Article

Not peer-reviewed version

Morphology and the Structural Preconditions of Basin Formation

Gabriel Axel Montes *

Posted Date: 7 April 2026

doi: 10.20944/preprints202604.0435.v1

Keywords: morphological participation index; basin formation; quantum Darwinism; prototime; spectral integration; architectural prior; witness redundancy; seam topology; temporal scaffolding; dissociation logic

Article

Morphology and the Structural Preconditions of Basin Formation

Gabriel Axel Montes

Neural Axis; ASI Institute; Center for the Future of AI, Mind, and Society, Florida Atlantic University; gabriel@neuralaxis.org

Abstract

The Quantum Darwinist Theory of Consciousness (QDT) and the Prototime Interpretation (PT) characterize localized conscious basins in terms of spectral integration, PT-participation, recursive coherence, witness redundancy, and temporally ordered record formation [7–9]. An upstream question has remained largely implicit in that program: what sort of morphology makes such dynamics structurally plausible? This paper argues that morphology constitutes the architectural precondition of basin formation, and that the Morphological Participation Index (MPI; Montes 5) can make that precondition operational. Architecture, realized integration, carrier structure, and witnessed temporality each answer a different question about the same candidate system. MPI contributes the architectural prior: it localizes where balanced seams lie, where redundant trace or witness surfaces are available, where carrier-sensitive assays are worth running, and where temporally thick, record-supported basins—integrated regimes whose stability depends on redundant internal records or traces that persist across behaviorally relevant time windows—are plausible. A structural factorization of candidate basins connects MPI's score bundle to downstream Φ_s, PT-participation, and clock indices. Expected dissociations—cases where high MPI coexists with low realized integration, or where trace-rich architectures lack the carrier geometry for PT-participation—sharpen experimental design and help distinguish genuine basin formation from structural mimics. The result is a bridge from morphology to the empirical core of QDT/PT, grounded in the same balanced-cut spectral formalism that underlies Φ_s itself.

Keywords: morphological participation index; basin formation; quantum Darwinism; prototime; spectral integration; architectural prior; witness redundancy; seam topology; temporal scaffolding; dissociation logic

1. Introduction: Why an Architectural Layer for QDT/PT

The QDT/PT research program [7–9] asks which emergent basins qualify as localized conscious subjects. The answer involves a conjunction of conditions: spectral integration (Φ_s) that resists balanced decomposition [1], PT-participation anchored in low-entropy coherence-sensitive carriers, recursive self-maintenance, redundant witness formation in the Quantum Darwinist sense [13,14], and temporal thickness sufficient to sustain a metastable regime over the relevant psychophysical integration window. Each of these conditions has been developed in recent work. What has remained less explicit is the role of architecture—the constraint geometry that determines whether such conditions are even structurally available.

Architectures differ in where their weakest seams lie, how their trace and witness surfaces are distributed, whether their geometry can protect organized carrier modes, and whether their temporal organization spans the range needed for thick basins. These are morphological questions, and they arise before any dynamical or carrier-sensitive assay is brought to bear. A system whose morphology lacks the seam structure for high Φ_s, or whose trace surfaces are too concentrated for redundant witnessing, or whose timescale ladder is too shallow for temporal thickness, is unlikely to satisfy the QDT/PT conjunction regardless of its microphysics.

MPI is developed here as a way to make that architectural question precise. The full MPI formalism—constraint hypergraphs, balanced-cut seam diagnostics, subscore definitions, calibration conventions, and toy benchmarks—is presented in Montes [5]. The present paper takes that formalism as given and asks what role an architectural prior should play within QDT/PT, and what bridge hypotheses, dissociations, and empirical consequences follow.

We begin by distinguishing four analytically separable layers—architecture, realized integration, carrier structure, and witnessed temporality—and locating MPI within that framework (Section 2). We then develop the bridge itself: how seam topology, witness redundancy, carrier geometry, temporal scaffolding, and local Booleanization connect MPI's outputs to QDT/PT (Section 3), and how candidate basins can be structurally factorized before downstream assays are applied (Section 4). Finally, we identify comparison classes and expected dissociations that follow from keeping these layers apart (Section 6).

