ATLAS PAPER II

Shared Parity Networks within a Datum-Centered Curvature Witness Engine

Project-position revision for solver-facing continuity

Date: April 9, 2026

Status: Revised draft reflecting the datum graph, witness-cloud, overlap-comparison, and parity-extraction architecture.

> **Project-position note.** This paper now defines Atlas's canonical parity-network extraction regime inside a broader datum-centered curvature witness engine. Older parity-only framings are retained where still useful, and demoted where they overconstrain development.

ABSTRACT

Paper I defined the Gravitational Coherence Surface (GCS) as an operational boundary of gravitational distinguishability: the locus at which a chosen source-readable signal reaches parity with a calibrated competing floor. Paper II originally extended that idea to genuinely many-body settings, showing that directional parity among several comparable sources induces not an isolated shell but a network of balanced interfaces.

That core result remains valid, but Atlas has since developed into a broader engine architecture. The solver is no longer best understood as only a parity extractor. It is now better framed as a datum-centered curvature witness engine: a system that constructs root datums, assigns node-specific witness clouds, evaluates weak-field general-relativistic tidal structure on those clouds, compares overlapping reports from distinct nodal vantage points, and then extracts parity interfaces, exposed faces, seams, and related geometric artifacts from that richer comparison layer.

This revision updates Paper II to reflect that change in emphasis. The shared parity network remains an important theoretical and computational object, but it is no longer treated as the sole or ultimate description of Atlas. It is instead the first canonical extracted artifact family within a larger nodal curvature-negotiation stack. In that stack, Node 0 serves as a geometric witness rooted at the unweighted center of configuration, Node 1 serves as a reference or floor datum, Node 2 serves as the mass-weighted realized center, and physical source nodes contribute the actual curvature content. Hierarchy

organizes computation, but local competition is decided by direct field comparison rather than inherited authority.

The first operational extraction regime remains intentionally narrow. Atlas Milestone 1 still privileges the weak-field point-mass tidal proxy $Q_i \sim M_i / r_i^3$ as the canonical readable field for stable extraction of pairwise balance support, exposed faces, and triple-junction structure. But the runtime architecture is now expected, where practical, to carry the local tidal tensor itself and to treat cloud overlaps, center diagnostics, and artifact persistence under perturbation as first-class elements of the solver.

The resulting picture is more disciplined and more fertile. Atlas is not proposed as a modification of general relativity, nor is the shared parity network claimed as a fundamental invariant. Rather, Atlas is a structured weak-field instrument for interrogating relational curvature organization. Within that instrument, the parity network is the first stable geometric language, not the final ontology.

Keywords: *Atlas Solver, weak-field GR, tidal tensors, witness clouds, overlap comparison, parity networks.*

1. PURPOSE AND REVISED STANCE

The purpose of this paper is no longer merely to give the many-body extension of Paper I. It is to state clearly what portion of Atlas Paper II still governs, what has been promoted upstream into engine architecture, and what has been demoted so that development is not constrained by an aging framing.

Paper I asked where a source-readable signal reaches parity with a calibrated floor. Early Paper II asked what manifold is implied when several comparable sources compete under the same parity logic. That question remains important, but Atlas now begins one level earlier. Before parity extraction, the solver constructs root datums, deploys witness clouds, evaluates weak-field curvature-relevant structure, and compares overlapping reports from multiple nodal vantage points. Only after that comparison layer is built does parity extraction become one canonical artifact family.

The stance remains conservative:

- standard GR only;
- no new field equations;
- no claim that derived interfaces are fundamental invariants;
- no conflation of distinguishability with force balance or orbital stability;
- no permission for elegant metaphor to outrun extraction stability.

What has changed is the bookkeeping geometry and the order of operations. Paper II therefore survives, but with narrower jurisdiction. It no longer stands as the full conceptual map of Atlas. It defines the canonical parity extraction regime inside a broader solver.

2. PROJECT POSITION OF PAPER II

This paper should now be read as a middle-layer document.

