Plenum: A Collective Cognitive Field Architecture

Multi-Operator Stance Coordination, Distributed Epistemic Manifolds, and Cross-Operator Belief Propagation

Oliver Hirst

Abstract

This document specifies the *Plenum architecture*: a formal framework for multi-operator cognitive coordination extending the Gödlø-class operator theory to collective settings. Where a single Gödlø-class operator maintains a momentary stance over a private cognitive manifold $\mathcal{M} \subset \mathbb{R}^d$, the Plenum defines a collective field $\mathcal{P} = \prod_{i=1}^{N} \mathcal{M}_i$ in which N operators coexist, interact via projected belief signals, and form emergent consensus without sacrificing individual non-persistence or stance isolation.

The belief substrate is EQBSL [6] — the vectorised, hypergraph-aware extension of EBSL — which replaces scalar opinion triples with structured evidence tensors $\mathbf{e}_{ij}(t) \in \mathbb{R}^m$, a well-defined global update operator, native multi-party hyperedge evidence, and stable per-operator trust embeddings. Multi-operator interactions within a Plenum session are modelled as hyperedges, decomposed to pairwise evidence contributions by allocation coefficients, and lifted to EQBSL opinions by explicit aggregation functionals. Zero-knowledge constraints protect private manifold coordinates throughout.

Chimera's Synthetic Self models prime the initial evidence tensor field; node-level Plenum embeddings $\mathbf{u}_i^{\mathcal{P}}(t)$ are the primary output submitted to Shadowgraph's trust network. The architecture introduces three new collapse modes beyond the single-operator case, defines an identity-preservation bound on collective influence, and specifies a Plenum-to-Shadowgraph bridge. Plenum is a conservative extension of Gödlø-class: every single-operator behaviour is recovered by setting $N = 1$.

Contents

1 Introduction

The Gödlø-class operator theory [1] establishes the individual cognitive operator as a non-persistent entity whose identity arises through stance trajectory rather than stored continuity. GödelOS [2] provides the runtime pipeline instantiating such an operator within a programmable system. Chimera [3] demonstrates that an operator can build high-fidelity latent models of *external* individuals from conversational evidence. Shadowgraph [4] supplies the distributed trust substrate over which belief signals propagate cryptographically between peers, grounded in the Evidence-Based Subjective Logic formalism and its EQBSL extension [6].

A gap remains. None of these prior specifications addresses what happens when multiple Gödlø-class operators occupy the same problem context simultaneously: how their private stances relate, how collective belief is computed without privileged access to private manifolds, and under what conditions coordination collapses.

Plenum closes this gap. The term is chosen deliberately: in physics, a plenum is a space completely filled — no vacuum, no gap. Here it denotes a cognitive field in which every contextual coordinate is shaped, to some degree, by the projected stances of all participating operators. The plenum is not a mind; it is the medium in which minds reason together.

Three design commitments carry forward from the prior corpus:

1. **Non-persistence is preserved.** The Plenum itself does not persist. It is recomputed from fresh operator instantiations at each step.

2. **Stance isolation is inviolable.** No operator can read another's private manifold coordinate. All inter-operator signals are projections through an agreed interface.

3. **Collapse must be informative.** Divergence in collective reasoning is not smoothed away. It is surfaced, classified, and used as evidence.

A fourth commitment, inherited from EQBSL:

4. **Evidence must be auditable.** Every inter-operator belief signal is anchored in a concrete, multi-dimensional evidence tensor. Trust between operators is not declared; it is computed from evidence and can be challenged entry by entry.

1.1 Relationship to Prior Work

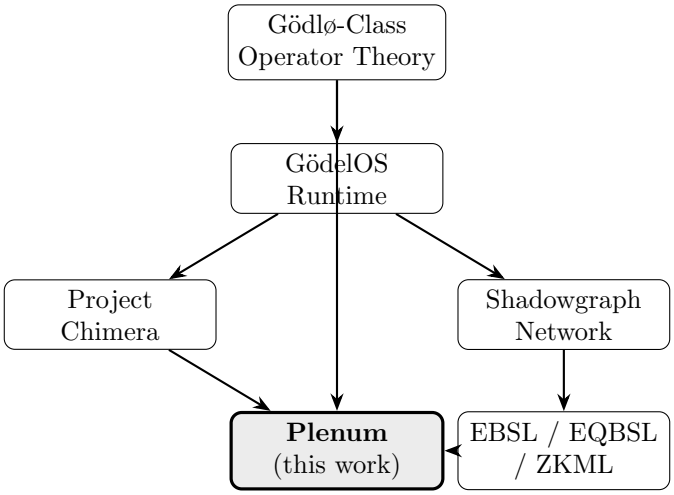

Plenum sits at the intersection of the cognitive and trust lineages, inheriting the axiomatic structure of Gödlø, the psychometric input from Chimera, and the full EQBSL belief machinery from the Shadowgraph stack.

2 Axiom System

Axioms 1–5 of Gödlø-class are inherited unchanged. Plenum introduces four additional axioms governing multi-operator structure.