2. Division of Labor Across Architecture, Dynamics, Carrier Structure, and Witnessed Temporality

QDT/PT implicitly distinguishes several levels of analysis. Making them explicit clarifies both what MPI contributes and where its scope ends.

Table 1. Division of labor across architecture, dynamics, and carrier structure. Each layer answers a different question about the same candidate system.

Layer	Question	Representative object
Morphology	What architectural resources are available for coupling, partitioning, trace registration, and temporal organization?	MPI [5]
Realized integration	Does the realized functional graph resist balanced decomposition over a specified window?	Φ_s [1]
Carrier structure	How much of that realized integration is borne by coherence-sensitive or low-entropy carriers?	PT-participation [9]
Witnessed temporality	Are integrated patterns redundantly recorded and stabilized into temporally thick basins?	Witness, record, and temporal-basin diagnostics [7,9]

"Participation" in this context means that morphology actively constrains which unified regimes can be formed, stabilized, and recovered after perturbation. A morphology participates when its seam structure, trace surfaces, recurrence loops, and timescale ladder shape which basins are accessible—not merely which basins happen to be realized at a given moment.

Systems can separate sharply along these layers, and those separations are empirically informative. A telemetry-rich infrastructure can build extensive trace surfaces and durable records while remaining entirely classical—high trace geometry, potentially high temporal scaffolding, but no carrier-sensitive integration. Conversely, a system may contain local coherence-sensitive structure yet lack the trace geometry or temporal scaffolding required for stable basin formation. MPI is designed to make such separations visible before downstream assays are applied.

These structural features also bear on system-level control—whether an architecture can sustain persistence across timescales, trace support sufficient for state registration, seam geometry that permits coordinated action without fragmentation, and multiscale organization that allows local processes to cohere into durable system-level regimes [11]. In that limited sense, the framework doubles as an agency-relevance prior.

doi:10.20944/preprints202604.0435.v1

3. Morphology as the Architectural Layer of Basin Formation

Each of the following subsections takes one dimension of the QDT/PT conjunction and asks what MPI contributes to it architecturally.

3.1. Seam Topology and Spectral Integration

Φ_s quantifies how strongly a candidate basin resists being informationally decomposed into independent parts [1]. In QDT, high Φ_s is one of the conditions on localized conscious basins [9]. MPI's integration component (s_{int}) imports the same balanced-cut mathematics one level upstream: it asks which morphologies contain seam structures hospitable to basin-level irreducibility before one has established that the relevant dynamics are present.

Conductance and normalized-cut objectives [3,6,10,12] penalize trivial separations and recover the standard spectral relation between weak seams and the second Laplacian eigenvalue. In MPI, the resulting seam maps are candidate fault lines for basin formation—places where a putative unified regime should fragment first under lesion, overload, or dynamical perturbation. If MPI's seam maps are doing real work, then seam lesions targeted by MPI should produce sharper Φ_s losses than lesions at control locations with matched edge weight but different seam status.

3.2. Witness Redundancy and Record Proliferation

In QDT, the mechanism by which basin identity becomes stable and objectively recoverable depends on redundant witnessing. Following the Quantum Darwinist account of classical objectivity [13,14], information about a system's pointer state must be redundantly recorded across multiple environmental fragments for that state to be intersubjectively accessible. In the Schneider–Bailey development, internal neural and microtubular populations function as witnesses to a common PT-coupled pattern, with the redundancy requirement contributing record-stability and the recursion requirement contributing temporal thickness and self-maintenance [9].

MPI's trace component (s_{trace}) captures the architectural side of that requirement. It asks how many independent paths connect core dynamics to trace surfaces and how evenly the burden of registration is distributed, without presupposing which dynamical patterns are actually being recorded. Concentrated trace geometry predicts choke points, selective registration failure, and basin brittleness; distributed trace geometry predicts more redundancy-rich support. If this matters empirically, systems with higher s_{trace} should exhibit stronger witness redundancy, more stable record-bearing dynamics, and more recoverable basin readouts once downstream assays are applied. In non-biological comparison cases, externalized records—logs, telemetry channels, cached state, pheromone-like fields [2]—play the same formal role and help separate trace richness from basin unity.

