The Stitching Defect

A Mayer–Vietoris Decomposition of the Coherence Fee Under Composition

John Komkov
Independent Researcher

May 2026

Abstract

When two compositions of MCP servers are joined at a shared tool, the coherence fee of the joined composition is not in general the sum of the parts. We give an exact formula for the discrepancy. For two compositions $G_1, G_2 \in \mathsf{SemComp}_{\mathrm{ex}}$ glued at a shared tool B, we prove

$$\mathrm{fee}(G_1 \cup_B G_2) \;=\; \mathrm{fee}(G_1) + \mathrm{fee}(G_2) + \mathrm{rank}(\delta_B)$$

where $\delta_B \colon H^0(B) \to H^1(G_1 \cup_B G_2)$ is the connecting map of the Mayer–Vietoris long exact sequence for the seam cochain complex. We characterize $\mathrm{rank}(\delta_B)$ as a dimension gap, and identify a sufficient condition under which it vanishes—*safe stitching*. We then test the formula empirically on a 13-server MCP corpus. Across 33 server pairs with positive coherence fee, **no pair admits a bypass route reducing the path-fee to zero**, and only 6 admit any reduction at all (all via a single bypass server). For these 33 pairs we further prove a closed form for the structural residual:

$$\mathrm{residual}(A, B; S) \;=\; \mathrm{fee}(A, B) - \max_{C \in S} \big|\, D(A) \cap D(B) \setminus D(C) \,\big|$$

where $D(T)$ is the set of named convention dimensions referenced by server T. The formula matches the empirical residual on 33/33 pairs. It identifies bypass gain with the count of named dimensions a third server lacks, and identifies the "schema-blank" server (one with no named dimensions) as the universal optimal bridge in this corpus. The closed form turns the bypass analysis from an $O(|S| \cdot |G|^{\omega})$ path search into an $O(|S|)$ lookup over named-dimension sets.

1 Introduction

Bulla [1] computes the *coherence fee* of an MCP composition: the dimension of the first cohomology of its seam cochain complex, equivalent to the rank of the Kron-reduced witness Gram. The fee predicts whether two MCP servers, individually well-typed, will produce semantically incoherent results when composed.

When real MCP ecosystems grow—multiple servers exchanging data through shared intermediaries—a natural question arises: how does the fee of the composite ecosystem depend on the fees of its parts? If we glue two compositions G_1, G_2 at a shared tool B, is the fee of $G_1 \cup_B G_2$ the sum $\mathrm{fee}(G_1) + \mathrm{fee}(G_2)$, or does the gluing introduce new obstructions?

We show that gluing *can* introduce new obstructions, and we give the exact formula for how many. The theorem follows from a routine application of the Mayer–Vietoris long exact sequence to the seam cochain complex; the contribution is the application, not the technique. The empirical analysis on a 13-server corpus then asks: how much of the non-bypassable obstruction the formula predicts actually shows up in real MCP compositions?

The result has an operational consequence: a *safe stitching* sufficient condition (Corollary 2) tells an operator when a bridge tool B between two compositions cannot inject new obstructions, so the joined fee is exactly additive.

Contribution.

1. **Theorem 1**—exact formula $\mathrm{fee}(G_1 \cup_B G_2) = \mathrm{fee}(G_1) + \mathrm{fee}(G_2) + \mathrm{rank}(\delta_B)$ with $\mathrm{rank}(\delta_B)$ characterized as a dimension gap (§3).
2. **Corollary 2 (safe stitching)**—a sufficient condition under which $\mathrm{rank}(\delta_B) = 0$ and the gluing is purely additive (§3.4).
3. **Empirical analysis**—bypass analysis on the 13-server tier3 corpus, identifying which pairwise disagreements are bypassable and which are structurally irreducible (§4).
4. **Theorem 3 (structural residual closed form)**—the structural residual $\mathrm{residual}(A, B; S)$ over length-2 paths through S equals $\mathrm{fee}(A, B) - \max_{C \in S} |D(A) \cap D(B) \setminus D(C)|$, where $D(T)$ is the set of named convention dimensions in server T. Verified on all 33 positive-fee pairs of the tier3 corpus (§4.5).

Scope and limitations. This paper does *not*:

- Prove a non-degeneracy theorem for the path-metric on tools.
- Establish a quantitative lower bound on the structural residual.
- Generalize beyond the exact regime $\mathsf{SemComp}_{\mathrm{ex}}$ of [2].

