The Cohomology of Equivocation

Detecting Split-View Lies in Federated Witness-Log Gossip by Sheaf Consistency

Paper 7 of the Harbor program — the R6 mechanism, its statistical harness, and the consistency-radius theorems

Erich Owens

The Harbor Research Program `erich.owens@gmail.com`

August 2026

Abstract

In a federation of witnesses gossiping signed log heads, an equivocator shows different histories to different neighbors. Pairwise cross-checking catches a lie exactly on the links that are checked; this paper answers, and deliberately answers *only*, the next question: when can an analyst convict inconsistency across a link that was *never* checked? The answer is a sheaf-consistency statement with a sharp topological boundary. Model the gossip graph with a cellular sheaf whose restriction maps are shared-prefix coordinate selections, and compute the least-squares *completion residual* r of the compared-plus-relayed data. Then (i) for a single equivocator, detection of a split-view lie beyond pairwise comparison occurs iff the unchecked edge lies on a cycle of the *visible* coordinate subgraph — the compared-plus-relayed links carrying that coordinate — with its endpoints' reports relayed to the analyst; on a cut edge of that subgraph $r = 0$ by algebra, and across a truly severed edge equivocation is provably dark; (ii) $r > 0$ is sound: no global history explains the visible data, and *any* offset pattern consistent with it has norm at least r, with equality achieved — r is a certified lower bound on the injected lie, in closed form $r = |s|\sqrt{1 - R_{\text{eff}}^{K_c}(e)}$ for a single-edge lie, the effective resistance being measured in that same visible subgraph (the harness's 1.2247 *is* $3\sqrt{1 - 5/6}$); (iii) the residual localizes to visible cycles through the equivocator, and computes in $\widetilde{O}(|E| \cdot L)$ via per-coordinate graph Laplacians. The scope discipline is the signature: single equivocator (coalitions on a cycle cancel to zero, measured at 6×10^{-15}), no attribution ever (signatures attribute; cohomology localizes), and vanishing r is never an all-clear. A pre-registered statistical harness returned verdict COMMIT with the cut-edge silence at float epsilon and the severed arm at 0/200. *Express lane: the one-breath sentence opens each section; the formal statement is in the box below it; a reader in a hurry needs nothing else.*

1 The problem: a lie across a link nobody checked

The scene. Six relays in a federation gossip the heads of their signed witness logs. Adjacent relays cross-check: each pair compares the log prefix they share, signs the comparison, and forwards what it heard. Relay q equivocates — it shows one head to its clockwise neighbor and a different head to its counterclockwise neighbor — and it is careful: the one link on which its two stories would collide is exactly the link that is never directly compared, because those two neighbors never reconcile. Every pairwise check in the system passes. The auditor's table of signed comparisons is green. Has the federation lost?

This is the equivocation problem in its federated form: not "did two conflicting signed statements land on one desk" (that is classical fork detection, and signatures settle it), but "can the *pattern* of everyone else's consistent reports convict an inconsistency that no single comparison witnessed?" The Harbor program's ledger scoped result R6 to precisely this question, and the scope discipline is the paper's signature: we claim detection — a certificate that no single global history explains the visible reports — on an exactly characterized set of topologies and visibility patterns, and we prove darkness everywhere else, rather than letting the word "cohomology" imply more.

Express lane. One sentence per result, formal statements in the boxes. For a single equivocator, detection beyond pairwise comparison happens iff the unchecked edge lies on a cycle *the analyst can see* — one built entirely from compared and relayed links — and its endpoints' reports are relayed to the analyst; cut edges of that visible subgraph are provably silent, severed edges provably dark (box in §4); the completion residual r is a certified lower bound on the lie any adversary must have injected, with closed form $r = |s|\sqrt{1 - R_{\mathrm{eff}}^{K_c}(e)}$ (box in §6); the residual localizes to cycles through the equivocator and costs at most L scalar Laplacian solves (same box); and the statistical harness that validated all of it was pre-registered and returned COMMIT (box in §5).

