

# The Witness Gram as a Kron-Reduced Laplacian of Hidden Convention Risk

John Komkov

Independent Researcher

<https://github.com/jkomkov/bulla>

April 2026

## Abstract

The witness Gram  $K(G) = H^\top(I - P_O)H$  of a Bulla tool composition is always a Kron-reduced graph Laplacian — literally the matrix Gabriel Kron introduced in 1939 for reducing electrical networks to sub-networks of interest. Under Dimension-Field Disjointness (DFD) it decomposes as a direct sum of Kron reductions of per-dimension carrier-graph Laplacians; without DFD it remains a single Kron-reduced Laplacian, of the merged coboundary-incidence graph. This identifies the canonical object sitting at the corank of the backbone theorem [2]: the coherence fee equals  $\text{rank}(K(G))$ , and the bases of  $K(G)$ 's column matroid reproduce the minimum disclosure sets of [1]. We show that the effective resistance between two hidden fields in the Kron-reduced witness geometry quantifies disclosure substitutability: the smaller the resistance, the more equivalent the two fields are as repair choices. All structural identities are verified on the 703-composition real-schema MCP registry corpus in exact rational arithmetic (703/703 rank identities, 240/240 leverage predictions including 37 natural partial-CHP regime breaks). The paper closes the program's structural theory and identifies the frontier where field-level diagnostics become discriminative. A companion signed-incidence note [8] supplies the pairwise endpoint-coupling boundary used here: same-dimension witness blocks are rank-bounded by construction, and their frequent rank-2 realization is evidence against separable typed-repair constraints, not a characterization of the true dimension-aware objective.

## 1 Introduction

**Master sentence.** *The witness Gram of a Bulla composition is always a Kron-reduced graph Laplacian. Dimension-Field Disjointness controls whether that Laplacian decomposes as a direct sum over dimensions; the Convention-Hidden Partition controls whether the Kron reduction is trivial.*

The coherence fee [1] of a tool composition  $G$  is an integer measuring the codimension of an observable coboundary inside a full internal-state coboundary. A companion result [2] shows that the fee equals the corank of the observable column set in a linear matroid on the columns of the full coboundary, and that the closed-form pairwise fee  $\text{fee}(A, B) = U(A) + U(B) + |\text{shared\_dims}(A, B)|$  holds under two structural hypotheses (DFD and CHP) on every one of 703 real-schema MCP compositions tested.

This paper asks a more pointed question. The backbone identifies the fee as *the rank of an object*; what is the object, and what can it tell us that the integer cannot? We give three answers, developed through one central theorem:

1. The object is the *witness Gram*  $K(G) = H^\top(I - P_O)H$  — a symmetric positive-semidefinite matrix on the hidden-field coordinate space, of exact rational entries, with rank equal to the fee.
2. Under DFD, the witness Gram decomposes as a direct sum of *Kron reductions* of per-dimension carrier-graph Laplacians — literally the matrices Gabriel Kron introduced in 1939 for reducing electrical networks to sub-networks of interest [3].
3. In the Kron-reduced geometry, the classical effective resistance between two hidden fields quantifies their *disclosure substitutability*: small resistance means disclosing either field largely resolves the other’s convention risk; large resistance (within a component) means the fields are structurally distinct witness obligations.

**Claim discipline.** What is established here is the object-level witness geometry and the Laplacian structure it carries. What failed productively was the older hope that same-dimension typed constraints might collapse to a separable partition rule. What remains open is the correct coupled-disclosure optimization problem; the signed-incidence note [8] supplies the honest middle-layer boundary for that statement, and this paper stays on the geometry side of the line. **A concrete example, before the machinery.** Consider two compositions in the 703-corpus with identical fee = 2: the `financial_pipeline` (3 tools, 4 hidden fields spanning 2 convention dimensions, each a 2-vertex carrier graph) and a subset of the `airtable-mcp + mcp-xmind` pairing (3 hidden `path`-related fields in one 5-vertex carrier graph with 2 observable carriers Kron-reduced out). The bare integer fee cannot tell these apart. The witness Gram does:

$$K(\text{financial\_pipeline}) = \begin{bmatrix} 1 & -1 & & & & \\ -1 & 1 & & & & \\ & & 1 & -1 & & \\ & & -1 & 1 & & \\ & & & & & & \end{bmatrix} = L(K_2) \oplus L(K_2),$$

$$K(\text{airtable+mcp-xmind}_{\text{path dim}}) = \frac{5}{3} \cdot L(K_3) = \text{Schur}(L(K_5) / O_d) \text{ for } |O_d| = 2.$$

Two different graph Laplacians, two different bases, two different effective-resistance geometries — and therefore two different disclosure prescriptions — at the same fee. In the first, any one field per dimension suffices (1 spanning tree per  $K_2$ -block, giving a single minimum basis `{day_convention, risk_metric}` up to per-block choice). In the second, any two of the three hidden `path`-fields form a valid disclosure set (3 bases on  $K_3$  by Matrix-Tree), and the three fields are equally good disclosure substitutes at effective resistance  $\frac{2}{3}$  apart.