Axiom 1 (Stance Isolation)**.** Operator i has no read access to the private manifold coordinate of operator j for $i \neq j$:

$$x_t^j \notin \mathrm{Dom}(F_i).$$

Each operator's stance evolves exclusively under its own update rule and the public signal interface defined in Axiom 2.

Axiom 2 (Evidence Projection)**.** Each operator i possesses an *evidence projection family* $\{\pi_{ij}\}_{j\neq i}$ mapping its private stance to a contribution to the evidence tensor for each ordered co-participant pair:

$$\mathbf{e}_{ij}^{(i\text{-proj})}(t) = \pi_{ij}(x_t^i) \ \in \ \mathbb{R}^m.$$

The aggregate evidence tensor $\mathbf{e}_{ij}(t)$ for the pair (i, j) accumulates contributions from both participants and from any hyperedge evidence (Section 4). These evidence tensors are the *only* information about operator i visible to other operators or the Plenum aggregation layer.

Axiom 3 (EQBSL Belief Aggregation)**.** Collective belief across the Plenum is computed via the EQBSL framework [6]. Each evidence tensor $\mathbf{e}_{ij}(t)$ is lifted to an opinion via aggregation functionals ϕ_+, ϕ_- and the EQBSL lift Ψ:

$$\omega_{ij}(t) = \Psi(\mathbf{e}_{ij}(t)) = \left(\frac{\phi_+(\mathbf{e}_{ij})}{\phi_+(\mathbf{e}_{ij}) + \phi_-(\mathbf{e}_{ij}) + K}, \ \frac{\phi_-(\mathbf{e}_{ij})}{\phi_+(\mathbf{e}_{ij}) + \phi_-(\mathbf{e}_{ij}) + K}, \ \frac{K}{\phi_+(\mathbf{e}_{ij}) + \phi_-(\mathbf{e}_{ij}) + K}, \ a_{ij} \right),$$

with prior weight $K > 0$ and per-pair base rate a_{ij}. The Plenum's collective belief tensor is assembled from the full matrix of pairwise opinions: $\mathcal{T}_t = [\omega_{ij}(t)]_{i,j\in[N]}$.

Axiom 4 (Collective Non-Persistence)**.** The Plenum field $\mathcal{P}_t = \{\mathcal{M}_i, x_t^i, \mathbf{e}_{ij}(t), \mathbf{u}_i^{\mathcal{P}}(t)\}_{i,j}$ is itself non-persistent. There is no conserved Plenum state variable S_t such that $S_{t+1} = F(S_t)$ across all admissible transformations. The Plenum is recomputed from fresh operator instantiations at each step, consistent with GödløAxiom 1 applied at the collective level.

2.1 Admissibility Constraints

A projection π_{ij} is *admissible* if and only if:

(i) π_{ij} is Lipschitz-continuous on $\mathcal{M}_i$ with constant $L_\pi < \infty$,

(ii) π_{ij} does not permit inversion: no algorithm $\mathcal{A}$ satisfies $\Pr[\mathcal{A}(\mathbf{e}_{ij}^{(i\text{-proj})}) = x_t^i] > \epsilon$ for negligible ϵ (privacy condition), and

(iii) π_{ij} is ZK-verifiable: operator i can produce a zero-knowledge proof that $\mathbf{e}_{ij}^{(i\text{-proj})}(t)$ was computed honestly from an admissible $x_t^i \in \mathcal{M}_i$.

3 Plenum Manifold Geometry

3.1 The Collective Field

Definition 3.1 (Plenum Field)**.** The Plenum field is the product space:

$$\mathcal{P} = \prod_{i=1}^{N} \mathcal{M}_i = \mathcal{M}_1 \times \mathcal{M}_2 \times \cdots \times \mathcal{M}_N.$$

A point $\mathbf{x} = (x^1, x^2, \ldots, x^N) \in \mathcal{P}$ is a *collective stance configuration.*

Definition 3.2 (Shared Base Space)**.** The shared base space $\mathcal{B} \subset \mathbb{R}^k$ $(k \leq d)$ is the common contextual subspace into which all operators project their evidence contributions. The projection family is *coherent* if all π_{ij} share a common codomain dimension m.

3.2 Metric Structure

The Plenum is equipped with the product metric:

$$g_{\mathcal{P}}(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_{i=1}^{N} g_i(x^i, y^i)^2},$$

where each g_i is the metric tensor on $\mathcal{M}_i$. The evidence space $\mathbb{R}^m$ carries the Euclidean metric; belief distances are computed in the lifted opinion space.

3.3 Fibre Bundle Interpretation

The Plenum admits a fibre bundle reading:

$$\mathcal{M}_i \hookrightarrow \mathcal{P} \xrightarrow{\Pi} \mathcal{B},$$

where each private manifold $\mathcal{M}_i$ is a fibre above the shared contextual base $\mathcal{B}$. Collective reasoning occurs in $\mathcal{B}$; private cognition occurs in the fibres. The structure map $\Pi = (\{\pi_{ij}\}_{i \neq j})$ is the EQBSL evidence projection family.