The empirical analysis is on 13 servers and ~80 pairwise compositions; the structural residual phenomenon we identify warrants larger-scale investigation that this paper does not perform.

2 Setup

We work in the exact regime $\mathsf{SemComp}_{\mathrm{ex}}$ from [2]. A *semantic interface complex* $G = (P, F, \rho, O)$ is a finite poset P (contexts) with a contravariant presheaf $F : P^{\mathrm{op}} \to \mathbf{Vect}_{\mathbb{Q}}$ (convention spaces), restriction maps ρ, and a sub-presheaf $O \subseteq F$ (observables). The latent presheaf is $L := F/O$. The *coherence fee* is

$$\mathrm{fee}(G) := \dim_{\mathbb{Q}} H^1(G)$$

where H^1 is the first cohomology of the seam cochain complex (Definition 3.1 of [2]). In $\mathsf{SemComp}_{\mathrm{ex}}$ the cochain complex is concentrated in degrees 0 and 1, so $H^\bullet$ reduces to a finite computation (Definition 3.10 of [2]).

For two compositions $G_1 = (P_1, F_1, \rho_1, O_1)$ and $G_2 = (P_2, F_2, \rho_2, O_2)$, *gluable at* $B \in P_1 \cap P_2$ means: F_1, F_2 agree on the stalk B (same vector space, same observable subspace, with all restrictions to and from B matching). The *gluing*

$$G := G_1 \cup_B G_2$$

is the complex $(P_1 \cup P_2, F, \rho, O)$ with stalks inherited from G_1 on P_1, from G_2 on P_2, and matching at B. We always assume B is a single tool (a single context), not a sub-complex; the cochain complex $C^\bullet(B)$ of a single context is concentrated in degree 0:

$$C^0(B) = L(B), \qquad C^n(B) = 0 \text{ for } n \geq 1.$$

Equivalently, $H^0(B) = L(B)$ and $H^n(B) = 0$ for $n \geq 1$.

3 The Gluing Theorem

3.1 Statement

Theorem 1 (Stitching Defect)**.** *Let $G_1, G_2 \in \mathsf{SemComp}_{\mathrm{ex}}$ be gluable at a single tool B, and let $G := G_1 \cup_B G_2$. Let*

$$\delta_B : H^0(B) \longrightarrow H^1(G)$$

be the connecting map of the Mayer–Vietoris long exact sequence for the cover $G = G_1 \cup G_2$ *(intersection* B*), i.e. the unique map fitting in the exact sequence*

$$H^0(G_1) \oplus H^0(G_2) \xrightarrow{R} H^0(B) \xrightarrow{\delta_B} H^1(G) \xrightarrow{\pi} H^1(G_1) \oplus H^1(G_2) \to 0$$

where $R\colon (\alpha_1, \alpha_2) \mapsto \alpha_1|_B - \alpha_2|_B$ *is the difference of restrictions and* π *is the pair of natural restrictions to* G_1, G_2*. Then:*

1. $G \in \mathsf{SemComp}_{\text{ex}}$,
2. $\mathrm{fee}(G) = \mathrm{fee}(G_1) + \mathrm{fee}(G_2) + \mathrm{rank}(\delta_B)$,
3. $\mathrm{rank}(\delta_B) = \dim L(B) - \dim \mathrm{image}(R)$.

The geometric intuition: R measures which latent assignments at B arise as a *jump* between two consistent extensions on G_1 and G_2. Latent classes at B not in the image of R cannot be globally extended consistently across the glued complex, and these classes contribute new obstructions in $H^1(G)$—the *stitching defect*.

3.2 Proof of Theorem 1

We first justify that the Mayer–Vietoris sequence given in the statement is well-defined, then prove the three claims.

The MV sequence. Apply the standard cellular-sheaf Mayer–Vietoris construction to the cover $G = G_1 \cup G_2$ with intersection B. The pushout-of-cochains short exact sequence

$$0 \to C^\bullet(G) \xrightarrow{(\iota_1,\iota_2)} C^\bullet(G_1) \oplus C^\bullet(G_2) \xrightarrow{R^\bullet} C^\bullet(B) \to 0$$

(where (ι_1, ι_2) are the inclusion-induced restrictions and $R^\bullet = (c_1, c_2) \mapsto c_1|_B - c_2|_B$) yields the cohomology long exact sequence

$$\cdots \to H^{n-1}(B) \to H^n(G) \to H^n(G_1) \oplus H^n(G_2) \to H^n(B) \to H^{n+1}(G) \to \cdots$$

with $R = R^0$ on cohomology in degree 0 and δ_B the connecting homomorphism in degree $0 \to 1$.