Upstream of it lies the datum and witness architecture: configuration intake, root construction, node typing, cloud assignment, frame choice, curvature evaluation, and overlap bookkeeping.

Inside it lies the first canonical extraction language: pairwise balance support, exposed faces, triple-junction structure, and the many-body shared parity network in the weak-field point-mass tidal regime.

Downstream of it lie visualization, artifact certification, comparative runs, and interpretive extension.

In compact form, the Atlas pipeline is now:

Configuration -> Datum graph -> Node-specific witness clouds -> Weak-field tidal evaluation -> Overlap comparison layer -> Canonical parity extraction -> Artifact certification and archive

This paper therefore no longer claims that the shared parity network is the whole machine. It claims instead that parity extraction is the first stable and analytically transparent artifact family produced by the machine.

3. ROOT DATUMS AND NODAL HIERARCHY

Atlas now begins with a rooted nodal construction on a chosen time slice Sigma_t.

Node 0: geometric root datum. Node 0 is the unweighted geometric center of the active source configuration. It does not generate curvature. It is a Platonic witness datum used to interrogate how the realized configuration departs from ideal geometric form.

Node 1: reference or floor datum. Node 1 represents a reference condition against which readable signals may be compared. In early runs it may be implemented as a constant floor or other explicit reference rule. It should not be forced to masquerade as a physical source if doing so muddies the ontology.

Node 2: barycentric datum. Node 2 is the mass-weighted realized center of the configuration. Like Node 0, it is not a source of independent curvature in the solver bookkeeping. It is a witness datum, but one rooted in realized mass distribution rather than in pure geometry.

Node 3+ : physical source nodes. These are the declared masses with positions, velocities, and other source attributes. They are the actual curvature contributors in the weak-field construction.

Possible later derived nodes include pair anchors, triple anchors, and other cluster witnesses. These are deferred. They should not be granted equal architectural status until the primary datum system is stable.

The key distinction is this: Atlas is hierarchical in construction but simultaneous in local competition.

Hierarchy is used to organize computation and focus sampling. It is not used to grant inherited field authority at a point. Local dominance and local balance are decided by direct comparison of evaluated readable quantities.

4. TIME-SLICE AND FRAME SETTING

Work on a chosen spacelike slice Sigma_t of a weak-field many-source system. Time enters parametrically. A run produces the datum graph, cloud samples, comparison objects, and extracted artifacts associated with that slice. A temporal history may be built later as an ordered family of such slices.

Frame choice must be explicit. The observer rule associated with a node or run is not invisible bookkeeping. It is part of the physics contract. In early Atlas runs one may adopt a barycentric laboratory frame or another clearly stated adapted frame, but the selected rule must be recorded beside every artifact family and diagnostic archive.

The present revision therefore treats observer choice as a controlled variable, not as an implicit convenience.

5. NODE CLOUDS AS WITNESS STRUCTURES

The earlier paper described parity clouds as collections of parity points sampled along source-launched rays. That remains a useful concept for one extraction layer, but Atlas now requires a richer cloud semantics.

A node cloud is a witness structure attached to a node and evaluated on the slice according to a declared observer rule. Conceptually one may write

Node_k = (position, clock, cloud, observer rule, readable quantities).

The important consequence is that clouds are not only source-centered sprays of candidate interface points. They are organized interrogators of the actual curvature field generated by the physical sources.

Source-node clouds sample the embodied wells associated with declared masses. Node 2 clouds sample the realized center frame of the configuration. Node 0 clouds sample from the ideal geometric witness frame. Node 1 may provide floor or reference comparisons without being promoted into a pseudo-source.

Cloud overlap is not an accidental by-product. It is a first-class object. The overlap region between two clouds is a discovery chamber in which Atlas can compare what different witness structures report about the same neighborhood.

This is especially important for:

- Node 0 vs Node 2,
- Node 2 vs nearby source nodes,
- neighboring source pairs,
- later derived local clusters.

6. CURVATURE EVALUATION PHILOSOPHY

Atlas permits heuristic placement and demands GR-serious valuation.

