semantic Projection File)**.** A *semantic projection file* for $x \in X$ under $\pi : X \to M$ is a data structure $\mathtt{f}(x)$ such that: (i) $\pi(x)$ is losslessly recoverable from $\mathtt{f}(x)$; (ii) for all constraints $c \in \mathcal{C}_{\text{rel}}$, membership $x \in \mathcal{A}(c)$ is reconstructable from $\mathtt{f}(x)$; and (iii) $|\mathtt{f}(x)|$ is minimized subject to (i) and (ii).

Proposition 11.21 (`.m8` is a Semantic Projection File)**.** *The `.m8` format stores the tuple* (id, A, ω, ϕ, I) *in 32 bytes per memory. Here id recovers* $\pi(x) = m$*, and* (A, ω, ϕ, I) *encodes the fiber coordinate sufficient for behavioral identity (Proposition 11.4). It therefore realizes Definition 11.20.*

Proof. Condition (i): *id* recovers m. Condition (ii): by Proposition 11.3, (A, ω, ϕ, I, D) encodes the five degrees of freedom of the constraint perturbation; by Proposition 11.4, this determines membership in all relevant admissible regions. Condition (iii): Theorem 3.3 guarantees that x contains information beyond m; the wave tuple is the minimal five-parameter encoding of that additional information. □ □

Remark 11.22*.* A conventional file stores content: elements of M, the projected space. An `.m8` file stores constraint structure: the fiber coordinate, the un-projected information. The 99% compression figure is therefore not primarily a technical achievement but an algebraic consequence: Theorem 3.3 guarantees that $|X| \gg |M|$, so representing only the five-parameter constraint tuple rather than the full $x \in X$ achieves compression proportional to the ratio $|X|/|M|$, which is large for high-entropy memories. The Marqant layer handles residual string-level redundancy of the projected content.

11.7 The Master Derivation

Theorem 11.23 (MEM|8 is a Partial Cognitive Runtime)**.** *MEM|8 is a partial cognitive runtime for* $(\mathcal{C}, \mathcal{A}, \pi)$*, with the following theory-to-runtime correspondences:*

Theory primitive	*Runtime implementation*
Constraint field $\mathcal{A}$	*Wave grid* $(256^2 \times 65536)$
Constraint perturbation	*Memory wave* $M_{xyz}(t)$
Overlapping constraint fields	*Wave interference (Thm. 11.8)*
Constraint relaxation / entropy	*Exponential decay* $D(t, \tau)$ *(Thm. 11.11)*
Admissibility stack	*Reactive layer hierarchy (Thm. 11.14)*
CLIO inference	*Noise floors, temporal blankets, attention weights*
Curvature control	*Custodian operator (Thm. 11.18)*
Semantic projection	`.m8` *format (Prop. 11.21)*
Chain of Memory	*Wave constraint preservation, not state replay*

Proof. Each row is established by the cited result. Condition (i) of Definition 11.5: every constraint is representable as a wave pattern. Condition (ii): $\mathcal{A}(c)$ is computable from the interference pattern in bounded time. Condition (iii): entropy is estimable from amplitude distribution (high amplitude = strong constraint = low entropy). Condition (iv): composition is superposition, relaxation is temporal decay, collapse is reactive bypass. The runtime is partial: RSVP's global scalar field is not fully numerically implemented; Spherepop's full containment calculus is approximated by the urgency-ordered special case; CLIO search is approximated by threshold filtering. □ □

Corollary 11.24 (Memory = Persistent Modification of Future Accessibility)**.** *The equation*

$$\textit{memory} = \textit{persistent modification of future accessibility}$$

is not a definition within the constraint-first framework. It is a theorem: the unique characterization of cognitively effective memory (Proposition 11.2), derived from the accessibility ontology (Definition 2.3). When Chenoweth and Chenoweth state this as a "final principle" [10], they are identifying—from an engineering starting point—the primitive that the present framework derives from ontological first principles.

Corollary 11.25 (Consciousness as Recursively Regulated Accessibility)**.** *Consciousness, modeled as global availability of constraint-weighted projections to deliberative processing, is in MEM|8 not the presence of waves alone. It is the system's capacity to regulate which wave-encoded constraints enter the active projection manifold:*

$$\textit{consciousness} = \textit{recursively regulated accessibility.}$$

The recursive qualifier reflects that the regulation mechanism is itself subject to the same constraint-accessibility structure, formalized by the chain of memory (Definition 11.1).