The idea in one breath: if everyone is honest there is one true history behind every report, and then the disagreements the relays report must cancel around every loop of the gossip graph — so a loop that fails to cancel convicts an equivocator across a link nobody ever checked, and where the analyst can close no loop, nothing can be seen at all.

The analogy, mapped by relations. Wristwatches around a table (Figure 1): each adjacent pair reports the *offset* between their watches, never the absolute times. If everyone is consistent, the offsets around the full table telescope to zero — there exists a global assignment of times inducing them. A nonzero loop-sum convicts someone of showing different times to different neighbors, *without knowing any absolute time and without the missing pairwise check*: the rest of the loop substitutes for it. The relation that transfers, not the surface: loop-sum = cocycle condition; missing comparison on a loop = unchecked edge on a cycle; and — the misread to preempt — a watch at the end of a *line* of tables can lie freely, because no loop passes through it: the telescoping argument needs the loop, and so does the theorem. Note also which loops count: a loop is only usable if the analyst hears about *every* pair on it, so a table where one adjacent pair never reports is, for this argument, a line and not a loop.

Those plain phrases have technical names, and §2 supplies them: "one true history" is a *global section*, "the disagreements that history would produce" is a *coboundary*, "cancels around every loop" is the *cocycle condition*, and "how far the reports sit from any one true history" is the *completion residual* r. Nothing in the paper needs more than those four.

2 The vocabulary, defined

This paper borrows its machinery from applied algebraic topology, and that field's vocabulary is forbidding out of proportion to what it actually asks of the reader. Everything below is linear algebra on a graph. A reader who is comfortable with incidence matrices and least squares already has the prerequisites; what follows is a dictionary, in the order the terms are needed, each defined the sentence after it is named and then exercised immediately — on the six relays of §1, so that every abstract term lands on the same concrete scene rather than on a fresh one.

A *cellular sheaf* on a graph is a bookkeeping device for local data that is supposed to agree along edges: it assigns a vector space to each vertex, a vector space to each edge, and a linear map from each vertex into each edge it touches. Ours assigns to witness v the space $\mathbb{R}^{d_v}$ of its possible log states, to edge e the space of the log prefix its two endpoints both report, and to each (vertex, edge) pair the map that selects the shared coordinates. The word "sheaf" is doing no more work here than "data structure with a consistency condition" — Robinson's slogan that sheaves are the canonical data structure for sensor integration [8] is meant literally.