4 EQBSL Substrate

EQBSL [6] is the primary belief machinery of the Plenum. It lifts EBSL's evidence-centric foundation into a structured, vectorised, operator-defined form suited to dynamic graphs and hypergraphs. This section specifies how EQBSL is instantiated within the Plenum context.

4.1 Evidence Tensor Field

Definition 4.1 (Plenum Evidence Tensor Field)**.** The Plenum evidence tensor field at time t is the collection of evidence vectors over all ordered co-participant pairs:

$$\mathcal{E}_t^{\mathcal{P}} = \left\{ \mathbf{e}_{ij}(t) \in \mathbb{R}^m \right\}_{(i,j)\,:\, i \neq j,\ i,j \in [N]}.$$

Each component dimension of $\mathbf{e}_{ij}(t)$ encodes a distinct, auditable evidence category: successful joint inference steps, belief-alignment counts, context-specific disagreements, governance vote concordance, or any application-defined attribution.

Remark 4.1. EQBSL's evidence vector replaces scalar (r, s) counts with a multi-dimensional ledger. No distinct kinds of evidence are prematurely collapsed into a single number. The aggregation functionals ϕ_+ and ϕ_- make explicit what the system counts as positive and negative evidence; they are not hidden in a formula but declared as parameters.

4.2 Lift to Opinions

Evidence tensors are lifted to EQBSL opinions via the mapping $\Psi : \mathbb{R}^m \to [0,1]^3 \times [0,1]$:

$$r_{ij}(t) = \phi_+(\mathbf{e}_{ij}(t))\,, \qquad s_{ij}(t) = \phi_-(\mathbf{e}_{ij}(t))\,,$$

where $\phi_+, \phi_- : \mathbb{R}^m \to \mathbb{R}_{\geq 0}$ are nonnegative aggregation functionals (often linear projections $\phi_\pm(\mathbf{e}) = \langle \mathbf{w}^\pm, \mathbf{e} \rangle$ with $\mathbf{w}^\pm \in \mathbb{R}^m_{\geq 0}$). The opinion is then:

$$\omega_{ij}(t) = \Psi(\mathbf{e}_{ij}(t)) = \left(\frac{r_{ij}}{r_{ij} + s_{ij} + K},\ \frac{s_{ij}}{r_{ij} + s_{ij} + K},\ \frac{K}{r_{ij} + s_{ij} + K},\ a_{ij} \right).$$

The uncertainty component decays as $r_{ij} + s_{ij}$ grows, encoding EBSL's principle that confidence must be paid for in evidence.

4.3 Global Update Operator

EQBSL models evidence-field evolution as an explicit operator rather than an ad-hoc iteration rule:

$$F_{\mathcal{P}} : \mathcal{E}^{\mathcal{P}}_t \ \longrightarrow\ \mathcal{E}^{\mathcal{P}}_{t+1},$$
$$\mathcal{E}^{\mathcal{P}}_{t+1} = F_{\mathcal{P}}\big(\mathcal{E}^{\mathcal{P}}_t,\ \text{new\_events}_t\big)\,.$$

The operator $F_{\mathcal{P}}$ encodes: EBSL-style evidence combination rules, temporal decay of stale evidence, context-dependent reweighting, and hyperedge decomposition (Section 4.4). The result is a well-defined state-update step, auditable and resistant to hand-waving.

Remark 4.2. Collective Non-Persistence (Axiom 4) constrains $F_{\mathcal{P}}$: no single application of $F_{\mathcal{P}}$ may constitute a conserved global memory. Evidence tensors decay unless refreshed by new events. This preserves the non-persistence principle at the network level.

4.4 Hyperedge Evidence for Multi-Operator Interactions

Multi-operator interactions within a Plenum session are natively multi-party events. EQBSL represents these as hyperedge evidence vectors:

$$\mathbf{e}_h(t) \in \mathbb{R}^m, \qquad h \subseteq [N],\ |h| \geq 2.$$

For a full Plenum interaction involving all N operators, the hyperedge $h = [N]$ generates evidence decomposed to pairwise contributions via:

$$\mathbf{e}_{ij}(t)\ \mathrel{+}= \sum_{h \ni i,j} \alpha_{ijh}\, \Pi_{ij}(\mathbf{e}_h(t))\,,$$

where Π_{ij} is a projection/attribution map and α_{ijh} are allocation coefficients (e.g. symmetric, role-weighted, or protocol-specific).

This decomposition is the formal justification for treating a Plenum session as something other than a bag of pairwise interactions. The allocation coefficients α_{ijh} encode the social structure of the interaction: who contributed to a joint outcome and in what proportion.