11.8 Critical Assessment

Theorem 11.23 establishes structural correspondence; it does not validate the empirical claims of [10]. The reported performance figures (973× insertion speedup, 292× retrieval speedup) lack disclosed experimental protocols and independent replication. The consciousness metrics are novel and non-standard; their relationship to peer-reviewed consciousness frameworks remains unestablished. The wave–neural oscillation connection is analogical: the computational wave grid does not demonstrably instantiate neural dynamics in any rigorous sense. The "sensory free will" framing conflates control over *access* with phenomenal consciousness; within the present framework, the architecture establishes that the system has partial control over which trajectories enter the projection manifold M—agency in the sense of Corollary 11.25 but not phenomenal experience without additional argument.

None of this diminishes the structural convergence. An engineering effort starting from neural oscillation theory and wave physics independently arrived at constraint-shaped forgetting, attention as accessibility modulation, reflexive collapse under narrow admissibility, curvature-controlled attractors, and identity as constraint rather than state preservation. This convergence constitutes evidence that the constraint-first architecture reflects something about the conceptual structure of cognition, not merely an idiosyncratic choice of ontology.

11.9 Distributed Mixture-of-Experts as Accessibility Compression

The preceding subsections established MEM|8 as a partial cognitive runtime for the constraint-first framework. A recent result in machine learning provides concrete empirical confirmation that large computational systems naturally organize around the same structure. The dMoE architecture of Feng et al. [9] reduces a diffusion language model's active expert count from 69.5 to 14.6 while retaining 99.11% of performance—a compression ratio of 4.77× with near-zero functional loss. This section argues that dMoE is not merely an engineering optimization but a computational instantiation of the Map-Territory Projection Theorem (Theorem 3.3) and the projection entropy framework (Section 4).

The conventional model as process ontology. A standard transformer with token-level MoE routing embodies a process ontology in the sense of Section 2: computation is a sequence of token-level operations, and the routing decision is an individual event attached to each token. The question the model asks at each step is *which computation should this token perform?*—a local, trajectory-centric question.

dMoE moves one level down, in exactly the sense of Proposition 2.4. Instead of asking which computation each token should perform, it first constructs a shared routing structure—the block-level expert distribution $\widetilde{S}_{\text{block}}$—for the entire block, and only then permits token trajectories inside that structure. The question becomes: *what is the admissible region of expert space for this block context?* Token routing is then constrained navigation within that region.

The block distribution as accessibility field. The block-level expert distribution is formed by aggregating token-level router scores:

$$S_{\text{block}} = \bigoplus_{i \in B} \hat{s}_i, \qquad \hat{s}_i = \text{TopKMask}(s_i, k),$$

and normalizing to a probability vector $\widetilde{S}_{\text{block}} \in \Delta^{|E|-1}$. The coreset is then selected by top-p thresholding: $\mathcal{C} = \text{Top-P}(\widetilde{S}_{\text{block}}, p)$.

This is an accessibility field in the sense of Definition 2.3: it assigns to the current block configuration an admissible set of computational resources $\mathcal{C}$, and all token-level computation is constrained to navigate within $\mathcal{C}$. The correspondence with the framework is exact:

Beyond Process	**dMoE**
Constraint field $\mathcal{A}$	Block-level expert distribution $\widetilde{S}_{\text{block}}$
Admissible trajectories	Expert routes within coreset $\mathcal{C}$
Object (invariant basin)	Stable expert subset under repeated blocks
Projection $\pi : X \to M$	Top-p coreset selection
Fiber degeneracy $\pi^{-1}(m)$	Multiple equivalent token-level routes
Projection entropy	Routing redundancy across the full expert space

The Map-Territory Theorem in a working model. The Map-Territory Projection Theorem (Theorem 3.3) states that any surjective map $\pi : X \to M$ with $|X| > |M|$ has nontrivial fibers: distinct elements of X map to the same element of M, and the information distinguishing them is invisible in M. dMoE's core empirical result is a direct computational demonstration of this theorem.