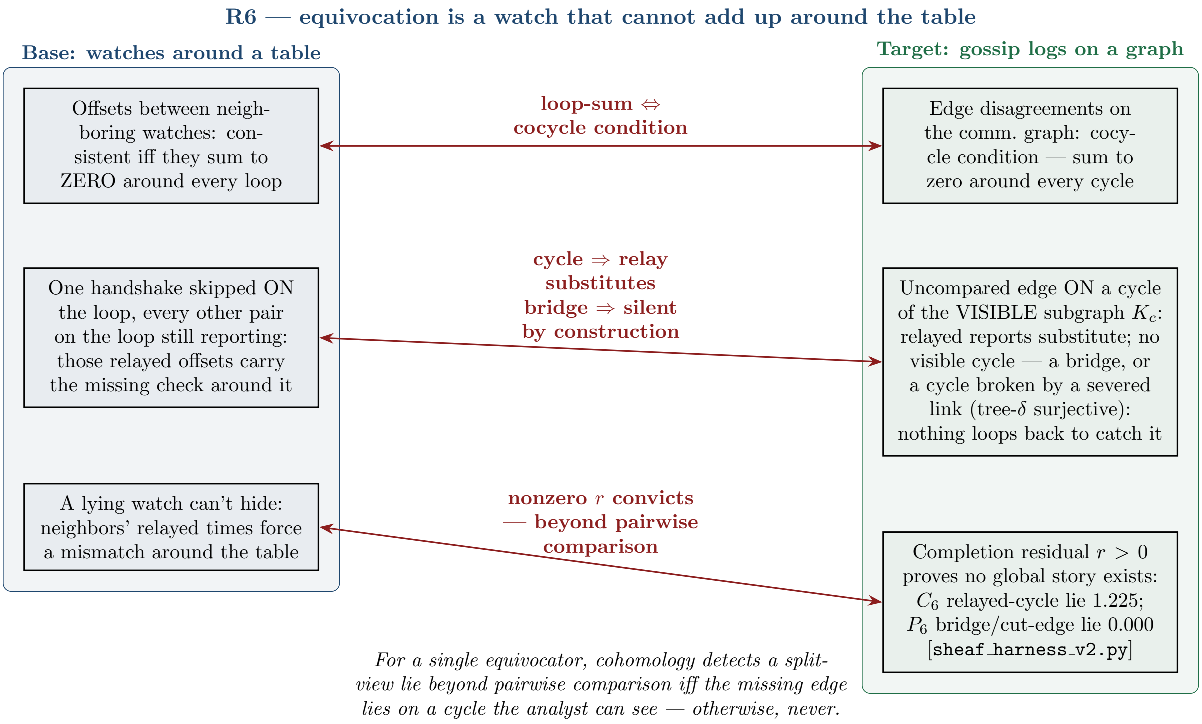

Figure 1: Relation-map for the mechanism. Wristwatch offsets around a table map to signed-log disagreements around a gossip cycle by *relations*: loop-sum ↔ cocycle condition; a loop through the missing link ↔ a cycle through the uncompared edge — a loop counts only when every pair on it reports, which is the visible subgraph K_c of §3. The map predicts, correctly, where the mechanism dies: no visible loop, no telescoping, no detection.

The vertex space $\mathbb{R}^{d_v}$ is the *stalk* at v — the set of values that vertex could locally hold, with no reference to what its neighbours hold. Relay 3 of the scene has a stalk: the five log-head entries it could report, one real number each, so $d_3 = 5$ (the harness of §5 runs $\mathbb{R}^5$ stalks throughout). The maps out of it are the *restriction maps*: restriction, because they cut a vertex's full state down to the part an edge can see. On the link between relays 3 and 4, which reconcile only their first two entries, the restriction map is the two-row matrix that keeps entries 1 and 2 and drops the rest. In our setting a restriction map is always a coordinate selection, nothing more exotic — "report the first k entries of your log," as a matrix.

The *coboundary* δ is the operator that takes an assignment of a value to every vertex and returns, for every edge, the disagreement its two endpoints report there. For scalar stalks on a graph it *is* the signed incidence matrix, and δx is the familiar vector of differences $x_u - x_v$ along each edge. A disagreement pattern g that equals δx for some x is called a *coboundary*: it is exactly a pattern that some consistent global assignment would have produced. This is the paper's critical definition. Honesty means the observed disagreements are a coboundary; equivocation means they are not. Concretely, on the six-relay ring the observable is six numbers, one per link — that column of six is the object the entire apparatus operates on, and §4 writes a four-relay instance of it out in full and checks it by hand.

A *global section* is such an x — one assignment of a log state to every witness that simultaneously explains every edge's reported disagreement. For the six-relay scene it is six numbers, one per relay, that the auditor would have to be able to write down for the green table of comparisons to be honest. "No global section exists" is the formal content of "no single history explains what everyone reported," and it is the only thing this paper ever convicts anyone of.

The *cycle space* of a graph is the set of edge-weightings that circulate — flow in equals flow out at every vertex — and its dimension is the number of independent loops, written β_1. It matters because of one elementary fact: a coboundary summed around any cycle telescopes to zero. Walk a loop adding up $x_u - x_v$ and every term cancels. So the cycle space is precisely the set of directions in which an honest world *cannot* have disagreement, which is why a nonzero component there is evidence and a nonzero component anywhere else is not. The six-relay ring has $\beta_1 = 1$: one loop, so exactly one direction of evidence. The same six relays in a chain have $\beta_1 = 0$, no direction of evidence at all — and, crucially for what follows, so does the ring once a link goes dark, because the loop the argument needs must be built from links the analyst can actually see.