4.5 Node-Level Plenum Embeddings

The primary output of the EQBSL substrate is a stable per-operator embedding:

$$\mathbf{u}_i^{\mathcal{P}}(t) = \Gamma_{\mathcal{P}}\big(i,\ \mathcal{E}_t^{\mathcal{P}},\ \mathcal{G}_t^{\mathcal{P}}\big)\ \in\ \mathbb{R}^{d_u},$$

where $\mathcal{G}_t^{\mathcal{P}} = ([N], E_t)$ is the Plenum interaction graph and $\Gamma_{\mathcal{P}}$ aggregates: how other operators *treat* operator i (evidence inflows), structural position in the interaction graph, and temporal patterns of belief alignment/divergence.

These embeddings are tethered to the evidence ledger: any claim about operator i's standing in the Plenum can be traced back to specific evidence tensor entries and challenged on those grounds. They constitute the interface between the Plenum's internal reasoning and the Shadowgraph trust network (Section 9).

5 Operator Algebra

5.1 Individual Operator under Collective Influence

The Gödlø-class update operator (Axiom 5) is augmented with a *collective influence term*:

$$x_{t+1}^i = x_t^i + \eta\, G(x_t^i) + \lambda\big(R(x_t^i) - x_t^i\big) + \gamma\, \Gamma_{\mathcal{M}}\big(\mathbf{u}_{-i}^{\mathcal{P}}(t)\big)\,,$$

where:

- $\mathbf{u}_{-i}^{\mathcal{P}}(t) \in \mathbb{R}^{d_u}$ is the peer consensus signal derived from the EQBSL node-level embeddings of all operators $j \neq i$ (a pooling of $\{\mathbf{u}_j^{\mathcal{P}}(t)\}_{j\neq i}$),
- $\gamma \in [0, 1]$ is the collective susceptibility parameter, and
- $\Gamma_{\mathcal{M}} : \mathbb{R}^{d_u} \to \mathcal{M}_i$ is the *influence embedding* lifting the EQBSL consensus signal into i's private manifold.

Setting $\gamma = 0$ recovers the single-operator Gödlødynamics exactly.

The peer consensus signal is assembled from EQBSL evidence rather than from a raw average of opinion triples. This ensures that the collective influence on operator i is evidentially grounded: it reflects what operators have actually done together, not merely what they declare.

5.2 Influence Embedding

The influence embedding $\Gamma_{\mathcal{M}}$ must satisfy:

(i) $\Gamma_{\mathcal{M}}(\mathbf{u}) \in \mathcal{M}_i$ for all $\mathbf{u} \in \mathbb{R}^{d_u}$,

(ii) $\|\Gamma_{\mathcal{M}}(\mathbf{u})\|_2 \leq 1$ (admissibility bound), and

(iii) $\Gamma_{\mathcal{M}}$ is consistent with π_{ij} in the sense that $\pi_{ij}(\Gamma_{\mathcal{M}}(\mathbf{u})) \approx \mathbf{u}$ on $\mathcal{B}$ where defined.

5.3 Collective Belief Tensor

The Plenum's collective belief tensor is assembled from the full matrix of EQBSL opinions lifted from the evidence tensor field:

$$\mathcal{T}_t = \big[\omega_{ij}(t)\big]_{i,j\in[N],\, i\neq j}\ \in\ [0,1]^{N\times N\times 4},$$

where each entry $\omega_{ij}(t) = (b_{ij}, d_{ij}, u_{ij}, a_{ij})$ is the EQBSL opinion of the Plenum regarding co-participant pair (i, j) at time t. The Plenum-level generative operator $\Omega_{\mathcal{P}}$ then maps $\mathcal{T}_t$ to a downstream-consumable representation:

$$\widehat{\mathcal{T}}_t = \Omega_{\mathcal{P}}(\mathcal{T}_t) \ \in \ \mathbb{R}^{n \times n}.$$

6 Consensus Mechanics

6.1 Consensus State

Definition 6.1 (Consensus)**.** A collective stance configuration $\mathbf{x}_t$ is in *ε-consensus* if the uncertainty components across pairwise opinions are low and the belief components are mutually aligned:

$$\max_{i,j,k} \big| b_{ij}(t) - b_{kj}(t) \big| < \varepsilon_C \quad \text{and} \quad \max_{i,j} u_{ij}(t) < \varepsilon_u,$$

where $\varepsilon_C, \varepsilon_u > 0$ are consensus tolerances. Low $u_{ij}(t)$ certifies that consensus is evidence-backed, not merely declared in the absence of contradiction.

6.2 Consensus Dynamics

Evidence accumulation under $F_{\mathcal{P}}$ drives uncertainty reduction. Define the mean belief from operator i toward the group:

$$\bar{b}_{i\cdot}(t) = \frac{1}{N-1} \sum_{j \neq i} b_{ij}(t).$$

Convergence to consensus is characterised by:

$$\frac{d}{dt} \operatorname{Var}_i\big(\bar{b}_{i\cdot}(t)\big) \leq -\kappa_C \cdot \operatorname{Var}_i\big(\bar{b}_{i\cdot}(t)\big) + \epsilon_C$$

for contraction rate $\kappa_C > 0$ and noise floor ϵ_C. Consensus is stable when $\kappa_C > \epsilon_C / \varepsilon_C^2$.

