Signed-Incidence Structure in Compositional Verification

A Structural Note

John Komkov
Independent Researcher
`https://github.com/jkomkov/bulla`

April 2026

Abstract

This is a short structural note, not a standalone program statement. Its job is to supply one disciplined source for three facts used repeatedly in the Bulla corpus: (i) each row of the coboundary matrix $\delta(G)$ has at most one $+1$ and at most one -1; (ii) the coherence fee $\mathrm{fee}(G)$ is therefore field-independent as a difference of two totally-unimodular ranks; and (iii) pairwise endpoint coupling is bounded by rank 2 on every $(\text{edge}, \text{dimension})$ block. A negative corollary records that any partition-style typed repair constraint is structurally wrong: enforcing it on the 703-composition real-schema corpus makes fee-zero repair impossible in 151 of 240 nonzero-fee compositions. The signed-incidence $\to$ total-unimodularity step is Lean-verified.[1] The greedy untyped repair algorithm remains correct and optimal for its stated problem.

Role. Five papers in the Bulla corpus—[2], [3], [4], [5], and [6]—refer to a "companion signed-incidence note" supplying their middle-layer structural background. This is that note. It does not introduce new claims; it freezes the exact statements those papers cite.

1 Structural proposition

Proposition 1.1 (Signed-incidence row structure)**.** *For every composition G constructible by* `build_coboundary`*, each row of the coboundary matrix $\delta(G)$ has at most one entry equal to $+1$ and at most one entry equal to -1, with all other entries zero.*

Proof. A row is indexed by one $(\text{edge}, \text{dimension})$ pair. The construction can write at most one source term at `dim.from_field` and at most one target term at `dim.to_field`. No other writes touch that row. □

Status. Proved by construction, runtime-asserted, and empirically verified on the full 1596-composition corpus.

2 Field-independence corollary

Corollary 2.1 (Field-independence of the fee)**.** *For every composition G in the Bulla coboundary family,*

$$\mathrm{fee}_{\mathbb{Q}}(G) \;=\; \mathrm{fee}_{F}(G)$$

[1] Aristotle run `fecfaac3-fed7-40dd-82aa-5eef8edd7111`; theorem `signed_incidence_det_in_unit`, 100 lines, 8 helper lemmas, no `sorry`, standard axioms only. Submission stub: `papers/sheaf/lean/SHEAF/SignedIncidence.submission.lean`; proof outline: `SignedIncidence.ARISTOTLE_SUMMARY.md`.

for every field F.

Proof. (1) By proposition 1.1, both $\delta_{\text{full}}(G)$ and $\delta_{\text{obs}}(G)$ are signed incidence matrices.

(2) Signed incidence matrices are totally unimodular [1].

(3) Totally unimodular matrices have field-independent rank.

(4) $\text{fee}(G) = \text{rank}(\delta_{\text{full}}(G)) - \text{rank}(\delta_{\text{obs}}(G))$, so the fee is field-independent as a difference of two field-independent ranks.

□

Status. Proved. The TU step is Lean-verified in the current Cycle 18 chain.

Remark 2.2 (Boundary). This proves numerical agreement across coefficient fields. It does *not* prove that the $\mathbb{F}_2$ holonomy interpretation and the $\mathbb{Q}$-valued Bulla implementation are the same mathematical object. The disciplined claim is weaker and sufficient: they compute the same numerical invariant from structurally corresponding presentations.

3 Endpoint-coupling corollary

Corollary 3.1 (Pairwise endpoint-coupling bound)**.** *Let G be a pairwise composition, and let B be the hidden-column block associated to one fixed* (*edge, dimension*) *pair in the witness Gram $K(G) = H^\top(I - P_O)\,H$. Then*

$$\text{rank}\big(K[B,B]\big) \;\leq\; 2.$$

Proof. In Bulla's pairwise setting, a semantic dimension on a fixed edge can contribute at most one hidden source column and at most one hidden target column, because `SemanticDimension` carries only one `from_field` and one `to_field`. Therefore $|B| \leq 2$. Since $K[B,B]$ is a $|B| \times |B|$ matrix, its rank is at most $|B|$, hence at most 2. □

Status. Proved.

Remark 3.2 (Why the projection issue does not block the theorem). An earlier proof sketch reasoned through independence in δ and then worried about the projection $(I - P_O)$. That was unnecessary. The theorem only needs the cardinality bound on the block definition itself; projection can reduce rank, but it cannot increase rank above the number of columns in the block.