3.3. Resonant-mode Support and Carrier Geometry

QDT treats consciousness as depending not only on integration but on the kinds of modes that carry it. PT-participation measures how deeply a basin's integration is anchored in low-entropy, coherently coupled modes that retain coupling to the prototemporal substrate [9]. The geometry that makes high PT-participation physically realizable includes structural features such as confinement, frustrated lattices, protected channels, and cavity-like architectures.

MPI's resonant-mode component (s_{res}) is the architectural counterpart of this carrier question. At lower tiers it uses structural proxies—cycle-richness, edge-connectivity, phase-consistency—to identify architectures that appear capable of protecting organized carrier modes. At higher tiers it absorbs direct phase, coherence, or quality-factor evidence [4]. The testable prediction is that variation in s_{res}, within families already showing comparable s_{int} and Φ_s, should predict the outcome of coherence- or resonance-sensitive perturbations—phase disruption, carrier-specific anesthetic perturbations, or other interventions targeted at the geometric carriers singled out by the relevant substrate.

3.4. Time, Records, and Local Booleanization

In the Prototime Interpretation, classical time emerges through decoherence, redundancy accumulation, and the stabilization of record-bearing sectors from a deeper prototemporal substrate [7,8]. QDT extends that picture: a conscious "moment" corresponds to a metastable high-Φ_s basin sustained over the relevant psychophysical integration window, within which multiple faster microstates and oscillatory cycles are nested [9].

MPI addresses the architectural side of this temporal question through two components. The temporal scaffolding component (s_{temp}) asks whether the morphology supports a nontrivial spread of internally relevant timescales and durable record loops rather than a single thin clock. The optional contextual patchiness module (s_{ctx}) asks whether recordable predicates form substantial local compatibility patches without collapsing into one trivially global Boolean order. Read together, these outputs give an architectural handle on local Booleanization: where stable contexts might form, how thick or thin they may be temporally, and whether sector boundaries align with seam structure. If temporal scaffolding matters for basin persistence, systems with higher s_{temp} should exhibit more durable clock signatures, longer dwell times in high-Φ_s plateaus, and more robust record-supported basins.

4. Structural Factorization of Candidate Basins

Let $B_{\sigma,t}$ denote a candidate basin specified at scale σ and time window t. Downstream QDT/PT analysis asks whether the realized dynamics on $B_{\sigma,t}$ exhibit high Φ_s, high PT-participation, and sufficient clock or record persistence. MPI contributes a structural prior at the same scale and window. Conceptually,

$$\pi_H(B_{\sigma,t}) := F(s_{\text{int}}, s_{\text{multi}}, s_{\text{res}}, s_{\text{trace}}, s_{\text{temp}}, s_{\text{rob}}, s_{\text{ctx}}),$$

with s_{ctx} included only when the contextual module is defined, and the research decomposition takes the form

$$\Pr(B_{\sigma,t} \text{ is QDT/PT-relevant} \mid H, D) \;\propto\; \pi_H(B_{\sigma,t})\, \mathcal{E}(\Phi_s(B_{\sigma,t}), P(B_{\sigma,t}), C_{\text{clock}}(B_{\sigma,t}) \mid D),$$

where D denotes available dynamical data, P denotes PT-participation, C_{clock} denotes a clock or record-persistence index, and $\mathcal{E}$ is an evidential term summarizing the dynamical and carrier-sensitive evidence.

This expression is a research decomposition rather than a full Bayesian model. Neither π_H nor $\mathcal{E}$ is a calibrated quantity at this stage; the decomposition formalizes the program's division of labor and makes explicit which questions belong to architecture and which to downstream dynamics. Architecture and realized dynamics answer different questions and should be reported separately. MPI turns the architectural side into explicit objects—seam partitions, trace maps, timescale diagnostics, and robustness reports—that can guide where Φ_s, PT-participation, and clock indices are measured in the first place.

Table 2 summarizes how MPI outputs enter this decomposition.

doi:10.20944/preprints202604.0435.v1

Table 2. How MPI outputs enter the QDT/PT research decomposition, with example downstream tests connecting architectural predictions to dynamical and carrier-sensitive evidence.