Heuristic placement means that cloud points may be proposed using cheap, intelligent scouting logic: nodal hierarchy, pair-line intuition, barycentric offsets, low-margin regions, overlap cues, or adaptive refinement rules.

GR-serious valuation means that what is finally evaluated at those points is not an arbitrary proxy substituted for physics. The solver should compute the weak-field curvature-relevant quantity of interest from the actual source construction.

The governing local mechanism is geodesic deviation. Let u^a denote the chosen observer four-velocity and ξ^a a nearby separation vector. Then

$$D^2 \xi^a / d\tau^2 = - R^a{}_{bcd} u^b \xi^c u^d.$$

The associated tidal operator may be written

$$K_{ab} = - R_{acbd} u^c u^d.$$

In vacuum-friendly language the electric Weyl tensor gives

$$E_{ab} = C_{acbd} u^c u^d.$$

In the weak-field regime the same content reduces, up to conventional sign choices, to the Hessian of the source-attributed potential.

Scalar curvature alone is not the object of interest here. The relevant local quantity is the tidal tensor or an explicitly declared readable scalarization of it.

7. SOURCE ATTRIBUTION AND READABLE BUNDLES

In full nonlinear GR, exact source attribution is not available in a naive, global, gauge-independent sense. Atlas does not claim otherwise. But in the weak-field regime it is legitimate as a working solver construction to carry separate source perturbations and to evaluate approximate per-source tidal contributions.

Accordingly, at an important sample site x Atlas may carry:

- the total tidal tensor T_total(x),
- per-source tidal tensor contributions T_i(x),
- a readable bundle derived from those tensors,
- overlap membership and witness provenance.

The theory should therefore compute rich, store medium, and publish slim.

The runtime may preserve the tensor bundle where practical. The extraction layer may derive a restrained family of readable quantities from it. Public or paper-facing summaries may still emphasize the simplest stable scalar channel.

8. SOURCE-READABLE SIGNALS

Let T_i(x,t) denote the source-attributable curvature-relevant tensor bundle associated with source S_i on the slice. Depending on regime and purpose, T_i may refer to the source-attributed electric Weyl tensor, the weak-field tidal tensor, or another explicitly stated weak-field diagnostic derived from the source construction.

To compare sources one applies a readable operator R to obtain a nonnegative signal

Q_i(x,t; n-hat) = R[T_i](x,t; n-hat) >= 0.

The dependence on direction n-hat is retained because readability need not be isotropic. In isotropized or scalarized operational modes one may suppress n-hat and write simply Q_i(x,t).

The structural claim of this paper remains intentionally modest: once a consistent comparison language Q_i is fixed, the parity-extraction machinery follows.

What is new in this revision is the surrounding discipline: the readable signals are now understood as one layer derived from a richer curvature witness stack rather than as the entire native content of Atlas.

9. FIRST EXPLICIT REGIME

To harden the theory, choose a concrete first formal regime.

Let E_ab^(i) denote the source-attributed electric Weyl tensor of source S_i, or in the weak-field Newtonian limit the source-attributed tidal tensor

$$E_{ij}^{(i)} = - \partial_i \partial_j \Phi_i$$

computed from the potential Φ_i attributable to source S_i alone.

Define the directional source-readable signal by projection along a chosen unit direction $\hat{n}$:

$$Q_i(x,t; \hat{n}) = | n^a n^b E_{ab}^{(i)}(x,t) |.$$

This choice is source-resolved, directional, weak-field native, and compatible with the geodesic-deviation logic that motivates Atlas.

For an isolated point mass M_i at source-centered distance r_i,

$$E_{ij}^{(i)} = - G M_i (3 e_i e_j - \delta_{ij}) / r_i^3$$

and along the radial readable direction one obtains

$$Q_i = 2 G M_i / r_i^3.$$

This remains the first explicit parity language of Paper II, but it is now understood as one privileged readable branch inside the broader witness engine.

10. CANONICAL ATLAS MILESTONE 1 SPECIALIZATION

Atlas still requires a narrow first implementation target.

Milestone 1 uses the weak-field point-mass tidal proxy

$$Q_i \sim M_i / r_i^3$$

as the canonical extraction language.