**Organization.** §2 proves the rank identity (Tier 1, no hypotheses). §3 proves the Kron-reduction theorem under DFD (Tier 2, the centerpiece). §4 states and proves the four corollaries that fall out of Tier 2 (leverage, no-co-loop, basis-count via Matrix-Tree, effective resistance as disclosure substitutability). §5 studies regime breaks: 37 natural partial-CHP compositions in the corpus, plus synthetic DFD and bridge-topology constructions. §6 reports the corpus-wide verification in exact rational arithmetic. §7 closes with operational interpretation.

## 2 The witness Gram and its rank identity

Fix a Bulla composition  $G = (T, E, D)$  with tool set  $T$ , directed edges  $E$ , and semantic dimensions  $D$ . For each tool  $t \in T$  let  $t.\text{int}$  and  $t.\text{obs} \subseteq t.\text{int}$  denote its internal and observable field sets. Write  $\delta_{\text{full}} : C_{\text{full}}^0 \rightarrow C^1$  and  $\delta_{\text{obs}} : C_{\text{obs}}^0 \rightarrow C^1$  for the coboundaries as constructed in [2, §2], which share the edge-dimension cochain space  $C^1$ .

By the Column-Compression Identity [2, Theorem 2.1],  $\delta_{\text{obs}} = \delta_{\text{full}}[:, \mathcal{O}]$  for a column index subset  $\mathcal{O} \subseteq \{1, \dots, n_{\text{full}}\}$ . Write  $H \subseteq \{1, \dots, n_{\text{full}}\}$  for the complementary *hidden* column set, and  $H_c := \delta_{\text{full}}[:, H]$  for the hidden-column block. Let  $P_O : C^1 \rightarrow C^1$  be the orthogonal projector onto  $\text{range}(\delta_{\text{obs}})$ .

**Definition 2.1** (Witness Gram). The *witness Gram* of  $G$  is the symmetric positive-semidefinite  $|H| \times |H|$  matrix

$$K(G) := H_c^\top (I - P_O) H_c.$$

Entries of  $K(G)$  are rational by construction:  $\delta_{\text{full}}$  is a signed incidence matrix (one  $+1$  and one  $-1$  per row, zeros elsewhere) and is therefore totally unimodular, so the Moore-Penrose pseudoinverse of any column submatrix exists over  $\mathbb{Q}$ .

**Theorem 2.2** (Rank identity). *For every Bulla composition  $G$ ,*

$$\text{rank}(K(G)) = \text{fee}(G).$$

*Moreover, the bases of the column matroid of  $K(G)$  coincide with the bases of the witness matroid  $M/\mathcal{O}$  in [2, §2.3], and therefore with the minimum disclosure sets of [1, Theorem 6.2].*

*Proof.* Write  $W := (I - P_O)H_c$ . Then  $K(G) = W^\top W$ , so  $\text{rank}(K) = \text{rank}(W)$ . Since  $(I - P_O)$  is the orthogonal projector onto  $\text{range}(\delta_{\text{obs}})^\perp$  and  $H_c$  spans (together with  $\delta_{\text{obs}}$ ) the whole range of  $\delta_{\text{full}}$ , we have  $\text{rank}(W) = \text{rank}(\delta_{\text{full}}) - \text{rank}(\delta_{\text{obs}}) = \text{fee}(G)$  by the backbone identity. The basis equivalence follows because column-bases of  $W$  are exactly column-subsets of  $H_c$  whose images under  $I - P_O$  are linearly independent — which is the matroid-contraction definition of bases in  $M/\mathcal{O}$ .  $\square$

*Remark 2.3* (Tier 1 is regime-independent). Theorem 2.2 makes no hypothesis on DFD or CHP. It is a structural identity between linear-algebraic objects. The regime hypotheses of §3 control not the *rank* of  $K(G)$  but its *shape* — specifically, when it decomposes as a direct sum.

### 3 Kron reduction under DFD

We recall DFD [2, Def. 2.7]: a composition  $G$  satisfies Dimension-Field Disjointness if no  $(t, f)$  pair is referenced by two distinct dimensions in  $D$ . Under DFD, the column partition  $\{1, \dots, n_{\text{full}}\} = \bigsqcup_{d \in D} V_d$  induces a block-diagonal structure on the full coboundary, with one block per dimension.