MPI output	**Role in QDT/PT**	**Example downstream test**
Seam partitions, s_{int}	Candidate unity topology; localizes balanced fault lines along which integrated basins are most likely to fragment	Compare seam lesions with changes in Φ_s
s_{multi}	Architecture for recursive re-instantiation across scales and nested basin stabilization	Test whether coarse-grainings with higher seam persistence support more durable basin structure
s_{res}	Structural support for carrier modes that remain coherence-sensitive near the classical boundary	Compare with PT-participation proxies or phase-sensitive assays
Trace maps, s_{trace}	Distribution of trace surfaces for redundant record instantiation	Test record redundancy, trace concentration, and basin-specific binding signatures
s_{temp}	Architecture for nested timescales, clock-like refresh, and temporally thick basins	Compare with clock indices, dwell times, or plateau durations
s_{ctx}	Local Booleanization and context boundaries within a broader non-Boolean setting	Test whether compatibility patches align with local record sectors
s_{rob}	Stability of the structural profile under nuisance perturbation	Test whether predicted basins survive modest lesions, noise, or relabeling

In practice, this means MPI outputs can *select* candidate basins before computationally expensive downstream analyses are run. Rather than computing Φ_s or estimating PT-participation across an entire system at every scale and window, one can first identify the scales and regions where the architectural profile is most favorable—where seams are hardest to cut, traces are most redundant, and temporal breadth is widest—and then focus dynamical assays on those candidates.

5. MPI as an Architectural Prior: A Compact Summary

The full MPI framework is developed in Montes [5]. For the convenience of readers encountering the framework here, a brief summary follows.

MPI represents morphology as a weighted constraint hypergraph $H = (V, E, w, \chi)$ or as an explicit multilayer family $\mathcal{H} = \{H^{(\ell)}\}$. It returns a score bundle $\text{MPI}(S) = (\mathbf{s}, \Pi, \mathcal{M}, s_{\text{ctx}})$, where $\mathbf{s} \in [0,1]^6$ is a vector of core subscores, Π is a multiscale family of seam partitions, and $\mathcal{M}$ collects seam maps, trace maps, and related diagnostics. The six core subscores are:

Table 3. The six core MPI subscores. Full definitions, calibration conventions, and implementation details appear in Montes [5].

Subscore	**What it measures**
s_{int}	Resistance to balanced decomposition (via Fiedler-vector sweep cuts)
s_{multi}	Persistence of seam structure across coarse-graining levels
s_{res}	Structural support for coherent or resonance-sensitive modes (tiered)
s_{trace}	Redundancy and distribution of trace or interface surfaces
s_{temp}	Breadth of the morphology's timescale ladder
s_{rob}	Stability of the profile under perturbation and relabeling

The integration component is anchored in balanced-cut spectral decomposition: it uses the normalized Laplacian, the Fiedler vector, and conductance-minimizing sweep cuts, aligning with the same spectral logic that underlies Φ_s [3,12]. An optional scalar roll-up $\text{MPI}_{\text{tot}} = \prod s_k^{\alpha_k}$ is available but secondary to the profile. An optional contextual patchiness module (s_{ctx}) captures local Booleanization when a defensible predicate family is available.

The primary output is the structural profile together with the accompanying maps and partitions. For QDT/PT applications, the most important outputs are seam maps (which localize candidate basin boundaries), trace maps (which identify where redundant registration can accumulate), and the multiscale partition hierarchy (which reveals where recursive re-instantiation is architecturally supported).

6. Comparison Classes and Expected Dissociations

If architecture, realized integration, carrier structure, and witnessed temporality are genuinely distinct layers, then they should be empirically dissociable. A well-designed benchmark suite should include cases that score high on some layers and low on others, producing informative contrasts rather than trivial correlations.