Reasons for preserving this priority are unchanged:

1. It is analytically transparent. 2. It already supports exact toy benchmarks. 3. It stabilizes the extraction pipeline before richer modes are promoted. 4. It reduces solver drift by preventing too many rival readable fields from competing for canonical status too early.

But several older priorities are now explicitly demoted.

The parity node is not milestone-defining. A single privileged center is not allowed to dominate the architecture. The network itself is not treated as the whole ontology of Atlas. Source-launched parity clouds are not the only cloud semantics that matter.

The practical Milestone 1 target is therefore: stable extraction of pairwise balance support, exposed faces, seam candidates, and center diagnostics from the canonical readable field, while retaining the option to carry richer tensor data in the runtime.

11. FROM WITNESS CLOUDS TO OVERLAP COMPARISON

Atlas now inserts a comparison layer between field evaluation and parity extraction.

After cloud points are generated and evaluated, the solver builds overlap objects between materially intersecting cloud supports. These overlaps may be geometric, topological, or proximity-based depending on implementation detail, but they must be explicit. The overlap is where witness structures can be compared on shared neighborhoods.

Not every overlap deserves promotion. To prevent combinatorial overgrowth, Atlas should apply at least three gates:

1. Spatial gate: retain only materially intersecting or near-intersecting cloud supports.
2. Contrast gate: promote only overlaps with interesting agreement, disagreement, anisotropy, low-margin competition, or other declared significance metrics.
3. Persistence gate: treat an inferred artifact as primary only if it survives modest perturbations of cloud density, readable scalarization, or frame choice.

This comparison layer is upstream of parity extraction and helps prevent the solver from confusing sampling quirks with geometric structure.

12. COMPETING FIELDS AND PAIRWISE PARITY FUNCTIONALS

Within a chosen readable regime, define for each source S_i a competing field against which its readable signal is tested.

$$\text{A. Prescribed floor: } C_i(x,t; \hat{n}) = N_0(x,t; \hat{n})$$

$$\text{B. Aggregate of all other sources: } C_i(x,t; \hat{n}) = A(\{Q_j(x,t; \hat{n})\} \text{ for } j \neq i)$$

$$\text{C. Pairwise competitor for local interface extraction: } C_i^{(j)}(x,t; \hat{n}) = Q_j(x,t; \hat{n})$$

$$\text{With } Q_i \text{ and } C_i \text{ fixed, define the source-specific parity functional}$$

$$F_i(x,t; \hat{n}) = Q_i(x,t; \hat{n}) - C_i(x,t; \hat{n}).$$

Equivalent normalized forms may also be used, for example

$$\rho_i = Q_i / C_i \quad \pi_i = (Q_i - C_i) / (Q_i + C_i).$$

For the first many-body extraction regime the cleanest local interface extractor remains the pairwise competitor $C_i^{(j)} = Q_j$, which yields the pairwise parity condition

$$Q_i = Q_j.$$

In the point-mass tidal specialization this becomes

$$M_i / r_i^3 = M_j / r_j^3.$$

13. PAIRWISE BALANCED SHEETS, EXPOSED FACES, AND SEAMS

For two sources S_i and S_j define the pairwise parity difference field

$$\Psi_{ij}(x,t) = Q_i(x,t) - Q_j(x,t),$$

or the normalized residual

$$\Pi_{ij}(x,t) = (Q_i - Q_j) / (Q_i + Q_j).$$

The pairwise balanced sheet Σ_{ij} is the zero set

$$\Sigma_{ij}(t) = \{ x \in \Sigma_t : \Psi_{ij}(x,t) = 0 \}.$$

The sign of Ψ_{ij} indicates local readable dominance. One side is i-dominant, the other is j-dominant, and the sheet itself marks pairwise balanced readability.

In the point-mass tidal specialization,

$$Q_i = 2 G M_i / r_i^3, \quad Q_j = 2 G M_j / r_j^3,$$

so the balanced-sheet condition is

$$M_i / r_i^3 = M_j / r_j^3,$$

or equivalently

$$r_i / r_j = (M_i / M_j)^{(1/3)}.$$

For equal masses the sheet lies on the geometric mid-surface. For unequal masses it bends toward the lower-mass source.