**Definition 3.1** (Carrier graph). For dimension  $d \in D$ , the *carrier graph*  $G_d$  is the multigraph with

- vertex set  $V_d = \{(t, f) : f = d.\text{from} \text{ or } f = d.\text{to} \text{ on tool } t\}$ ,
- edge set indexed by rows of  $\delta_{\text{full}}$  in dimension  $d$ , each row connecting the columns with nonzero entries  $(-1, +1)$ .

Let  $L_d$  be the combinatorial Laplacian of  $G_d$ . Under DFD each column of  $\delta_{\text{full}}$  belongs to exactly one  $V_d$ , so the hidden/observable split refines as  $V_d = H_d \sqcup O_d$  with  $H_d = V_d \cap H$  and  $O_d = V_d \cap \mathcal{O}$ .

**Definition 3.2** (Kron reduction). The *Kron reduction* of  $L_d$  with respect to  $O_d$  is the Schur complement

$$\text{Schur}(L_d / O_d) := L_d[H_d, H_d] - L_d[H_d, O_d] L_d[O_d, O_d]^{-1} L_d[O_d, H_d],$$

whenever  $L_d[O_d, O_d]$  is invertible. When  $O_d = \emptyset$ , the Kron reduction is defined to be  $L_d$  itself (no absorption). In degenerate cases — e.g.  $O_d$  vertices form a disconnected subgraph causing  $L_d[O_d, O_d]$  to be singular — we replace the inverse by the Moore-Penrose pseudoinverse; for connected carrier graphs with  $|O_d| \geq 1$  and any non-grounded vertex outside  $O_d$ , the principal submatrix  $L_d[O_d, O_d]$  is nonsingular [4, Prop. 3].

The name ‘‘Kron reduction’’ is not decorative: this is the matrix introduced by Gabriel Kron in 1939 for eliminating interior nodes from a power-system network while preserving the effective behavior at the boundary [3]. In our setting the ‘‘boundary’’ is the hidden carrier columns  $H_d$  and the ‘‘interior’’ is the observable carriers  $O_d$ , which are grounded by the projection  $I - P_O$ . The same matrix appears in spectral graph theory [4], random-walk analysis, and the theory of effective resistance [5].

**Theorem 3.3** (Kron reduction of the witness Gram). *If  $G$  satisfies DFD, then*

$$K(G) = \bigoplus_{d \in D} K_d, \quad K_d = \text{Schur}(L_d / O_d).$$

*When additionally  $O_d = \emptyset$  for every  $d$  (full CHP), the Kron reduction is trivial and  $K_d = L_d$  itself (the pure carrier-graph Laplacian).*

*Proof. Block decomposition.* Under DFD, the column partition  $\{V_d\}_d$  induces a block-diagonal structure on  $\delta_{\text{full}}$ : each column has nonzero entries only in rows of the single dimension that references its  $(t, f)$  pair (because row indexing splits by  $(e, d.\text{name})$ , and only the unique dimension containing  $(t, f)$  contributes). Hence  $\delta_{\text{full}} \cong \bigoplus_d \delta_{\text{full},d}$ , and the hidden/observable split respects the dimension blocks. The witness Gram inherits the block decomposition:  $K(G) = \bigoplus_d K_d$ , with cross-block entries zero.

*Within-block identification.* Each dimension- $d$  row of  $\delta_{\text{full}}$  has exactly two nonzero entries ( $-1$  at source,  $+1$  at destination), making  $\delta_{\text{full},d}$  the signed incidence matrix of the carrier multigraph  $G_d$ . A classical identity (Godsil-Royle [6, Ch. 8]) gives

$$\delta_{\text{full},d}^\top \delta_{\text{full},d} = L_d,$$

the combinatorial Laplacian of the underlying unsigned multigraph.

*Within-block Kron identification.* Write  $\delta_{\text{full},d}[H_d]$  and  $\delta_{\text{full},d}[O_d]$  for the hidden and observable column blocks of  $\delta_{\text{full},d}$ . Because  $\delta_{\text{full}}$  is block-diagonal under DFD, the range of  $\delta_{\text{obs}}$  restricted to dimension- $d$  rows is exactly  $\text{range}(\delta_{\text{full},d}[O_d])$  — observable columns from other dimensions produce zero entries on dimension- $d$  rows and therefore cannot contribute to that row subspace. Consequently  $P_O$  restricted to dimension- $d$  rows acts as the within-dimension observable projector onto  $\text{range}(\delta_{\text{full},d}[O_d])$ , which by the Laplacian identity above has Gramian  $L_d[O_d, O_d]$ . The standard formula  $P = A(A^\top A)^{-1}A^\top$  for the projector onto a column span gives

$$K_d = \delta_{\text{full},d}[H_d]^\top (I - P_O) \delta_{\text{full},d}[H_d] = L_d[H_d, H_d] - L_d[H_d, O_d] L_d[O_d, O_d]^{-1} L_d[O_d, H_d],$$

which is  $\text{Schur}(L_d / O_d)$ . When  $O_d = \emptyset$  the Schur term vanishes and  $K_d = L_d[H_d, H_d] = L_d$ .  $\square$