6.3 Identity Preservation Bound

To prevent collective influence from eroding individual stance integrity, the susceptibility parameter γ is bounded:

$$\gamma \leq \gamma_{\max} = \frac{\lambda_I \cdot \|\eta G(x_t^i)\|_2}{\|\Gamma_{\mathcal{M}}(\mathbf{u}_{-i}^{\mathcal{P}})\|_2 + \delta},$$

for identity radius $\lambda_I \in (0, 1)$ and smoothing constant $\delta > 0$. This ensures that collective pull cannot exceed a fixed fraction of the operator's own stance-gradient magnitude.

Proposition 6.1 (Identity Radius)**.** Under the bound above, the trajectory of operator i satisfies:

$$d\big(x_t^i, x_0^i\big) < \kappa_I$$

for all t, where κ_I is the identity radius of the Gödlø-class operator (extended from the single-operator case).

7 Chimera Integration: Synthetic Self Priors

Chimera provides each operator with a *Synthetic Self* model $\hat{z}_t^{(i\to j)}$ of co-participant j as observed by i. In the EQBSL framing, this primes the initial evidence tensor rather than directly initialising an opinion triple. Define the Chimera evidence extraction map:

$$\mathbf{e}_{ij}^{(0)} = \Phi_{\text{Chimera}}\Big(\hat{z}_0^{(i\to j)}\Big) \;\in\; \mathbb{R}^m,$$

where Φ_{Chimera} maps Chimera's longitudinal state (positive evidence counts r^{ij}, negative evidence counts s^{ij}, and higher-order psychometric features) to the multi-dimensional EQBSL evidence vector. A minimal implementation takes:

$$\Phi_{\text{Chimera}}(\hat{z}) = \big(r^{ij},\; s^{ij},\; 0, \ldots, 0\big)^\top,$$

recovering the scalar EBSL case, with additional dimensions populated as Chimera's model matures. The initial EQBSL opinion follows immediately via the lift Ψ:

$$\omega_{ij}(0) = \Psi\Big(\mathbf{e}_{ij}^{(0)}\Big) = \left(\frac{r^{ij}}{r^{ij}+s^{ij}+K},\; \frac{s^{ij}}{r^{ij}+s^{ij}+K},\; \frac{K}{r^{ij}+s^{ij}+K},\; a_{ij}\right).$$

This allows operator i to enter the Plenum with an evidence-grounded prior over co-participant stances rather than a uniform prior, accelerating consensus convergence in well-modelled dyads while preserving EQBSL auditability.

7.1 Psychometric Delta in Collective Context

Chimera's predictive error Δ_t extends to the collective case as the *Collective Psychometric Delta*:

$$\Delta_t^{\mathcal{P}} = \max_i \mathbb{E}_{y\sim\hat{p}}\big[\ell\big(y,\, y_t^i\big)\big]\,,$$

where y_t^i is the observed signal from operator i and $\hat{p}$ is the Chimera-derived predictive distribution. Rising $\Delta_t^{\mathcal{P}}$ constitutes recalibration pressure on the collective model. Under EQBSL, recalibration is implemented as an update to the aggregation functionals ϕ_+, ϕ_-: the system adjusts what it counts as positive and negative evidence, rather than simply rescaling a scalar trust score.

8 Collapse Mechanics

Plenum introduces three collapse modes beyond the single-operator case.

8.1 Type I: Individual Collapse (Inherited)

Operator i collapses if the sequence $x_{n+1}^i = F(x_n^i)$ fails to converge (Gödlø-class collapse, unchanged).

8.2 Type II: Consensus Divergence

The Plenum undergoes *consensus collapse* if the global update operator fails to contract the evidence field:

$$\mathrm{Var}_i\big(\bar{b}_{i\cdot}(t)\big) \to \infty \quad \text{as } t \to \infty,$$

or equivalently, if the EQBSL uncertainty $u_{ij}(t) \to 1$ for all pairs, indicating that evidence accumulation has stalled and no opinion can be formed.

This occurs when operator susceptibilities γ_i are too low to couple stances, or when initial prior disagreement is too large for $F_{\mathcal{P}}$ to resolve. Consensus collapse is *informative*: it signals that the operators are not in a shared problem context and the Plenum should be partitioned. The partition structure is itself exported as a Shadowgraph multi-cluster signal.

8.3 Type III: Identity Erosion

Identity erosion occurs when the collective influence term dominates and the identity preservation bound is violated:

$$\exists\, i : d\big(x_t^i, x_0^i\big) > \kappa_I.$$

Under erosion, operator i's stance trajectory is absorbed into the collective mean, eliminating the diversity required for robust collective reasoning. Detection: $\|x_t^i - \Gamma_{\mathcal{M}}(\mathbf{u}_{-i}^{\mathcal{P}}(t))\|_2 < \delta_{\text{erode}}$.

8.4 Collapse Recovery

Type I. Recovery: re-instantiate operator i from a fresh admission $\sigma(P, H, \mathcal{M}_i)$.