Let X be the space of all possible token-level routing assignments across a block (combinatorially large, scaling as $\binom{|E|}{k}^{|B|}$) and M be the space of block-level routing assignments constrained to the coreset (much smaller, scaling as $\binom{|\mathcal{C}|}{k}^{|B|}$ with $|\mathcal{C}| \approx 14$ versus $|E| \approx 70$). The projection $\pi : X \to M$ maps full routing assignments to their coreset-constrained counterparts.

The nontrivial fibers are the distinct full-routing assignments that produce the same coreset-constrained output: they are *functionally equivalent* under the projection, because the model cannot distinguish between them at the block level. The performance result (99.11% retention) confirms that these fibers are not only large but *functionally degenerate*: the information that distinguishes elements within a fiber carries almost no functional significance. Discarding the fiber coordinate—collapsing to the coreset manifold—loses less than 1% of performance.

Proposition 11.26 (dMoE Instantiates the Projection Theorem)**.** *The dMoE result is a computational instantiation of Theorem 3.3: the full routing space X projects onto the coreset space M via* $\pi = \text{Top-P} \circ \text{Normalize} \circ \bigoplus$*, and the fibers* $\pi^{-1}(m)$ *are large but functionally degenerate, as confirmed by near-zero performance loss under coreset restriction.*

Proof. By Theorem 3.3, $|X| > |M|$ guarantees nontrivial fibers. The empirical result establishes that the fibers are *functionally* nontrivial in size (many token routes map to the same coreset) but *informationally* degenerate (they are functionally equivalent for the model's task). This is exactly the projection theorem's content: the information distinguishing fiber elements—which specific routes tokens took within the coreset—is invisible in the observable output. □ □

Memory as coreset storage, not route storage. The connection to MEM|8 is now precise. Theorem 11.23 established that MEM|8 stores constraint structure (the fiber coordinate sufficient for behavioral identity) rather than full state trajectories. dMoE suggests what this means concretely in a large language model context: effective memory stores the block accessibility field $(S_{\text{block}}, \mathcal{C})$ rather than the individual routing decisions $\{R_i\}_{i \in B}$. Future inference reconstructs the admissibility manifold from the stored coreset, not from replaying stored routes.

This is the computational realization of the Chain of Memory principle (Definition 11.1): identity persists through constraint-preserving reconstruction, not state preservation. The stored coreset $\mathcal{C}$ carries the constraint structure—which experts are admissible—while the specific routes within $\mathcal{C}$ are fiber elements: interchangeable, degenerate, and safely discarded.

Corollary 11.27 (Effective Memory = Admissibility Structure, Not Route History)**.** *A system that stores the block-level accessibility field* $(S_{\text{block}}, \mathcal{C})$ *rather than individual token routes* $\{R_i\}_{i \in B}$ *retains the functionally relevant constraint structure while discarding fiber-degenerate routing details. The dMoE result confirms empirically that this compression is achievable with* $4.77\times$ *reduction in stored routing complexity and near-zero functional loss.*

Top-p as CLIO in expert space. The top-p coreset selection is formally isomorphic to the CLIO sparse admissibility operator (Theorem 7.3). CLIO restricts inference to admissible queries given a document context. dMoE restricts routing to admissible experts given a block context. Both are instances of the same operation—sparse projection onto a tractable admissible subset—applied to different substrates: inferential state space in CLIO, expert index space in dMoE.

The threshold p in dMoE corresponds to the coherence threshold θ in Chain of Memory: it determines how much of the accessibility structure must be preserved for the projection to be considered semantically coherent. A lower p means more aggressive compression; a lower θ means looser identity preservation. Both trade fiber detail for computational tractability.

12 The Implementation Tower

The full derivation assembles into a layered tower in which each level is the unique realization of the level above under tractability and continuity constraints.

Definition 12.1 (The Constraint-First Implementation Tower)**.**

Level	Framework	Domain
L1	Semantic Infrastructure	Ontology: constraints, accessibility, projection
L2	CLIO	Inference: sparse admissibility search
L3	Spherepop	Evaluation: urgency-ordered containment collapse
L4	Chain of Memory	Continuity: constraint-preserving reconstruction
L5	RSVP	Field dynamics: scalar/vector accessibility evolution
L6	MEM|8	Runtime: wave grids, decay, reactive layers, compression

Each level is a projection of L1 onto a specific computational or physical screen.