**Definition 3.4** (Coboundary-incidence graph). The *coboundary-incidence graph*  $G^*$  of a composition  $G$  has vertex set  $\{(t, f) : f \in t.\text{int}\}$  (all internal column indices) and edge set the rows of  $\delta_{\text{full}}$ , each row connecting the two columns with nonzero entries  $(-1, +1)$ . Under DFD,  $G^*$  splits as a disjoint union of carrier graphs,  $G^* = \bigsqcup_{d \in D} G_d$ .

**Theorem 3.5** (Kron reduction without DFD). *For every Bulla composition  $G$  (with or without DFD),*

$$K(G) = \text{Schur}(L(G^*) / \mathcal{O}),$$

*where  $L(G^*)$  is the combinatorial Laplacian of the coboundary-incidence graph. Under DFD this reduces to the direct-sum form of Theorem 3.3; without DFD,  $G^*$  is a single multigraph spanning multiple dimensions and  $K(G)$  is the Kron reduction of  $L(G^*)$  as a whole.*

*Proof.* The argument of Theorem 3.3 applied to the full coboundary — without the DFD-driven block decomposition — gives  $K(G) = H_c^\top(I - P_O)H_c$ , with  $H_c, P_O$  as before. Identifying  $\delta_{\text{full}}^\top \delta_{\text{full}} = L(G^*)$  (each row a signed incidence edge in  $G^*$ , Godsil-Royle [6]) and applying the same projector formula  $P_O = \delta_{\text{full}}[\mathcal{O}]L(G^*)[\mathcal{O}, \mathcal{O}]^{-1}\delta_{\text{full}}[\mathcal{O}]^\top$  (with pseudoinverse in degenerate cases) yields the Schur-complement expression. Under DFD,  $L(G^*)$  and thus the Schur complement are block-diagonal, recovering Theorem 3.3.  $\square$

*Remark 3.6.* Theorem 3.5 is the regime-independent form of the main theorem: the witness Gram is *always* a Kron-reduced graph Laplacian, while DFD controls whether the graph splits by dimension. Both regimes are demonstrated empirically and synthetically in §5.

**Example 3.7** (Worked example on `financial_pipeline`). The `financial_pipeline` composition has 3 tools and 4 hidden fields split between two dimensions (`day_convention` and `risk_metric`), each with carrier graph  $K_2$  on 2 hidden vertices (no observable carriers; full CHP). The witness Gram is

$$K = L(K_2) \oplus L(K_2) = \begin{bmatrix} 1 & -1 & & & & \\ -1 & 1 & & & & \\ & & 1 & -1 & & \\ & & -1 & 1 & & \\ & & & & & & \\ & & & & & & \end{bmatrix},$$

with  $\text{rank}(K) = 2 = \text{fee}$ , leverage  $\frac{1}{2}$  on every hidden field, and basis count  $\tau(K_2) \cdot \tau(K_2) = 1 \cdot 1 = 1$  per dimension-pair: the minimum disclosure set is unique up to the dimension-level choice.

## 4 Four corollaries

All four corollaries fall out of Theorem 3.3 plus standard graph Laplacian theory. Write  $C$  for a connected component of the effective graph  $\tilde{G}_d$  whose Laplacian is  $K_d$ . (Under full CHP,  $\tilde{G}_d = G_d$ .)

**Corollary 4.1** (Leverage per component). *Under DFD, the leverage score of a hidden field  $j$  in component  $C$  of  $\tilde{G}_d$  is*

$$\ell_j = \frac{|H_C| - 1}{|H_C|},$$

*provided  $|H_C| \geq 2$ . Leverage is uniform within a component; leverage values may differ across components and across dimensions.*

*Proof.* Kron reduction preserves connectivity of the effective graph: when  $L_d$  is the Laplacian of a connected carrier graph  $G_d$ , the Kron-reduced Laplacian  $K_d = \text{Schur}(L_d/O_d)$  is the Laplacian of a connected effective graph  $\tilde{G}_d$  on  $H_d$  [4, Theorem 3.1] — Kron reduction adds “virtual edges” between hidden vertices connected through observable paths, so the effective graph is at least as connected as the induced subgraph  $G_d[H_d]$  (and strictly more connected when observable vertices were cut-vertices of  $G_d$ ). Consequently  $K_C$  for each component  $C$  of  $\tilde{G}_d$  inherits null space  $\text{span}(\mathbf{1}_{H_C})$ , so  $\text{rank}(K_C) = |H_C| - 1$ . The orthogonal projector onto  $\text{col}(K_C) = \mathbf{1}_{H_C}^\perp$  is  $P = I - \frac{1}{|H_C|}\mathbf{1}_{H_C}\mathbf{1}_{H_C}^\top$ , whose diagonal is  $(|H_C| - 1)/|H_C|$  at every coordinate in  $H_C$ . The leverage of field  $j$  is by definition this diagonal entry.  $\square$