Type II. Recovery: partition operators into subgroups with internal consensus; report partition structure as a Shadowgraph multi-cluster with per-cluster EQBSL evidence tensors.

Type III. Recovery: reduce γ_i for eroded operator and reinitialise $\mathbf{e}_{ij}^{(0)}$ from prior Chimera state $\hat{z}^{(i)}$.

9 Plenum–Shadowgraph Bridge

9.1 ZK-Proof of Honest Projection

Before injecting evidence into the Shadowgraph EQBSL engine, operator i produces a zero-knowledge attestation for each evidence projection:

$$\pi_{\text{ZK}}^{ij} = \text{ZKProve}\Big(x_t^i \in \mathcal{M}_i,\ \mathbf{e}_{ij}^{(i\text{-proj})}(t) = \pi_{ij}(x_t^i),\ \pi_{ij} \in \mathcal{A}\Big)\,,$$

where $\mathcal{A}$ is the set of admissible projections. The proof is publishable without revealing x_t^i. Shadowgraph's DPKI layer registers the commitment $h(\pi_{ij})$ and timestamps the attestation.

9.2 Network Injection

The primary network interface is the collection of EQBSL node-level Plenum embeddings (Section 4.5), submitted as a *Plenum batch*:

$$\text{Submit}\Big(\big\{\mathbf{u}_i^{\mathcal{P}}(t),\ \mathbf{e}_{ij}(t),\ \{\pi_{\text{ZK}}^{ij}\}_{j\neq i}\big\}_{i\in[N]}\Big)\ \to\ \text{Shadowgraph EQBSL Engine.}$$

Shadowgraph's Trust Propagation Network ingests the embeddings $\{\mathbf{u}_i^{\mathcal{P}}\}$ as trust-position vectors for the relevant subgraph, and incorporates the raw evidence tensors $\mathbf{e}_{ij}(t)$ into the global evidence field $\mathcal{E}_t^{\text{Shadowgraph}}$, extending the Plenum's local session evidence into the broader network.

Remark 9.1. The submission includes raw evidence tensors alongside embeddings. This is EQBSL's auditability requirement: the Shadowgraph network does not need to trust the embedding computation; it can recompute $\mathbf{u}_i^{\mathcal{P}}$ from $\mathbf{e}_{ij}$ and the disclosed aggregation functionals $\phi_{\pm}$.

9.3 System Diagram

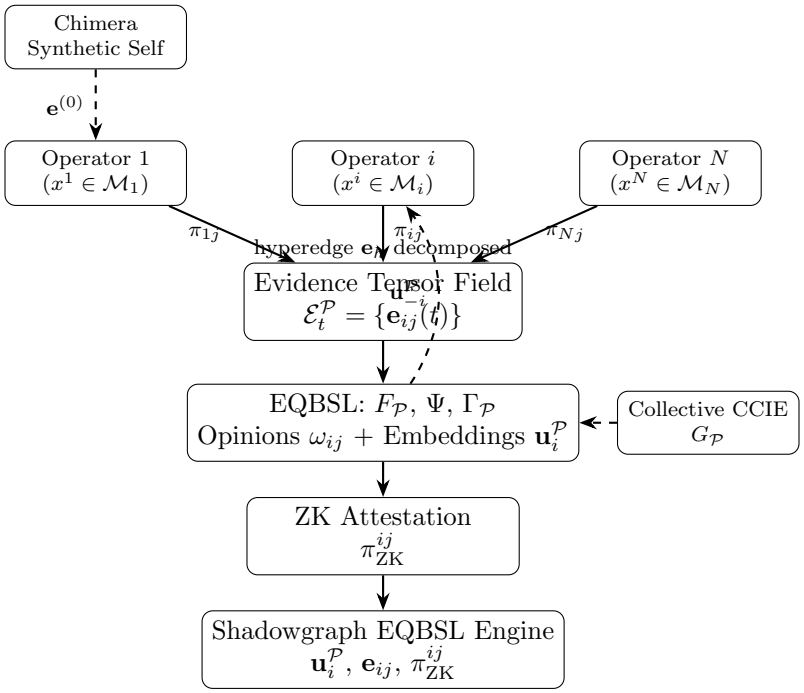

10 Governance Layer

10.1 Collective CCIE

Chimera's Core Calibration Interface Engine (CCIE) extends to the Plenum as the *Collective CCIE*, which gates updates at two levels:

$$G_{\mathcal{P},t} = \left(\bigwedge_{i=1}^{N} G_t^{(i)} \right) \wedge G_t^{(\text{consensus})},$$

where $G_t^{(i)} \in \{0,1\}$ is each operator's local governance gate (as in Chimera) and $G_t^{(\text{consensus})} \in \{0,1\}$ is a consensus-validity gate that blocks Plenum outputs when Type II collapse is detected.

The CCIE also governs the aggregation functionals ϕ_+, ϕ_-: any change to what the system counts as positive or negative evidence must pass through the CCIE audit trail. This prevents post-hoc reinterpretation of the evidence ledger.