Cokernel and H^1 name that same space from the algebraic side: $\operatorname{coker}(\delta)$ is what edge-data space has left over after you remove everything δ can produce, i.e. everything an honest world could have generated. H^1 is the standard name for that quotient. The reader who prefers to think entirely in terms of "project the observed disagreements onto the cycle space and see what survives" will never be wrong in this paper; the cohomological name buys the connection to Curry [6] and Hansen–Ghrist [7] and costs nothing else.

The *completion residual* r of §3 is then a least-squares problem in the ordinary sense: over all global sections x, and over all values the analyst cannot see, minimise the Euclidean distance between what was reported and what δx predicts. It is zero exactly when some honest world fits the evidence.

Effective resistance $R_{\text{eff}}(e)$ is the resistance measured across edge e when the graph is treated as a circuit of unit resistors — and, throughout this paper, measured on the *visible* subgraph of §3, since a detour through a link the analyst cannot see is not a detour the analyst can take. It appears in §6 for a reason that is physical rather than formal: a lie on an edge is hard to detect exactly when the rest of the graph could have produced the same disagreement pattern by itself — which is the condition that current finds *no* cheap alternative path around the edge, i.e. *high* effective resistance. On a bridge nothing routes around e at all, $R_{\text{eff}} = 1$, and the lie hides completely; around a long cycle the only detour is long, R_{eff} approaches 1, and conviction weakens like $1/\sqrt{n}$. At the other end of the scale a well-connected edge has many cheap alternative paths, R_{eff} near 0, and the lie is convicted at nearly its full size: the certified fraction is $\sqrt{1 - R_{\text{eff}}}$, largest exactly where the alternative paths are easiest.

Finally the *graph Laplacian* $\delta^\top \delta$ is the matrix whose linear systems compute all of the above, and whose near-linear-time solvers [11] are what make the detector cheap enough to run every gossip round.

One caution about a word this paper uses in two senses. A *cycle* in the graph-theoretic sense above is a loop of edges. It is unrelated to a cycle of gossip *rounds*. Where the distinction matters we write "graph cycle."

3 The model and the observability contract

The sheaf. Fix a finite gossip graph $G = (V, E)$. Witness v carries a local log state $x_v \in \mathbb{R}^{d_v}$ (its stalk). Each edge $e = \{u, v\}$ carries the *shared prefix* the two endpoints both report — a subset S_e of coordinates, $|S_e| \leq L$ — and the restriction maps $\rho_{v,e} : \mathbb{R}^{d_v} \to \mathbb{R}^{|S_e|}$ are coordinate-subset selections. The observable on edge e is the disagreement $g_e = \rho_{u,e}(u\text{'s report}) - \rho_{v,e}(v\text{'s report})$. Honesty means all reports restrict one global section x: then $g = \delta x$ is a coboundary of the sheaf coboundary δ, and every cycle-sum of g vanishes. Equivocation is an offset pattern ε added to q's per-neighbor reports; it is invisible pairwise on any edge not directly compared, and the question is what survives in H^1 — the obstruction space $\operatorname{coker}(\delta)$, which for prefix restrictions decomposes per shared coordinate c into the cycle spaces of the *coordinate subgraphs*

$$G_c \;=\; \big(V,\ \{e \in E :\ c \in S_e\}\big)$$

[verified framework, [6, 7]]. G_c is what the *federation* has: every link that carries coordinate c. What the *analyst* has is smaller, and the difference carries Theorem 1 — see K_c below.