**Corollary 4.2** (No coloops under connected components). *Under DFD, if every hidden component  $C$  of  $\tilde{G}_d$  has  $|H_C| \geq 2$ , then the witness matroid  $M/\mathcal{O}$  has no coloops.*

*Proof.* A coloop is a field  $j$  appearing in every basis of  $M/\mathcal{O}$ , which corresponds to leverage  $\ell_j = 1$ . By Corollary 4.1,  $\ell_j = (|H_C| - 1)/|H_C| < 1$  whenever  $|H_C| \geq 2$ . Singleton components ( $|H_C| = 1$ ) correspond to hidden fields isolated in  $\tilde{G}_d$  and have leverage zero (loops), not one.  $\square$

**Corollary 4.3** (Basis count via Matrix-Tree). *Assume DFD and that each effective graph  $\tilde{G}_d$  is connected. Then the witness matroid has*

$$|\mathcal{B}(M/\mathcal{O})| = \prod_{d \in D} \tau(\tilde{G}_d),$$

where  $\tau(\tilde{G}_d)$  is the number of spanning trees of  $\tilde{G}_d$ . Under full CHP,  $\tilde{G}_d = G_d$  and the formula reduces to  $\prod_d \tau(G_d)$ .

*Proof.* Under DFD the witness matroid decomposes as a direct sum  $M/\mathcal{O} = \bigoplus_d (M/\mathcal{O})_d$  with  $(M/\mathcal{O})_d$  the graphic matroid of  $\tilde{G}_d$ . By the Matrix-Tree Theorem applied to each connected  $\tilde{G}_d$ , the bases of  $(M/\mathcal{O})_d$  are the spanning trees of  $\tilde{G}_d$ , counted by  $\tau(\tilde{G}_d) = \text{any cofactor of } K_d$ . Product across dimensions gives the stated formula.  $\square$

*Remark 4.4.* If some  $\tilde{G}_d$  is disconnected (not observed in the 703-corpus but possible in principle), the Matrix-Tree formula gives the number of *spanning forests* with one tree per component; the basis count of  $(M/\mathcal{O})_d$  is then the forest count. See [6].

Concrete spanning-tree counts for graph families appearing in the 703-corpus: a single bilateral edge gives  $\tau = 1$  (a unique minimum disclosure set); the complete graph  $K_n$  gives  $\tau = n^{n-2}$  (Cayley’s formula); the complete bipartite  $K_{m,n}$  gives  $\tau = m^{n-1} n^{m-1}$ . For most 703-corpus dimensions the carrier graph has a single bilateral edge and the disclosure set is unique.

**Corollary 4.5** (Effective resistance as disclosure substitutability). *By Theorem 3.3,  $K_d$  is a graph Laplacian (specifically, the Laplacian of the effective graph  $\tilde{G}_d$ ), so the classical effective resistance of [5] is well-defined on hidden fields. For hidden fields  $i, j$  in the same component  $C$  of  $\tilde{G}_d$ ,*

$$R_{\text{eff}}(i, j) = K_d^+[i, i] + K_d^+[j, j] - 2K_d^+[i, j].$$

Effective resistance in the Kron-reduced witness geometry quantifies disclosure substitutability: the smaller the resistance between two hidden fields, the more equivalent they are as repair choices.

Concretely: small  $R_{\text{eff}}(i, j)$  means disclosing either field substantially resolves the other’s convention risk; large  $R_{\text{eff}}(i, j)$  within a component means the fields are structurally distinct witness obligations;  $R_{\text{eff}}(i, j) = \infty$  (across components of  $\tilde{G}_d$  or across dimensions under DFD) means the fields are structurally independent and no disclosure substitution exists. The classical graph interpretation [5] identifies  $R_{\text{eff}}$  with the voltage drop between  $i$  and  $j$  when unit current flows through the effective network; the Bulla interpretation reads that voltage drop as governance cost of disclosure substitution.

## 5 Regime breaks

The two regime hypotheses of Theorem 3.3 — DFD and CHP — can fail independently, and within any regime the effective-graph topology (bridges, cut vertices) carries additional structure. Table 1 organizes the picture. The 703-corpus already contains the most important partial-CHP failure naturally; we complete the picture with two synthetic constructions, one per uncovered regime.