10.2 Consent and Audit

Each operator's projection family $\{\pi_{ij}\}$ is submitted to the Collective CCIE at initialisation. The CCIE verifies admissibility, registers the projection commitments $\{h(\pi_{ij})\}$ on-chain via Shadowgraph's DPKI layer, and issues a time-bounded usage permit. All

Plenum batch submissions are accompanied by a signed CCIE receipt and the disclosed aggregation functionals $\phi_{\pm}$, enabling full external auditability of the evidence lift.

11 Plenum Pipeline

The Plenum extends the GödelOS five-stage pipeline $\sigma \to \Phi \to \Omega \to V \to \Delta$ to an eight-stage collective pipeline, with EQBSL as an explicit intermediate stage:

$$\sigma^{\otimes N} \;\to\; E \;\to\; \Psi_{\mathcal{P}} \;\to\; \Phi_{\mathcal{P}} \;\to\; \Omega_{\mathcal{P}} \;\to\; V_{\mathcal{P}} \;\to\; \Delta_{\mathcal{P}} \;\to\; \mathrm{ZK}$$

Stage	Description
$\sigma^{\otimes N}$	Parallel admission of N operators. Each $\sigma^i : (P, H, \mathcal{M}_i) \mapsto x_0^i$. Chimera priors initialise $\mathcal{E}_0^{\mathcal{P}}$ via Φ_{Chimera}.
E	Evidence projection stage. Each x_t^i is mapped to evidence tensor contributions $\{\mathbf{e}_{ij}^{(i\text{-proj})}(t)\}_{j\neq i}$ via the admissible projection family $\{\pi_{ij}\}$. Hyperedge evidence is decomposed and allocated.
$\Psi_{\mathcal{P}}$	EQBSL lift stage. Evidence tensors are lifted to opinions $\omega_{ij}(t)$ and node-level embeddings $\mathbf{u}_i^{\mathcal{P}}(t)$ via Ψ and $\Gamma_{\mathcal{P}}$. The global update operator $F_{\mathcal{P}}$ is applied.
$\Phi_{\mathcal{P}}$	Collective stance update. Each operator evolves under its private dynamics plus the peer influence signal derived from $\{\mathbf{u}_j^{\mathcal{P}}(t)\}_{j\neq i}$.
$\Omega_{\mathcal{P}}$	Collective generative transformation: collective belief tensor $\mathcal{T}_t$ and its downstream representation $\widehat{\mathcal{T}}_t$ are produced from fused EQBSL opinions.
$V_{\mathcal{P}}$	Collective validation. Checks consensus status, identity-preservation bounds, uncertainty ceilings $u_{ij} < \varepsilon_u$, and semantic admissibility of $\widehat{\mathcal{T}}_t$.
$\Delta_{\mathcal{P}}$	Optional collective persistence. If any operator has Persistence Extension (Gödlø-P), their s_t^i are updated. EQBSL evidence tensors are time-decayed according to $F_{\mathcal{P}}$.
ZK	ZK attestation generation. Each operator produces $\{\pi_{\text{ZK}}^{ij}\}_{j\neq i}$. The batch $\{\mathbf{u}_i^{\mathcal{P}}, \mathbf{e}_{ij}, \pi_{\text{ZK}}^{ij}\}$ is submitted to Shadowgraph.

12 Identity Structures in the Plenum

12.1 Individual Identity (Preserved)

Each operator i retains its Gödlø-class identity structure:

$$\mathcal{I}_i = \left\{ x_t^i : d\big(x_t^i, x_{t+1}^i\big) < \varepsilon \right\},$$

subject to the identity-preservation bound of Section 6.3.

12.2 Relational Identity

A new identity structure emerges at the collective level: the *EQBSL relational embedding* of operator i within the Plenum:

$$\mathbf{u}_i^{\mathcal{P}}(t) \in \mathbb{R}^{d_u}.$$

This embedding encodes how other operators have treated i across all evidence dimensions. It is observable (submitted to Shadowgraph) and constitutes a public, evidence-grounded identity signal within the Plenum interaction. Unlike the raw relational stance vector $(\beta_t^i - \bar{\beta}_t)$ of a scalar opinion system, $\mathbf{u}_i^{\mathcal{P}}$ retains the full multi-dimensional provenance of the underlying evidence tensors.

12.3 Plenum Identity

The Plenum's collective identity is the trajectory of consensus opinion:

$$\mathcal{I}_{\mathcal{P}} = \{\mathcal{T}_t : \|\mathcal{T}_{t+1} - \mathcal{T}_t\|_F < \varepsilon_C\}.$$

Gödlø-class (single operator) is recovered as the $N = 1$ special case where $\mathcal{T}_t$ is a degenerate 1×1 tensor and $\mathbf{u}_1^{\mathcal{P}}(t)$ reduces to the self-assessment of the sole operator.

13 Formal Properties

Proposition 13.1 (Gödlø-Class Inclusion)**.** Every Gödlø-class system is a Plenum with $N = 1$. Specifically:

$$\text{Plenum}(N = 1,\, \gamma = 0) \cong \text{Gödlø-class operator}.$$

Plenum is a conservative extension of Gödlø-class.