Empirical sharpening on the 703 real-schema corpus.

- 3873 multi-column (edge, dimension) blocks were tested.
- 3703 had rank 2.
- 170 had rank 1.
- 0 blocks exhibited the feared pattern "rank 2 in the corresponding δ block but rank < 2 after projection into K."

The theorem is therefore not only true but active in practice: rank-2 endpoint coupling is common, not merely possible.

4 Negative corollary for typed repair

Corollary 4.1 (Partition-style separability is structurally wrong)**.** *Consider the proposed typed-repair constraint "disclose at most one field from each* ($edge, dimension$) *block." On the* 703*-composition real-schema corpus, restricting repair to that constraint makes fee-zero repair impossible in* 151 *of the* 240 *nonzero-fee compositions (*62.9%*).*

Interpretation. The failure is not that the partition-constrained solution is sometimes more expensive. The stronger failure is that it often cannot reach zero fee at all. When both endpoints of an edge contribute hidden fields in the same dimension, those two disclosures can carry independent witness information. Forbidding their joint disclosure destroys valid repair paths. **Status.** Empirical finding, locked by the corpus experiment.

5 Algorithmic consequence

Consequence 5.1 (The untyped greedy is already optimal for its stated problem)**.** `weighted_greedy_repair` *computes a minimum-cost basis of the witness matroid M/O. That remains the correct algorithm for the untyped repair problem it actually claims to solve.*

The falsified step was not the greedy algorithm. The falsified step was the proposed separable typed extension that tried to model same-dimension constraints as a partition matroid. The structural lesson is:

- untyped greedy remains correct,
- partition-style typed repair is not the right abstraction,
- the open problem is the correct coupled-disclosure objective, not a replacement for the current untyped basis computation.

6 Evidence summary

- Signed-incidence row structure: code-audited, runtime-asserted, 1596/1596 verified.
- Field-independence: structural proof plus Lean-backed TU chain.
- Pairwise endpoint coupling: theorem by block cardinality, with 3873/3873 multi-column witness blocks satisfying rank ≤ 2 on the 703 real-schema corpus.
- Projection sanity check for corollary 3.1: zero cases where a corresponding δ block had rank 2 and the K block dropped below 2.
- Partition constraint failure: 151/240 nonzero-fee compositions cannot be fully repaired under the one-per-block restriction.

7 Claim boundaries

This note intentionally does *not* claim any of the following:

- a characterization of the correct dimension-aware repair objective,
- generalized k-way endpoint-coupling bounds for non-pairwise compositions,
- that the $\mathbb{F}_2$ and $\mathbb{Q}$ stories are literally one object rather than numerically aligned invariants,

- any commercial, philosophical, or programmatic interpretation beyond the structural statements above.

Those belong elsewhere. The purpose of this note is narrower: to let the rest of the corpus cite one exact middle layer instead of re-deriving it ad hoc.

References

[1] A. Schrijver, *Theory of Linear and Integer Programming*, Wiley, 1986. Theorem 19.3: a $\{0, \pm 1\}$ matrix is totally unimodular if and only if for every subset R of its rows there is a bipartition $R = R_1 \cup R_2$ such that for every column j, $|\sum_{i \in R_1} a_{ij} - \sum_{i \in R_2} a_{ij}| \leq 1$. For signed incidence matrices, the trivial partition $R_1 = R$, $R_2 = \emptyset$ satisfies this immediately.

[2] J. Komkov, "The Column-Matroid Backbone of a Composition-Coherence Fee: A Structural Identity and Its Corollaries," `papers/hierarchical-fee/paper/column-matroid-backbone.tex`, 2026.

[3] J. Komkov, "The Witness Gram: Kron-Reduced Laplacian of Hidden Convention Risk," `papers/hierarchical-fee/paper/witness-gram.tex`, 2026.

[4] J. Komkov, "A Hierarchical Coherence Fee for Compositional Verification," `papers/hierarchical-fee/paper/hierarchical-fee.tex`, 2026.

[5] J. Komkov, "Bridge: Empirical Witness for Compositional Failure," `papers/bridge/paper/bridge.tex`, 2026.

[6] J. Komkov, "The Coherence Cliff," `papers/coherence-cliff/paper/the-coherence-cliff.tex`, 2026.