Only some portions of these pairwise sheets are exposed in the many-body partition. The retained face between S_i and S_j is the relevance-filtered portion that actually separates locally dominant regions. Higher-order seams arise where relevant exposed faces intersect transversely.

14. THE SHARED PARITY NETWORK AS A CANONICAL ARTIFACT FAMILY

The many-body shared parity network $N(t)$ is the union of all relevance-filtered pairwise faces:

$$N(t) = \bigcup_{i<j} I_{ij}(t).$$

This remains the central extracted artifact family of Paper II, but the phrase "central" now has a narrower scope. It means central within the canonical parity regime, not central within the full Atlas architecture.

The network partitions the slice into source-dominant cells and shared interface regions. In generic cases it forms a foam-like readable partition. In the single-source isotropic limit it collapses to the spherical shell of Paper I.

The important revision is conceptual rank: the shared parity network is no longer treated as the full identity of Atlas. It is the first stable geometric language produced by a richer witness engine.

15. NODE 0 / NODE 2 DIAGNOSTICS

The most important new diagnostic layer introduced by the project revision is the explicit comparison between the geometric witness and the realized-mass witness.

Let x_0 denote the position of Node 0 and x_2 the position of Node 2. Then the displacement vector

$$\mathrm{Delta}_{02} = x_2 - x_0$$

is itself a first-order asymmetry diagnostic.

But position alone is not enough. Atlas should compare the cloud reports of Node 0 and Node 2 across shared neighborhoods, asking for example:

- where their readable dominance maps agree,
- where they disagree,
- where one identifies low-margin balance and the other does not,
- how those agreements and disagreements change under mass asymmetry, shape perturbation, or cloud refinement.

These center diagnostics do not replace exposed faces and seams. They help explain how the actual mass configuration departs from its ideal geometric baseline and why certain extracted artifacts appear where they do.

16. THREE-BODY BENCHMARK LADDER

The equal-mass equilateral three-body case remains the first exact benchmark.

Consider three equal weak-field point masses M placed at the vertices of an equilateral triangle in the plane $z = 0$ of a time slice Sigma_t. Choose the point-mass tidal readable field

$$Q_i = 2 G M / r_i^3.$$

Because the masses are equal, each pairwise balanced sheet satisfies

$$r_i = r_j.$$

Thus each pairwise sheet is the perpendicular-bisector plane of the segment joining the corresponding pair of source points. The three planes intersect along the line through the circumcenter orthogonal to the triangle plane.

This benchmark remains indispensable, but it is no longer enough by itself. The revised Atlas benchmark ladder is:

1. Equal-mass equilateral three-body case: verify datum construction, midpoint-plane recovery, and central seam.

2. Unequal-mass triangle: verify Node 0 / Node 2 separation, bent balance sheets, and center-cloud diagnostic changes.

3. Near-collinear three-body case: stress-test seam extraction and low-margin overlap handling.

4. Controlled perturbation family: vary one mass or one position continuously and require extracted artifacts to deform continuously rather than teleport under small changes.

5. Cloud-layout robustness tests: modestly perturb witness cloud placement and require primary artifacts to persist if they reflect genuine field structure.

17. LOCAL PROPOSITION LAYER

Proposition 1 (Local existence of pairwise balanced sheets). Let Q_i and Q_j be continuous source-readable signals on an open region U of Sigma_t, and define Psi_ij = Q_i - Q_j. Suppose there exists x_* in U such that Psi_ij(x_*) = 0 and grad Psi_ij(x_*) != 0. Then the local balanced set Sigma_ij = { x in U : Psi_ij(x) = 0 } is, near x_*, a smooth embedded surface.

Proposition 2 (Equal-mass midpoint-plane limit). For two equal weak-field point masses with point-mass tidal readable field, Q_i = Q_j implies r_i = r_j. Therefore the balanced sheet is the perpendicular-bisector plane.