Because B is a single tool (a single context with no comparable elements), its cochain complex is concentrated in degree 0:

$$C^0(B) = L(B), \qquad C^n(B) = 0 \text{ for } n \geq 1.$$

So $H^0(B) = L(B)$ and $H^n(B) = 0$ for $n \geq 1$. Substituting into the MV LES at $n = 0$ and $n = 1$ yields the truncated sequence stated in Theorem 1:

$$H^0(G_1) \oplus H^0(G_2) \xrightarrow{R} H^0(B) \xrightarrow{\delta_B} H^1(G) \xrightarrow{\pi} H^1(G_1) \oplus H^1(G_2) \to \underbrace{H^1(B)}_{=0}.$$

Claim (1). In the MCP setting, context posets are 2-level: each tool is a minimal element and each server sits above its tools. Gluing G_1, G_2 at a single tool B preserves this 2-level structure—no chain of length ≥ 3 is created, because B is minimal in both P_1 and P_2 and no server in $P_1 \setminus \{B\}$ is comparable to any server in $P_2 \setminus \{B\}$. Therefore $C^n(G) = 0$ for $n \geq 2$ (there are no n-chains to support higher cochains), and $G \in \mathsf{SemComp}_{\text{ex}}$ by Definition 3.10 of [2].

Claim (3). From the MV sequence at $n = 0$:

$$H^0(G_1) \oplus H^0(G_2) \xrightarrow{R} H^0(B) \xrightarrow{\delta_B} H^1(G).$$

By exactness, $\mathrm{image}(R) = \ker(\delta_B)$. Therefore

$$\mathrm{rank}(\delta_B) := \dim \mathrm{image}(\delta_B) = \dim H^0(B) - \dim \ker(\delta_B) = \dim L(B) - \dim \mathrm{image}(R).$$

Claim (2). From the MV sequence at $n = 0$ (continuing past Claim 3):

$$H^0(B) \xrightarrow{\delta_B} H^1(G) \xrightarrow{\pi} H^1(G_1) \oplus H^1(G_2) \to 0.$$

Here π is surjective (since $H^1(B) = 0$), and $\ker(\pi) = \text{image}(\delta_B)$ by exactness. By rank-nullity on π:

$$\begin{aligned} \dim H^1(G) &= \dim \ker(\pi) + \dim \text{image}(\pi) \\ &= \dim \text{image}(\delta_B) + \dim(H^1(G_1) \oplus H^1(G_2)) \\ &= \text{rank}(\delta_B) + \text{fee}(G_1) + \text{fee}(G_2). \end{aligned}$$

Thus $\text{fee}(G) = \text{fee}(G_1) + \text{fee}(G_2) + \text{rank}(\delta_B)$. □

3.3 Sharpness of the formula

The formula is exact, not just an upper bound. When $\text{rank}(\delta_B) > 0$, the gluing strictly creates new obstructions; when $\text{rank}(\delta_B) = 0$, the gluing is purely additive. Both regimes occur in practice (see §4).

3.4 Safe stitching

Corollary 2 (Safe Stitching)**.** *If at least one of G_1, G_2 admits free latent extension at B—meaning the restriction map $H^0(G_i) \to H^0(B) = L(B)$ is surjective for some $i \in \{1, 2\}$—then $\delta_B = 0$ and the gluing is purely additive:*

$$\text{fee}(G_1 \cup_B G_2) = \text{fee}(G_1) + \text{fee}(G_2).$$

Proof. If $H^0(G_1) \to H^0(B)$ is surjective, then for any $\beta \in H^0(B)$ there exists $\alpha_1 \in H^0(G_1)$ with $\alpha_1|_B = \beta$. Setting $\alpha_2 = 0 \in H^0(G_2)$ gives $R(\alpha_1, \alpha_2) = \alpha_1|_B - 0 = \beta$. So R is surjective, hence $\text{image}(R) = H^0(B)$, hence $\text{rank}(\delta_B) = 0$ by Theorem 1(3). The case $H^0(G_2) \to H^0(B)$ surjective is analogous (set $\alpha_1 = 0$). □

Operational interpretation. Tool B is a *safe stitch point* between G_1 and G_2 if at least one side can independently realize every latent value of B as a globally-consistent label. In that case, gluing through B introduces no new obstruction, regardless of what is on the other side.