6.1. Expected Architectural Families

Four families provide a minimal comparison set (see Montes 5 for detailed benchmark construction):

1. **Centralized or controller-like architectures.** Strong global integration (s_{int}), relatively concentrated traces, limited multiscale depth. These are natural conjunction tests: if Φ_s is also high, the architectural and dynamical layers agree; if Φ_s is low despite high s_{int}, the realized dynamics fail to exploit the available architecture.
2. **Federated or cephalopod-like distributed architectures.** Weak global integration but strong multiscale nesting (s_{multi}), distributed traces, and broader temporal scaffolding. These can dissociate multiscale nesting from globally concentrated integration—a separation that matters for understanding whether recursive re-instantiation requires global unity or can be sustained through modular coordination.
3. **Stigmergic or collective architectures.** Distributed coordination through externalized traces and low-bandwidth local rules [2]. These provide classical control cases: they may be trace-rich and temporally persistent without exhibiting the carrier geometry or the balanced-seam structure associated with the full QDT/PT conjunction.
4. **Telemetry-rich technical infrastructures.** High observability, rich logging, often strong trace geometry but mixed internal unity. These stress-test the distinction between coordination surfaces and deeper integrative architecture, and help separate witness richness from basin unity.

6.2. Key Dissociations

The layered framework predicts several empirically testable dissociations:

1. **High MPI, low realized Φ_s.** An architecture with strong seam structure, distributed traces, and broad temporal scaffolding may nonetheless fail to realize high spectral integration—because the dynamics are too noisy, too weakly coupled, or too far from the relevant attractor regime. This dissociation would confirm that architecture sets an envelope rather than guaranteeing a particular dynamical outcome.
2. **High s_{trace}, low s_{res}.** A trace-rich architecture whose geometry does not support protected carrier modes. Such a system might satisfy the witness-redundancy requirement of QDT while failing the PT-participation requirement. Stigmergic collectives and telemetry-rich infrastructures are natural candidates.
3. **High s_{int} and s_{res}, low s_{temp}.** An architecture that supports both integration and carrier geometry but lacks temporal breadth. This would produce basins that are spectrally integrated and carrier-sensitive but temporally shallow—punctiform transitions rather than the thick, metastable regimes that QDT associates with conscious episodes.

4. **High s_{multi}, low s_{int}.** Strong multiscale nesting without strong global integration. This is the federated-architecture case: stable modular structure at every resolution but easy global cuts. It tests whether recursive re-instantiation can be sustained through federated coordination or whether it requires globally concentrated integration.

These dissociations matter regardless of how the QDT/PT conjunction is ultimately evaluated, because they identify which architectural dimensions vary independently and which tend to co-occur.

7. Future Directions

A plausible future program has four tracks.

Track 1: architectural screening.

Compute MPI on a benchmark family of structurally diverse systems in structural-only and, where data permit, structural-plus-dynamic mode. MPI should yield distinct score profiles and seam maps across architectural families, and those distinctions should not collapse into trivial correlates of size, density, or degree heterogeneity. Within the present program, this same step identifies which systems are plausible candidates for downstream Φ_s, PT-participation, and clock/record analysis, and which are better treated as control cases.

Track 2: structural–spectral link.

Test whether higher s_{int} and stronger seam persistence predict stronger realized Φ_s in measured data, by comparing structural MPI diagnostics with realized Φ_s from neural, organoid, simulation, or other time-series sources. The hypothesis is that morphology constrains the envelope of accessible basin-level integration: the architectural profile should indicate where strong spectral integration is easiest to stabilize and hardest to preserve under perturbation.

Track 3: carriers, traces, and temporal organization.

Test the remaining bridge claims directly. Does higher s_{res} predict stronger PT-participation proxies or other coherence-sensitive readouts? Do higher s_{trace} and s_{temp} predict stronger witness redundancy, more stable record-bearing dynamics, or more durable clock signatures? Does s_{ctx} align with local compatibility sectors rather than with arbitrary graph partitions? Each of these bears on whether the architectural profile can localize where carrier-sensitive integration, redundant registration, and temporally thick basin structure are likely to appear.

Track 4: lesion and intervention studies.

Use MPI outputs to localize interventions. Lesioning seam edges should produce sharper Φ_s losses than control lesions at matched edge weight. Perturbing trace surfaces should degrade record stability preferentially where trace geometry is concentrated. Altering phase structure should degrade PT-participation proxies in systems with otherwise similar seam topology. Varying coarse-graining choices should reveal whether predicted sectors persist. If MPI is to function as the architectural layer of the QDT/PT program, its components must track the right failures when the architecture is deliberately degraded.