Leverage is  $(|H_C| - 1)/|H_C|$  per component  $C$  of the effective graph in all four regimes (Corollary 4.1); what changes is which graph the components live in.

Table 1: Regime taxonomy: witness Gram form under DFD  $\times$  CHP.

DFD	CHP	Witness Gram form	Example
✓	full	$K = \bigoplus_d L_d$	203 corpus; <code>financial_pipeline</code>
✓	partial	$K = \bigoplus_d \text{Schur}(L_d/O_d)$	37 corpus; <code>mcp-xmind</code>
×	full	$K = L(G^*)$	synthetic DFD violation
×	partial	$K = \text{Schur}(L(G^*)/O)$	synthetic only
(sub-str.)	bridge;	$R_{\text{eff}}$ distinguishes	synthetic bridge topology

### 5.1 Partial CHP in the 703-corpus (natural regime break)

In the `airtable-mcp+mcp-xmind` composition, the dimension `path_convention_match` has 5 carrier columns: 3 hidden and 2 observable, split by whether each tool classifies its `path` field as a convention dimension or a primary input. The classifier disagreement across tools is structural, not accidental: `read_xmind` exposes `path` as a primary tool input (observable by design), while `extract_node` classifies its own `path` as a convention dimension (hidden). The same field name, different semantic roles.

The carrier graph is  $K_5$  (complete on the 5 tool-field pairs, because every pair of `mcp-xmind` tools in the composition has a bilateral row). Absorbing the 2 observable vertices via Kron reduction gives

$$K_d = \text{Schur}(L(K_5) / \{41, 42\}) = \frac{5}{3} \cdot L(K_3) = \begin{bmatrix} 10/3 & -5/3 & -5/3 \\ -5/3 & 10/3 & -5/3 \\ -5/3 & -5/3 & 10/3 \end{bmatrix}.$$

Leverage is  $\frac{2}{3}$  on all three hidden fields (per Corollary 4.1 with  $|H_C| = 3$ ), matching the observed value exactly. All three pairwise effective resistances are  $\frac{2}{5}$ : the three fields are equally good disclosure substitutes.

Thirty-seven compositions in the 703-corpus exhibit this partial-CHP regime; every one of them involves `mcp-xmind` paired with another server. The pattern is a single-server property propagated across all pairings, not a combinatorial interaction between server pairs.

### 5.2 DFD violation (synthetic)

The synthetic composition `regime_break_dfd_violation.yaml` has two tools with all fields hidden (full CHP) and a single edge whose dimensions share a `from`-field: both `dim1` and `dim2` reference `tool_a.x`, violating DFD. The coboundary is

$$\delta_{\text{full}} = \begin{bmatrix} -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \\ 0 & -1 & 1 & 0 \end{bmatrix} \quad (\text{rows: dim1, dim2, dim3; cols: x, y, u, v}),$$

and the witness Gram is

$$K(G) = \begin{bmatrix} 2 & 0 & -1 & -1 \\ 0 & 1 & -1 & 0 \\ -1 & -1 & 2 & 0 \\ -1 & 0 & 0 & 1 \end{bmatrix}.$$

$K(G)$  is *not* block-diagonal: the column `x` has off-diagonal entries  $-1$  to both `u` and `v` (two dimensions' worth). But  $K(G)$  *is* still a graph Laplacian — specifically,  $K(G) = L(G^*)$  where  $G^*$  is the coboundary-incidence graph on 4 vertices with edges  $\{(x, u), (x, v), (y, u)\}$  (the three dimensions' bilateral pairs). The leverage is  $\frac{3}{4}$  uniformly across all four hidden fields (because  $G^*$  is connected with  $|H_C| = 4$ ), and the fee is 3.

This is the non-DFD case of Theorem 3.5: block decomposition fails but the Laplacian structure persists on the merged graph.

### 5.3 Bridge topology (synthetic)

The synthetic composition `regime_break_bridge_topology.yaml` has three tools  $\{A, B, C\}$  each carrying `path_convention_match` via their `path` field, with bilateral edges  $A \rightarrow B$  and  $B \rightarrow C$  (no  $A \rightarrow C$  edge). Full CHP and DFD hold. The carrier graph is the path  $P_3$  on three vertices; the witness Gram is

$$K(G) = L(P_3) = \begin{bmatrix} 1 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 1 \end{bmatrix}.$$

Leverage is  $\frac{2}{3}$  uniformly (Corollary 4.1 with  $|H_C| = 3$ ), but effective resistances distinguish the bridge:  $R_{\text{eff}}(A, B) = R_{\text{eff}}(B, C) = 1$  while  $R_{\text{eff}}(A, C) = 2$ . The vertex  $B$  has lower average effective resistance ( $\bar{R} = 1$ ) than  $A$  or  $C$  ( $\bar{R} = 1.5$ ); disclosing  $B$  substitutes for either of the other two, while disclosing  $A$  barely substitutes for  $C$ .