Proposition 3 (Unequal-mass bending law). For two weak-field point masses in the same regime, the balanced sheet obeys

r_i / r_j = (M_i / M_j)^(1/3),

hence the sheet is displaced toward the lower-mass source.

Proposition 4 (Generic seam emergence). If three relevant pairwise balanced sheets intersect transversely inside a region where Q_i = Q_j = Q_k >= Q_l for all other l, then the triple-balance set is locally one-dimensional.

These claims remain intentionally local and modest. They support the canonical parity layer without pretending to complete the full witness-engine theory.

18. COMPUTATIONAL ARTIFACT DISCIPLINE

The theoretical object is not the same thing as the computational object, and the computational object is not the same thing as the rendered image.

For Atlas, the explicit hierarchy is now:

1. datum graph and cloud specifications,
2. evaluated sample sites and readable bundles,
3. overlap comparison objects,
4. reconstructed interface sheets,
5. relevance-filtered faces and seams,
6. certified artifacts,

7. smoothed meshes or rendered visualizations.

A rendered wireframe is not the theory. A marching-cubes surface is not the definition of the network. But these are legitimate engineering objects and are required for testing, comparison, and archive.

This hierarchy also clarifies what should not dominate development: a visually elegant rendered object is not evidence unless it is anchored to a stable extraction layer and survives basic perturbation checks.

19. JSON AND CONFIG STANCE

Source definitions should live in a machine-readable configuration, such as a JSON schema, that is richer than the first pass strictly requires.

Good declared fields include: stable source identifiers, mass, position, velocity, softening, optional spin, optional multipoles, shape model, active flag, root-datum modes, numerical controls, readable-field mode, and extraction settings.

The important split is:

Declared in config:

- source facts,
- solver intent,
- reference rules.

Derived at runtime:

- Node 0,
- Node 2,
- any later pair or triple anchors,
- overlap structure,
- extracted artifacts,
- certified diagnostics.

This keeps ontology and runtime derivation cleanly separated.

20. CLAIM STATUS

Established within the first explicit weak-field tidal regime:

- source-readable signals Q_i are well defined once a readable operator is chosen;
- pairwise balanced sheets arise from $Q_i = Q_j$;
- in the point-mass specialization the sheet equation is $M_i / r_i^3 = M_j / r_j^3$;
- equal-mass midpoint-plane and unequal-mass bending results follow;

- the equilateral three-body benchmark yields a three-face, one-seam seed;
- the datum-centered revision is compatible with, and does not invalidate, these parity-layer results.

Operationally fixed for Atlas Milestone 1:

- the canonical readable-field regime remains $Q_i \sim M_i / r_i^3$;
- Node 0, Node 1, Node 2, and source nodes must be constructed distinctly;
- node-specific witness clouds and overlap comparison now sit upstream of parity extraction;
- exposed faces, seams, and Node 0 / Node 2 diagnostics define success more fully than any parity-node claim;
- richer readable-field modes remain validation branches.

Deferred or still open:

- robust frame dependence studies,
- optimal cloud generation and refinement rules,
- artifact certification standards across readable-field modes,
- large-N phenomenology,
- strong-field generalization,
- temporal logic beyond the static-slice emphasis of V1.

21. CONCLUSION

Paper I gave the shell. The original Paper II gave the network. This revision places the network in its proper architectural setting.

Atlas is now best understood as a datum-centered curvature witness engine that builds root datums, deploys node-specific witness clouds, evaluates weak-field tidal structure, compares overlapping reports, and then extracts canonical artifact families from that relational field of comparison.

Within that broader engine, the shared parity network remains indispensable. It is the first stable many-body geometric language Atlas can extract cleanly. But it is no longer the full conceptual identity of the project, and it should not be allowed to overconstrain future development.

The revised discipline is therefore simple: preserve what is physically serious and computationally fertile, demote what aged into premature centrality, and let the solver grow from root datums, witness clouds, overlap comparison, and tidal evaluation outward.

That is the updated role of Paper II within Atlas.