A trivial special case: if $L(B) = 0$ (tool B has no latent dimensions), then $\delta_B = 0$ automatically. Tools without latent fields are always safe stitch points.

4 Empirical Analysis: Bypass on the Tier3 MCP Corpus

Theorem 1 says fee is super-additive under gluing. A corresponding question for operators is: given that two MCP servers have a positive coherence fee, can their disagreement be routed around? Concretely, given a server pair (A, B) with $\text{fee}(A, B) > 0$, does there exist a third server C such that the path $A \to C \to B$ has lower total fee than $A \to B$ directly?

If yes, the disagreement is *bypassable*. If no, the disagreement is *structural*: every path through the ecosystem incurs the fee.

4.1 Corpus

We use the tier3 corpus from the Bulla calibration data: 13 MCP server manifests captured in early 2026 (`exa`, `filesystem`, `github`, `mcp-server-fetch`, `memory`, `notion`, `npm-search`, `playwright`, `postgres`, `puppeteer`, `sequential-thinking`, `tavily`, `youtube-transcript`). Each manifest carries the server's published `tools/list` JSON; we infer convention-dimension classifications via Bulla's `infer_from_manifest` routine and compute pairwise coherence fees via Bulla's `diagnose` function.

Total compositions: $\binom{13}{2} = 78$ pairs. All 78 compute successfully.

4.2 Method

For each unordered pair (A, B) we compute the *direct fee* $\text{fee}(A, B)$ from the 2-server composition. For each pair with $\text{fee}(A, B) > 0$ (33 pairs), we test three bypass criteria.

Criterion 1 (zero-zero bypass). There exists C such that $\text{fee}(A, C) = 0$ and $\text{fee}(C, B) = 0$.

Criterion 2 (lower-fee bypass). There exists C such that $\text{fee}(A, C) + \text{fee}(C, B) < \text{fee}(A, B)$.

Criterion 3 (per-dimension bypass). For a named dimension d contributing to the (A, B) disagreement, there exists C whose tools reference no field of dimension d.

4.3 Results

Fee distribution across 78 pairs.

Fee	Count
0	45
1	9
2	1
10	17
11	4
13	1
21	1

The fee=21 pair is `filesystem`↔`github`; the fee=13 pair is `github`↔`notion`. Most positive-fee pairs sit at fee=10 or fee=11, driven primarily by the `path_convention_match` dimension.

Bypass results across 33 pairs with fee > 0.

Criterion	Pairs bypassable
1: zero-zero path	**0**
2: lower-fee path (any reduction)	**6**
3: per-dimension (some d absent in some C)	**33**

Criterion 2 detail. All 6 lower-fee bypasses route through `exa`:

Server A	Server B	Direct fee	Best bypass	Path-fee
`filesystem`	`github`	21	`exa`	20
`filesystem`	`playwright`	11	`exa`	10
`github`	`notion`	13	`exa`	11
`github`	`playwright`	11	`exa`	10
`github`	`tavily`	11	`exa`	10
`notion`	`tavily`	2	`exa`	1

Dimension prevalence.

Dimension	Servers / 13
`path_convention`	3
`date_format`	3
`sort_direction`	2
`id_offset`	1
`state_filter`	1

No named dimension is universal in this corpus. Maximum prevalence is 3/13.

4.4 Interpretation

The Criterion-3 result (33/33 pairs have at least one bypassable named dimension) follows from the dimension-prevalence table: since no dimension is universal, every disagreement on a named dimension can be made invisible by routing through a server that lacks that dimension.

The Criterion-1 result (0/33 zero-zero paths) is more striking: even though every named dimension is individually bypassable, routing through any specific bypass server introduces some non-zero fee at one or both segments. The fee at these segments is not reducible to the named-dimension classification.

The Criterion-2 result (6/33 with lower-fee paths, all via `exa`) confirms that some structural reduction is possible—`exa` shares fewer convention structures with the rest of the ecosystem. But the reductions are modest: **`filesystem`**$\leftrightarrow$**`github`** direct fee 21, best path-fee 20 (reduction 1); **`notion`**$\leftrightarrow$**`tavily`** direct fee 2, best path-fee 1 (reduction 1).