This example demonstrates that effective resistance (Corollary 4.5) carries structural information that leverage alone cannot: the three fields have identical leverage but qualitatively different roles in the witness geometry.

## 6 Corpus confirmation

All claims below are verified in exact rational arithmetic on the 703-composition real-schema MCP registry corpus of [2, §1.1]. Source: `bullacalibration/witness_geometry_sweep.py` and `verify_laplacian_collapse.py`.

Table 2: Structural identities verified (regime-independent).

Identity	Pass	Fail
$\text{rank}(K(G)) = \text{fee}(G)$	703/703	0
$\sum_j \ell_j = \text{fee}(G)$ (exact $\mathbb{Q}$ )	703/703	0
$0 \leq \ell_j \leq 1$ for all hidden fields	703/703	0

Table 3: Leverage-prediction match via Kron formula (nontrivial compositions only).

Prediction match	Count
$\ell_j$ matches $( H_C  - 1)/ H_C $ per component (exact $\mathbb{Q}$ )	240/240
Mismatches	0

Table 4: Kron three-case taxonomy on the 703-corpus.

Case	Description	Count	Leverage formula
1	Full CHP: $O_d = \emptyset$ ; $K_d = L_d$	203	$( V_d  - 1)/ V_d $
2	Trivial Kron: $O_d \neq \emptyset$ but $L_d[O_d, O_d] = 0$	0	$( H_d  - 1)/ H_d $
3	Genuine Kron reduction: $L_d[O_d, O_d] \neq 0$	37	$( H_d  - 1)/ H_d $
	Trivial ( $ H  = 0$ )	300	n/a
	Fee zero ( $ H  > 0$ , $\text{fee} = 0$ )	163	n/a

Three corpus-wide observations deserve comment.

**Case 2 is empty.** On the 703-corpus every partial-CHP composition exercises the full Kron reduction (nonzero inter-observable edges). There is no degenerate middle case: either full CHP

(Case 1, 203 compositions) or genuine Kron absorption (Case 3, 37 compositions). This gives the corpus a clean shape and simplifies §3’s narrative.

**Intra-dimension leverage uniformity.** No composition in the corpus exhibits non-uniform leverage within a single component. The theorem predicts this (Corollary 4.1: uniform within a component); the empirical confirmation rules out cases where disconnected hidden-carrier components might split a dimension into sub-components with distinct leverage values.

**Case 3 concentrates on a single server.** All 37 partial-CHP compositions involve `mcp-xmind`. The classifier’s disagreement on `path` fields across `mcp-xmind`’s tools propagates across every pairing with another server, producing 37 Case-3 compositions from a single structural source. The frontier where field-level diagnostics become discriminative is thus already present in the corpus, but it is a single-server phenomenon — not a combinatorial artifact of server pairs.

## 7 Discussion

The witness Gram gives the Bulla coherence fee a canonical form that the integer alone cannot: a Kron-reduced graph Laplacian whose spectrum, bases, and effective-resistance structure encode the composition’s hidden convention geometry. Four operational consequences:

1. **Closed-form leverage.** Within each connected component of the Kron-reduced effective graph, leverage is  $(|H_C| - 1)/|H_C|$  — a function of component size alone. Bulla’s CLI can report leverage by tabulating component sizes, no matrix operations required.
2. **Spanning-tree-many bases.** The number of minimum disclosure sets is  $\prod_d \tau(\tilde{G}_d)$ . A product of spanning-tree counts. Most 703-corpus compositions have  $\tau = 1$  per dimension (unique disclosure set); `mcp-xmind` compositions have  $\tau > 1$  (multiple valid substitutes).
3. **Weighted repair is matroid optimization.** The minimum-cost full repair is a minimum-weight basis of the witness matroid, solved by the classical greedy algorithm [7]. Not a heuristic.
4. **Effective resistance ranks substitutes.** Within a component, the effective resistance between two hidden fields measures how much one disclosure substitutes for the other. Disclosure planning on real compositions can use this to rank alternative repair choices when privacy, cost, or operational constraints rule out the greedy-optimal field.

**The regime-dependent vs regime-independent split.** Three of the witness Gram’s invariants — rank, leverage, and basis count — are closed-form functions of the effective-graph topology alone. On the 703-corpus this means that under the current Bulla classifier, they are computable by hand from carrier-count vectors without any matrix work. The fourth invariant — effective resistance (Corollary 4.5) — is where the geometry earns its keep: it distinguishes bridge structures that leverage cannot (§5.3), ranks disclosure substitutes when multiple minimum bases exist, and exposes the local-vs-distributed shape of hidden convention risk.