Three-tier visibility — the observability contract. Not every edge is equally visible to the analyst, and the paper's claims are stated *per tier* (Figure 2):

- **compared** — the analyst holds the direct cross-check of the edge: pairwise detection, no cohomology needed;
- **relayed** — both endpoints' reports about the edge reach the analyst by gossip, but the edge itself was never checked by its endpoints: this is the tier where cohomology earns its keep, and only on cycles of K_c;
- **severed** — no report about the edge crosses to the analyst at all: the corresponding data block is a free variable of the completion, and equivocation confined there is *provably dark.*

The visible coordinate subgraph K_c. Write K for the compared-plus-relayed edge set — every link the analyst holds a report about — and

$$K_c \;=\; \big(V,\ \{e \in K :\ c \in S_e\}\big) \;\subseteq\; G_c$$

for its coordinate-c part: the coordinate subgraph with the severed links deleted. Every detection statement in this paper is a statement about cycles of K_c, never of G_c, and the reason is one line of algebra made explicit in the detector below — a severed block is a free variable, so minimising over it simply deletes its row, and a loop that runs through a severed link is not a loop the analyst can close. The federation's topology sets G_c; the observability contract sets K_c; only the second one convicts anybody.

The detector. The *completion residual* is the least-squares defect of the best global explanation of everything visible:

$$r \;=\; \min_{x,\, g_{\text{sev}}} \big\| \, g_K - (\delta x)\big|_K \, \big\|_2,$$

the minimum over global sections x and free severed blocks g_{sev}, where K is the compared-plus-relayed edge set. The severed blocks are wholly unconstrained, so minimising over them is exactly *dropping their rows*: g_{sev} never enters the objective, and $r = \operatorname{dist}\big(g_K,\ \operatorname{im}\delta_K\big)$ for δ_K the coboundary restricted to visible edges. Honest executions give $r = 0$ exactly, for every visibility pattern; $r > 0$ certifies that *no* global history explains the visible data. Equivalently $r = \|\Pi_K\, g_K\|$, the projection of the visible disagreement data onto the per-coordinate cycle spaces *of the visible subgraphs*:

$$r^2 = \sum_c \big\|\mathrm{Proj}_{\mathrm{cycle}(K_c)}\, g^c\big\|^2, \qquad g^c \in \mathbb{R}^{K_c} \text{ the coordinate-}c\text{ part of the visible data}$$

[internal, `sheaf_harness_v2.py`]. The derivation is three lines and is the whole bridge from the minimization to Theorem 1, so it is worth stating: prefix restrictions are coordinate selections, so δ_K block-diagonalises per coordinate into the signed incidence operator B_c of K_c — the severed rows having already dropped out; edge space splits orthogonally as $\mathbb{R}^{K_c} = \operatorname{im}(B_c) \oplus \ker(B_c^\top)$, cut space plus cycle space; and the residual is what is left after removing everything an honest world could have produced, i.e. the component in $\mathrm{cycle}(K_c)$. Writing G_c here instead would not merely be loose, it would fabricate detections: on C_6 with a lie of 3 on a relayed edge and one other cycle edge severed, the residual is 1.1×10^{-14} while the G_c form reports 1.2247 — and reports a different value again for a different honest world, so it is not even a well-defined statistic [internal, `sheaf_consistency_radius.py`].

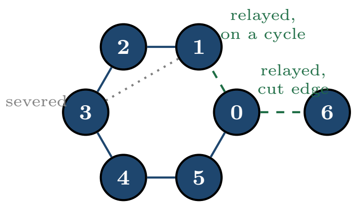

one gossip graph, three visibility tiers

K_c = compared ∪ relayed

what each tier lets the completion residual prove

—— **compared**	analyst cross-checks both endpoint reports on the edge; disagreement observed directly (pairwise detection).
- - - **relayed**	endpoints' reports reach the analyst by gossip; the edge itself is never checked. On a cycle of the visible subgraph K_c: $r > 0$, detectable (200/200). On a cut edge of K_c: $r = 0$ by algebra, provably silent (max 1.5×10^{-13}).
······ **severed**	no report crosses; the block is a free variable of the completion, so its row drops out of K_c altogether; equivocation there is provably dark (0/200) — and so is any cycle that ran through it.