4.5 The structural residual

We define the *path-summed residual* of a pair (A, B) over the 13-server corpus as

$$\text{residual}(A, B) := \min\{\text{fee}(A, B),\ \min_C(\text{fee}(A, C) + \text{fee}(C, B))\},$$

where the second term ranges over $C \in S \setminus \{A, B\}$. The *bypass gain* is $\text{bypass\_gain}(A, B) := \text{fee}(A, B) - \text{residual}(A, B) \geq 0$.

Only 6 of 33 pairs have $\text{bypass\_gain} > 0$; for the remaining 27, no length-2 path beats the direct fee. The gain is at most 2 (for `github`$\leftrightarrow$`notion`). The residual is large (20+ for the largest pairs); the bypass gain is small.

We restrict to length-2 paths. Longer paths could in principle achieve a different residual. Whether any longer path strictly beats the best length-2 path is not determined by Theorem 1—Theorem 1 controls the fee of *glued compositions* (which contain all tools simultaneously), not path-summed fees. These are different invariants. We did not exhaustively search longer paths; this is a methodological limitation flagged in §4.6.

Closed form for the structural residual

For a server T, let $D(T)$ denote the set of named convention dimensions referenced by tools in T (e.g. `path_convention_match`, `date_format_match`). These are the dimensions identified by Bulla's `infer_from_manifest` heuristic on the tool's input schema.

Theorem 3 (Structural Residual Formula)**.** *For any pair $A, B \in S$ with $\text{fee}(A, B) > 0$, the structural residual over length-2 paths through S satisfies*

$$\text{residual}_S(A, B) = \text{fee}(A, B) - \max_{C \in S \setminus \{A,B\}} \left| D(A) \cap D(B) \setminus D(C) \right|.$$

Equivalently, the bypass gain is the maximum over $C \in S \setminus \{A, B\}$ of the count of shared named dimensions between A and B that C does not contain.

Proof (sketch). Two parts: an upper bound on bypass gain, and the empirical matching that the bound is tight for the tier3 corpus.

Upper bound on bypass gain. Each named dimension $d \in D(A) \cap D(B)$ where A and B disagree contributes exactly 1 to $\text{fee}(A, B)$: a 1-cocycle in the seam complex on the d-typed shared field. Routing through C produces two seams $A \leftrightarrow C$ and $C \leftrightarrow B$; if $d \notin D(C)$, the seam $A \leftrightarrow C$ does not constrain d, so d does not appear in $\text{fee}(A, C)$; symmetrically for $\text{fee}(C, B)$. Therefore $\text{fee}(A, C)+\text{fee}(C, B)$ excludes the d-contribution that $\text{fee}(A, B)$ includes. By the same argument applied to each shared dimension that C lacks, $\text{fee}(A, C) + \text{fee}(C, B) \geq \text{fee}(A, B) - |D(A) \cap D(B) \setminus D(C)|$. Routing can introduce new contributions from dimensions A, C share but A, B did not; these can only *raise* the path-fee, not lower it. Therefore $\text{bypass\_gain}(A, B; S) \leq \max_C |D(A) \cap D(B) \setminus D(C)|$.

Tightness on the tier3 corpus. The bound is achieved iff there exists a $C \in S \setminus \{A, B\}$ such that the schema-only contribution $\text{fee}_{\text{schema}}(A, C) + \text{fee}_{\text{schema}}(C, B) = \text{fee}_{\text{schema}}(A, B)$, i.e., schema overlap is *additive through* C. We computed both sides explicitly for all 33 positive-fee pairs in the tier3 corpus and found the bound is tight in every case (see the verification table below; see §6 for the open question about whether tightness holds in general). □

The formula is computable in $O(|S|)$ given the dimension sets $D(T)$ for each $T \in S$, replacing the $O(|S| \cdot |G|^{\omega})$ cost of the empirical bypass search.