8. Limitations

MPI is a structural framework, and its contribution to the QDT/PT program is correspondingly bounded. Morphology can be over-read: some systems will be dominated by microphysical details that a coarse hypergraph fails to capture. The bridge may fail if the true carrier story is largely microphysical and only weakly constrained by the morphology that MPI captures—in which case the architectural prior would be too coarse to usefully localize where PT-participation is likely.

A high s_{res} remains a structural proxy unless higher-tier measurements are available, and strong trace geometry does not by itself show that the recorded pattern is a PT-participating one. Calibration matters: the sigmoid transforms and weighting conventions are placeholders until benchmark corpora are broad enough to support principled tuning (see Montes [5] for the calibration protocol). The contextual patchiness module depends on a defensible predicate family and remains more modeling-sensitive than the core bundle.

The framework will naturally generate boundary cases: systems that are trace-rich yet weakly integrated, or integration-strong yet temporally shallow, or carrier-favorable yet lacking witness redundancy. These are not failures of MPI. They are the conditions under which keeping architecture, realized dynamics, carrier structure, and temporal organization apart does the most empirical work.

9. Conclusion

The QDT/PT program stands to gain from an explicit architectural layer. Φ_s, PT-participation, and temporal-basin diagnostics tell us whether particular basins are realized, carrier-coupled, and temporally thick—but they leave open whether the morphology could have supported such basins in the first place. MPI addresses that gap with seam maps, trace maps, multiscale partitions, and temporal-breadth diagnostics that can direct where downstream assays are applied and how their results are read.

What this paper contributes is the bridge: the structural factorization of candidate basins, the four-layer decomposition, and the dissociations between those layers. Together these give a way to sharpen experimental design rather than leaving architectural questions as unexamined background assumptions. The payoff is more targeted empirical work—not just knowing that a basin has high Φ_s, but knowing *why* the architecture makes that value accessible and *where* it is most likely to break.

A natural direction for future work is to put these predictions to use: select candidate basins using MPI's structural outputs and ask whether the signatures of spectral integration, PT-participation, and temporal stabilization appear where the architecture says they should.

References

1. Mark Bailey and Susan Schneider. *When Wholes Resist Decomposition: A Spectral Measure of Epistemic Emergence*. PhilArchive, 2025.
2. Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.
3. Fan R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.
4. Michaël Fanuel, Carlos M. Alaíz, Ángela Fernández, and Johan A. K. Suykens. Magnetic eigenmaps for the visualization of directed networks. *Applied and Computational Harmonic Analysis*, 44:189–199, 2018.
5. Gabriel Axel Montes. Morphology, seam topology, and temporal scaffolding in complex systems. Preprint, 2026.
6. M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.
7. Susan Schneider and Mark Bailey. The prototime interpretation of quantum mechanics. In Dean Rickles and Hatam Elshatlawy, editors, *Quantum Gravity and Computation: Information, Pregeometry, and Digital Physics*. Routledge, forthcoming.
8. Susan Schneider and Mark Bailey. Superpsychism. *Journal of Consciousness Studies*, 33(1):13–39, 2026. doi: 10.53765/20512201.33.1.013.
9. Susan Schneider and Mark Bailey. The quantum Darwinist theory of consciousness: Resonance, space-time emergence, and a metric for what makes something conscious: A reply to critics. *Journal of Consciousness Studies*, 33(1):278–341, 2026. doi: 10.53765/20512201.33.1.278.
10. Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.
11. Herbert A. Simon. The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6):467–482, 1962.
12. Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. doi: 10.1007/s11222-007-9033-z.
13. Wojciech H. Zurek. Quantum Darwinism. *Nature Physics*, 5(3):181–188, 2009.

14. Wojciech H. Zurek. Quantum theory of the classical: quantum jumps, Born's rule and objective classical reality via quantum Darwinism. *Philosophical Transactions of the Royal Society A*, 376(2123):20180107, 2018.