Theorem 12.2 (Each Level is Uniquely Derived from the Level Above)**.** *Each level $L(n+1)$ is the unique (up to implementation choice) realization of $L(n)$'s abstract principles subject to the constraints of tractability, continuity, and minimality.*

Proof. L1 $\rightarrow$ L2. Given that cognition is accessibility navigation, inference must restrict to admissible regions (exhaustive search over X is intractable by cardinality). CLIO is the minimal inference principle preserving semantic coherence (Theorem 7.3) while bounding search.

L2 $\rightarrow$ L3. Given CLIO-style restricted inference, evaluation requires a total order. Urgency-ordered collapse is the unique latency-optimal order: when $|\mathcal{A}(c)| = 1$, no other order can improve the output (Proposition 11.15).

L3 $\rightarrow$ L4. Given ordered collapse, persistence through collapse requires constraint-based identity (states are destroyed by collapse; Proposition 11.4). Chain of Memory is the minimal continuity structure satisfying this.

L4 $\rightarrow$ L5. Given constraint-preserving reconstruction, a field language is needed for how the accessibility landscape evolves between collapse events. RSVP is the unique Lorentz-compatible scalar-vector formalism for accessibility evolution (Proposition 8.2).

L5 $\rightarrow$ L6. Given a field language, a runtime must encode constraint perturbations numerically. Theorem 11.8 establishes that complex wave representation is the unique natural encoding for overlapping constraint fields. Theorem 11.11 establishes exponential decay as the entropy-optimal relaxation schedule. Together they uniquely determine the wave grid architecture. Reactive layers, Custodian, and `.m8` follow from Theorems 11.14, 11.18, and Proposition 11.21. □ □

Remark 12.3 (The Tower as a Processor–Language Stack). The six-level relationship is precisely analogous to a compiler/processor stack:

CS analogy	**Constraint-first tower**
Specification language	Semantic Infrastructure (L1)
Type system / inference	CLIO (L2)
Evaluation strategy	Spherepop (L3)
Memory model	Chain of Memory (L4)
Instruction set / ABI	RSVP (L5)
Processor microarchitecture	MEM|8 (L6)

Just as a program written in a high-level language is ultimately executed by transistor-level operations, a cognitive process described at the level of accessibility ontology is ultimately executed by wave-grid operations. The levels do not contradict each other; each is a compiled form of the one above. This is what it means to say that MEM|8 is a *partial machine realization* of the constraint-first framework—not a compatible or analogous system, but an execution model for its primitives.

The claim in the preface of *Semantic Infrastructure*—that the frameworks are "five projections of one theory onto five different observable screens"—is therefore a theorem. Theorem 12.2 establishes that the projections are deductive, not analogical: each screen is the unique realization of the theory's primitives at that computational level.

13 Conclusion: Everything Navigates

The arguments of the preceding sections can now be gathered into a single proposition.

Theorem 13.1 (Master Theorem)**.** *Let any system be described by a triple $(\mathcal{C}, \mathcal{A}, \pi)$ where $\mathcal{C}$ is a configuration space, $\mathcal{A} : \mathcal{C} \to \mathcal{P}(\mathcal{C})$ a constraint field, and $\pi : \mathcal{C} \to M$ a projection to an observable space M. Then:*

(i) All stable structures are basins of $\mathcal{A}$ (Proposition 2.4).
(ii) All processes are trajectories through $\mathcal{A}$ (Definition 2.3).
(iii) All information loss is projection entropy (Proposition 4.2).
(iv) All coordination is stigmergic modification of $\mathcal{A}$ (Proposition 9.5).
(v) All inference is admissibility-constrained trajectory selection (Theorem 7.3).

Therefore, the appropriate governing slogan is not Heraclitus's "everything flows" *but:*

everything navigates a changing landscape of admissible possibilities.