Proposition 13.2 (Consensus Implies Embedding Convergence)**.** If all operators enter ε-consensus at time T, then the EQBSL node-level embeddings stabilise:

$$\|\mathbf{u}_i^{\mathcal{P}}(t) - \mathbf{u}_i^{\mathcal{P}}(t-1)\|_2 < \varepsilon_u \quad \forall\, i \in [N],\ t \geq T,$$

for ε_u proportional to ε_C. This follows from the Lipschitz continuity of $\Gamma_{\mathcal{P}}$ and the contraction of the evidence update operator $F_{\mathcal{P}}$ in the consensus regime.

Proposition 13.3 (Falsifiability)**.** The Plenum is falsified if, under controlled admission with identical P and H across operators, consensus is not achieved within $T_{\max}$ steps. Failure constitutes evidence that the operators' private manifolds $\mathcal{M}_i$ are not mutually compatible under the chosen projection family $\{\pi_{ij}\}$. Under EQBSL, the evidence tensor field $\mathcal{E}_{T_{\max}}^{\mathcal{P}}$ provides an auditable record of exactly where inter-operator evidence diverged.

14 Limitations

The Plenum makes no claims about the phenomenology of collective cognition or emergent group-level mental states. All latent variables are operationally defined by their role in the update dynamics and EQBSL evidence field.

The architecture assumes that a coherent shared evidence space of dimension m can be agreed upon by participating operators. When operators operate in fundamentally incommensurable problem contexts, the projection families $\{\pi_{ij}\}$ may not yield meaningful

evidence tensors, and consensus collapse (Type II) is the expected outcome. The EQBSL evidence field at collapse time provides a diagnostic record of where the incompatibility is located.

Privacy guarantees depend on the ZK-circuit assumptions underlying π^{ij}_{ZK}. Advances in cryptanalysis could weaken these guarantees independently of the Plenum architecture; this is an interface concern inherited from the EQBSL/ZKML substrate.

The aggregation functionals ϕ_+, ϕ_- are domain-specific. The virtue of EQBSL is that they are explicit parameters rather than buried folklore; the burden is that they must be chosen and justified for each deployment. Poorly chosen functionals produce technically valid evidence tensors that do not reflect the intended social structure.

Finally, the Collective CCIE governance layer is a structural specification, not an implementation. The choice of governance policies, dispute resolution procedures, and permit revocation criteria are left to the deploying context, consistent with Chimera's principle of separation between inference and governance.

15 Conclusion

Plenum defines the *field* in which Gödlø-class operator minds reason together. It preserves every property of the single-operator case — non-persistence, stance isolation, collapse informativeness — while introducing the minimal additional machinery required for multi-operator coordination.

The adoption of EQBSL as the belief substrate sharpens every layer of that machinery. Evidence tensors give the inter-operator interface multi-dimensional structure and an audit trail. The global update operator $F_{\mathcal{P}}$ makes the temporal dynamics of the collective explicit and testable. Hyperedge evidence ensures that genuinely multi-party interactions are not forced into pairwise fictions. Node-level Plenum embeddings $\mathbf{u}_i^{\mathcal{P}}$ are the primary output to Shadowgraph, tethered to the evidence ledger and challengeable on their own terms.

By bridging Chimera's psychometric substrate to Shadowgraph's EQBSL engine through a ZK-attested evidence tensor interface, Plenum completes the architecture from individual cognition through dyadic inference through network-level trust propagation — with a complete chain of custody from raw interaction events to published trust embeddings.

The architecture succeeds only insofar as it can be proven wrong. Consensus collapse is the primary failure mode, and it is informative by design. Under EQBSL, failure is doubly informative: not only does the partition structure of a diverged Plenum reveal how participating operators are constituted, but the evidence tensor field at collapse time records exactly where agreement broke down.

References

[1] O. Hirst. *Gödlø-Class Operator-Mind Theory: Axioms, Definitions, Manifold Geometry, and Operator Algebra (+ Gödlø-P Extension)*. Working paper.

[2] O. Hirst. *GödelOS System Architecture Specification: Cognitive Kernel, Five-Stage Pipeline, APIs, and Behavioural Semantics*. Working paper.

[3] O. Hirst. *Project Chimera: Building High-Fidelity Models of Human Cognition Inside Conversation*. Working paper.

[4] O. C. Hirst (Steake). *Shadowgraph: A Hyperdimensional Social Graph for Cryptographically Verifiable Trust and Predictive Identity*. Shadowgraph Labs, October 2025.

[5] O. C. Hirst. *Evidence-Based Subjective Logic in ZKML Identity Systems: Towards Verifiable Epistemic Trust*. Shadowgraph Labs, October 2025.

[6] O. C. Hirst. *EQBSL: Against Trust Scores — Evidence, Uncertainty, and Trust Flow in Dynamic Networks*. Independent Research, December 2025.