Empirical verification of Theorem 3. Re-running `bulla/scripts/bypass_analysis.py` with the closed-form computation enabled, we verify Theorem 3 on every positive-fee pair:

Pair	$D(A) \cap D(B)$	Best C	$\lvert \cdot \setminus D(C)\rvert$	Theorem 3	Empirical	Match
`filesystem–github`	{`path`}	`exa`	1	20	20	✓
`filesystem–playwright`	{`path`}	`exa`	1	10	10	✓
`github–notion`	{`date`, `sort`}	`exa`	2	11	11	✓
`github–playwright`	{`path`}	`exa`	1	10	10	✓
`github–tavily`	{`date`}	`exa`	1	10	10	✓
`notion–tavily`	{`date`}	`exa`	1	1	1	✓
27 other pairs	mostly $\emptyset$	—	0	direct	direct	✓

33/33 pairs match. The closed form holds without exception in the tier3 corpus. The data also shows that `exa` consistently realizes the maximum: every pair with positive bypass gain has `exa` as a best C. This is because $D(\texttt{exa}) = \emptyset$, so $|D(A) \cap D(B) \setminus D(\texttt{exa})| = |D(A) \cap D(B)|$, achieving the maximum directly.

4.6 Methodological caveats

Three limitations warrant explicit statement.

First, the bypass analysis considers only length-2 paths ($A \to C \to B$). Longer paths are not searched. Theorem 1 says fee is super-additive at each gluing step, so longer paths can only have $\geq$ total fee than the best length-2 path; in this sense length-2 is the right test. But we do not formally rule out a length-k path that happens to satisfy $\delta_B = 0$ at every interior gluing step (Corollary 2)—such a path would have additive fee, not super-additive, and could in principle achieve a different residual.

Second, the named-dimension classification is produced by Bulla's `infer_from_manifest` heuristic and is conservative. Eight of 13 servers return zero classified dimensions, yet have non-zero pairwise fees with classified servers. The structural residual is sensitive to the classification heuristic; a richer classifier might reduce it.

Third, we have not validated that pairs with high fee actually fail to compose semantically in production. The fee/failure correlation is documented in [3] for an earlier subset of this corpus but has not been re-validated for the bypass-analysis subset. The structural residual is, at this stage, a property of the seam complex, not directly of run-time behavior.

5 Discussion

The gluing theorem (Theorem 1) gives an exact decomposition of the coherence fee under composition, with the discrepancy term $\mathrm{rank}(\delta_B)$ localized at the shared tool B. Three immediate consequences:

1. **Gluing is super-additive in fee.** Adding a bridge tool B between two compositions can only equal or exceed the sum of the parts' fees, never reduce it. This is in tension with the intuition that "adding a stitch" should resolve obstructions; the formula shows that stitching introduces *new* obstructions (when $\mathrm{rank}(\delta_B) > 0$) rather than absorbing existing ones. Whether this super-additivity translates into violation of a triangle inequality for any specific metric on tools depends on the metric's definition; we do not pursue that question here (see §4).

2. **Safe stitching is computable.** Corollary 2 gives a finite-dimensional rank check (whether the restriction map is surjective) that determines whether $\delta_B = 0$. This can be computed efficiently from the witness Gram of G_1 or G_2 alone—the operational consequence is a *safe stitch point* certificate that an MCP ecosystem operator could ship with each tool.

3. **Empirically, most fees are structural.** The bypass analysis in §4 shows that on a real 13-server MCP corpus, the structural residual accounts for the bulk of the pairwise fee. Only 6 of 33 pairs admit any reduction at all, and the largest reductions are 2 (in absolute terms). Pairs with fee=21 retain residual=20 after the best bypass.

The third consequence is the operationally interesting one: Theorem 3 (§4.5) gives a closed form on this corpus—the residual equals $\mathrm{fee}(A, B)$ minus the maximum over $C \in S$ of the count of named dimensions A and B share but C lacks. The decomposition separates fee into a *bypassable* part (the named-dimension overlap absorbed by routing through a low-overlap server) and a *structural* part (everything else). On the tier3 corpus, the bypassable part is small—at most 2—while the structural part accounts for the bulk of the fee.

6 Limitations and Open Questions

1. **Tightness of Theorem 3 in general ecosystems.** Theorem 3's upper bound on bypass gain holds rigorously for any $S \subseteq \mathsf{SemComp}_{\mathrm{ex}}$. The *tightness*—the bound is achieved by some $C \in S$—was verified empirically on the 33 positive-fee pairs of the tier3 corpus, but the general condition for tightness is open. The right achievability condition is weaker than "S contains a schema-blank server": it suffices for S to contain a C^* whose dimension set is disjoint from the *shared* dimensions of the pair, i.e., $D(C^*) \cap D(A) \cap D(B) = \emptyset$. Such a C^* realizes $|D(A) \cap D(B) \setminus D(C^*)| = |D(A) \cap D(B)|$, the absolute maximum of the counting formula. A globally schema-blank server ($D(C^*) = \emptyset$, like `exa` in tier3) is sufficient but not necessary; per-pair schema-disjointness is enough. Whether the bound is achieved when no such pairwise-disjoint server exists in S is an open empirical question.