**Research frontier.** The foundational theorem — that hidden convention debt is a rank, and that rank has a canonical electrical-network interpretation — is now complete. The natural open questions shift from “what is the witness geometry?” to two more pointed ones. First: what happens when future MCP server designs, or future Bulla classifier revisions, partially expose convention fields? Theorem 3.5 covers all such cases in closed form, but the empirical phase diagram is open — which regime cells (DFD  $\times$  CHP) do evolving MCP ecosystems populate? Second: does the Kron-reduction framing open formal bridges to the electrical-network, random-walk, and resistance-distance literatures whose results (e.g., network simplification, Markov-chain commute times) could specialize to Bulla with non-trivial consequences? Both are outside the scope of this paper but within the scope of the program.

## A Lean verification status

The two regime-independent identities of this paper — the rank identity (Theorem 2.2) and leverage conservation (the foundation of Corollary 4.1’s sum) — are formally verified in Lean 4 against Mathlib (Aristotle/Harmonic project UUID 3c1a38f9-a823-4b80-ae5d-7e4dfaacad85). Two theorems closed, zero `sorry`, standard axioms only:

Theorem	Lean name	Paper reference	Key Mathlib
1	<code>witness_gram_rank_eq_fee</code>	Theorem 2.2 (rank identity)	<code>finrank_quotient_add_finrank</code>
2	<code>leverage_conservation</code>	Corollary 4.1 (basis of $\Sigma \ell_j = \text{fee}$ )	<code>trace_comp_comm'</code> , <code>trace_id</code>

Theorem 1 establishes the rank-difference identity in finrank-quotient form: under the column-compression hypothesis  $\delta_{\text{obs}} = \delta_{\text{full}} \circ \iota$ , the witness-Gram rank equals  $\text{finrank}(\text{range } \delta_{\text{full}}) - \text{finrank}(\text{range } \delta_{\text{obs}})$ . Theorem 2 establishes that for any idempotent endomorphism  $P$  of a finite-dimensional  $\mathbb{Q}$ -vector space,  $\text{trace } P = \text{finrank}(\text{range } P)$ , which directly gives leverage conservation when  $P = K^+K$  is the projection onto  $\text{col}(K)$ .

Combined with the twelve Lean-verified theorems of the backbone paper [2] (UUIDs 044a00b1, ff78c99c, edb6c604, 01333bb4), the program now has fourteen Mathlib-verified theorems with zero `sorry`. The remaining paper-specific claims — the Kron-reduction theorem (Theorem 3.3) and the closed-form leverage formula  $\ell_j = (|H_C| - 1)/|H_C|$  (Corollary 4.1) — are deferred to follow-on Lean work, since each requires formalizing Kron-reduction preservation of connectivity (Dörfler-Bullo Theorem 3.1) and the signed-incidence Laplacian identity, neither currently in Mathlib.

## B Code and data artifacts

All code, data, and verification scripts are available at <https://github.com/jkomkov/bulla>. Canonical artifacts:

- `bulla/src/bulla/witness_geometry.py` (module)
- `bulla/calibration/witness_geometry_sweep.py` (703-corpus sweep)
- `bulla/calibration/verify_laplacian_collapse.py` (240/240 leverage verification)
- `bulla/calibration/kron_three_cases.py` (Case 1/2/3 classification)
- `bulla/compositions/regime_break*.yaml` (synthetic regime-break examples)
- `papers/hierarchical-fee/laplacian_collapse.md` (extended proof memo)
- `papers/hierarchical-fee/THEOREM-LOCK.md` (vocabulary and statement lock)

## References

- [1] J. Komkov, “The Coherence Fee: Measuring Hidden Convention Risk in Hierarchical Tool Compositions,” 2026.
- [2] J. Komkov, “The Column-Matroid Backbone of a Composition-Coherence Fee: a Structural Identity and Its Corollaries,” 2026.
- [3] G. Kron, “Tensor Analysis of Networks,” John Wiley & Sons, 1939.
- [4] F. Dörfler and F. Bullo, “Kron Reduction of Graphs With Applications to Electrical Networks,” *IEEE Trans. Circuits and Systems I*, vol. 60, no. 1, pp. 150-163, 2013.
- [5] D. J. Klein and M. Randić, “Resistance Distance,” *J. Math. Chem.*, vol. 12, pp. 81-95, 1993.

- [6] C. Godsil and G. Royle, “Algebraic Graph Theory,” Springer, 2001.
- [7] J. Edmonds, “Matroids and the greedy algorithm,” *Math. Programming*, vol. 1, pp. 127-136, 1971.
- [8] J. Komkov, “Signed-Incidence Structure in Compositional Verification: A Structural Note,” [papers/signed-incidence/paper/signed-incidence.pdf](#), 2026.