Proof. Each clause follows from the corresponding theorem or proposition cited. Together, they constitute an exhaustive account of structure, process, information, coordination, and inference within the framework. The slogan collects the common pattern: in each case, the relevant phenomenon is not the trajectory taken but the accessibility landscape within which trajectory selection occurs. □ □

Corollary 13.2 (Unity of the Frameworks)**.** *Spherepop, RSVP, CLIO, Admissibility Lab, Yarncrawler, Semantic Infrastructure, and MEM|8 are projections of a single underlying theory—a geometric theory of constrained accessibility—onto seven different observable screens: computation, cosmology, inference, document structure, narrative coherence, civilizational organization, and artificial cognitive runtime.*

Proof. Each framework instantiates the triple $(\mathcal{C}, \mathcal{A}, \pi)$: Spherepop with $\mathcal{C}$ = containment structures, $\mathcal{A}$ = admissible Pop operations, π = string linearization; RSVP with $\mathcal{C}$ = field configurations, $\mathcal{A}$ = accessibility gradients, π = cosmological observables; CLIO with $\mathcal{C}$ = document states, $\mathcal{A}$ = admissible queries, π = answer space; Admissibility Lab with $\mathcal{C}$ = semantic neighborhoods, $\mathcal{A}$ = perturbation admissibility, π = embedding clusters; Yarncrawler with $\mathcal{C}$ = narrative trajectory space, $\mathcal{A}$ = coherent continuations, π = surface text; Semantic Infrastructure with $\mathcal{C}$ = full knowledge-state space, $\mathcal{A}$ = admissible inferential moves, π = observable propositional content. The structural isomorphism of all six instantiations confirms a common origin. □ □

The history of philosophy has debated being versus becoming for two and a half millennia. The present proposal does not resolve that debate by choosing a side. It relocates the question. Before there can be stable beings or flowing becomings, there must be a geometry—an accessibility landscape—that makes some states persistent and some transitions possible. That geometry is the ground of both. It is not a thing, not an event, not a flow. It is the structure of what can happen from here.

Principal sources for the process philosophy sections: A. N. Whitehead, *Process and Reality* (1929); Nicholas Rescher, *Process Philosophy* (1996); Heraclitus, fragments (DK 22). For information geometry and constraint ontology: S.-i. Amari, *Information Geometry and Its Applications* (2016); J. Ladyman and D. Ross, *Every Thing Must Go* (2007). For stigmergy: P.-P. Grassé (1959); E. Bonabeau et al., *Swarm Intelligence* (1999). For MEM|8: C. M. Chenoweth and A. A. Chenoweth, "MEM|8: A Wave-Based Cognitive Architecture for Multimodal Memory Integration and Consciousness Simulation," Zenodo preprint, DOI: 10.5281/zenodo.16436298 (2025). All formal frameworks (RSVP, Spherepop, CLIO, Admissibility Lab, Yarncrawler, Semantic Infrastructure) are the author's own; see associated monographs at `github.com/standardgalactic`.

References

[1] A. N. Whitehead, *Process and Reality: An Essay in Cosmology*. Macmillan, New York, 1929.

[2] N. Rescher, *Process Metaphysics: An Introduction to Process Philosophy*. SUNY Press, Albany, 1996.

[3] S. Amari, *Information Geometry and Its Applications*. Springer, Tokyo, 2016.

[4] J. Ladyman and D. Ross, *Every Thing Must Go: Metaphysics Naturalised*. Oxford University Press, Oxford, 2007.

[5] P.-P. Grassé, "La reconstruction du nid et les coordinations inter-individuelles chez *Bellicositermes natalensis* et *Cubitermes* sp. La théorie de la stigmergie," *Insectes Sociaux*, 6:41–83, 1959.

[6] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.

[7] B. J. Baars, *A Cognitive Theory of Consciousness*. Cambridge University Press, Cambridge, 1988.

[8] S. Franklin et al., "LIDA: A systems-level architecture for cognition, emotion, and learning," *IEEE Transactions on Autonomous Mental Development*, 6(1):19–41, 2014.

[9] S. Feng, Z. Chen, G. Fang, X. Ma, and X. Wang, "dMoE: dLLMs with Learnable Block Experts," arXiv:2605.30876, May 2026.

[10] C. M. Chenoweth and A. A. Chenoweth, "MEM|8: A Wave-Based Cognitive Architecture for Multimodal Memory Integration and Consciousness Simulation," Zenodo preprint, DOI: 10.5281/zenodo.16436298, July 2025.