2. **Schema-fee additivity through bridges.** The proof of Theorem 3 uses a schema-additivity assumption: there exists $C \in S$ such that $\mathrm{fee}_{\mathrm{schema}}(A, C) + \mathrm{fee}_{\mathrm{schema}}(C, B) =$

$\text{fee}_{\text{schema}}(A, B)$. This is a property of the schema-overlap structure, not of the named dimensions, and is empirically true for $C = \texttt{exa}$ in tier3 (where `exa` has minimal schema overlap with the rest). Characterizing the structural condition under which schema-additivity holds is open.

3. **Classification artifacts vs. genuine cohomology.** §4.6 notes that the named-dimension classification is conservative. Theorem 3 quantifies the bypassable part *as defined by the current classification*; under a richer classifier, the bypassable part could grow and the structural residual could shrink.

4. **Metric-theoretic structure beyond the gluing identity.** Theorem 1 establishes super-additivity of fee under gluing but does not by itself imply violation or satisfaction of any triangle inequality. The naive infimum metric $d(A, B) := \inf \text{fee}(G : A, B \in V(G))$ is bounded above by the direct $A \leftrightarrow B$ 2-tool fee; whether triangle inequality holds for d depends on properties of the composition-graph structure, not on Theorem 1 alone. A path-summed alternative $d_{\text{path}}(A, B) := \inf \sum \text{fee}(\text{segment}_i)$ satisfies triangle inequality trivially via path concatenation, but is a different invariant.

5. **No validation against runtime outcomes.** We compute a structural fee from manifests but do not verify that high-fee compositions actually fail in production.

6. **Restriction to the exact regime.** Theorem 1 assumes $G_1, G_2 \in \mathsf{SemComp}_{\text{ex}}$. Outside this regime, the cochain complex is not bounded in degrees $[0, 1]$ and the Mayer–Vietoris argument requires an additional step.

7 Conclusion

We have proved an exact decomposition formula for the coherence fee under gluing of MCP compositions (Theorem 1), identified an operational sufficient condition under which the gluing is purely additive (Corollary 2), and given a closed-form characterization of the structural residual on length-2 paths in terms of named convention dimensions (Theorem 3). On the 13-server tier3 corpus, the closed form holds on every positive-fee pair (33/33), with the schema-blank server `exa` realizing the maximum bypass for every pair where it is positive.

The combined result is modest in scope but operationally meaningful. The decomposition makes safe stitch points checkable in closed form; Theorem 3 turns the bypass analysis from an $O(|S| \cdot |G|^{\omega})$ path search into an $O(|S|)$ lookup over named-dimension sets; the empirical analysis confirms that the structural residual—the fee that no length-2 routing can absorb—is the bulk of the pairwise fee in a real MCP corpus.

References

[1] J. Komkov. *Bulla: Witness Kernel for Agentic Compositions.* `bulla` package, 2026.

[2] J. Komkov. *A Witness Logic for Semantic Composition: An Axiomatic Characterization of Witness Rank.* Composition Doctrine paper, 2026.

[3] J. Komkov. *State of Agent Coherence, Q2 2026.* Bulla calibration report, 2026.

[4] A. Hatcher. *Algebraic Topology.* Cambridge University Press, 2002.

[5] J. Hansen and R. Ghrist. Toward a spectral theory of cellular sheaves. *J. Applied and Computational Topology*, 2019.

[6] J. Curry. *Sheaves, Cosheaves and Applications.* PhD thesis, University of Pennsylvania, 2014.

A Reproducibility

The empirical analysis in §4 is reproducible from the public Bulla package and tier3 calibration data. The script `bulla/scripts/bypass_analysis.py` regenerates the tables in §4.3 and §4.5 deterministically from the 13 manifest JSONs in `bulla/calibration/data/tier3/manifests/`.

```
$ python3.11 bulla/scripts/bypass_analysis.py
```

Outputs:

- `bypass_analysis.json`—raw distance matrix and bypass tables
- `bypass_analysis.md`—human-readable summary

Random seeds: none (all computations are deterministic).