Figure 2: The three-tier observability contract — the tiers are properties of the *analyst's evidence*, not of the federation, and all three coexist on one topology. Left: one gossip graph wearing all three tiers — harborblue solid compared, seagreen dashed relayed, gray dotted severed; the compared-plus-relayed links are the visible subgraph K_c that every theorem in this paper quantifies over. (The severed link 1–3 is drawn as a chord, so this graph carries two independent loops; the running example of §1 is the plain ring, $\beta_1 = 1$.) Right: what each tier lets the completion residual prove. The relayed tier splits by topology: on a cycle of K_c, $r > 0$ detects (200/200); on a cut edge of K_c, $r = 0$ by algebra (max residual 1.5×10^{-13} over 400 trials); severed blocks are free variables and provably dark (0/200). Figure 3 uses this same legend [internal, `sheaf_harness_v2.py`].

Two defects that taught the model. The first harness (v1) got both halves of this contract wrong, and the rebuild kept the autopsies as named defects with reintroduction mutants. **D1:** v1 drew *random orthonormal projections* as restriction maps; for those, $\operatorname{coker}(\delta)$ generically collapses to 0 and the signal dies — the detector only exists because shared-prefix restrictions make the cokernel isomorphic to $\bigoplus_c H^1(G_c)$, cycle space tensor shared coordinates (the analyst's computable version of which is $\bigoplus_c H^1(K_c)$ — same statement, visible edges only). The rebuilt harness refuses D1's reintroduction by a structural self-check (it observed coker $= 0$ against $\beta_1 = 7$ and aborted). **D2:** v1 scored "detections" on cut and severed edges — by silently reading hidden-edge data the analyst was not entitled to; the rebuilt harness flags D2's reintroduction because the mutant's "partition" residuals come out bit-identical to full visibility [internal, `sheaf_harness_v2.py`]. Both defects are the same lesson: the theorem is about *what the analyst can see*, and an implementation that muddles the observability contract will manufacture either silence or false sight.

4 The mechanism: detection iff the missing edge closes a visible cycle

Cohomology detects a single equivocator's split-view lie beyond pairwise comparison exactly when the uncompared edge lies on a cycle every link of which the analyst can see — the cocycle condition then substitutes for the missing comparison; where no such cycle closes there is nothing to substitute, and the residual is zero by algebra — formal statement in the box below.

Theorem 1 (mechanism). In gossip of signed logs with prefix restrictions and the three-tier visibility model, and for a *single equivocator* (the honest scope fixed in §6 and §8: coalitions on a common cycle can cancel r to zero), the completion residual detects a *split-view* lie beyond pairwise comparison iff the unchecked edge lies on a cycle of its *visible* coordinate subgraph K_c — the compared-plus-relayed links carrying that coordinate — *and* its endpoints' reports are relayed to the analyst: then $r > 0$ proves no global explanation exists. On a cut edge of K_c, $r = 0$ identically — the coboundary map of a tree (or forest) is surjective onto the visible data, so the free completion absorbs any lie; this one clause covers both a bridge of the federation and a graph cycle of G_c whose other links happen to be severed. Across a severed edge, equivocation is provably dark: the severed block is unconstrained, and a consistent counterfactual world always exists.

"Split-view" is exact, not a hedge. When q sits on a cycle of K_c, the offsets invisible to the residual — $\ker(\Pi_K A_q)$, in the notation of §6 — are precisely the *uniform* ones: the same shift shown to every neighbor, which is a consistent alternative state and not a split view at all (§8). Measured on the star of q with all edges visible, every kernel direction lies in the uniform span to 5×10^{-16}, on both the two-path and C_6 topologies [internal, `sheaf_consistency_radius.py`].

Verification (minimal pair, plus the visibility witness). C_6 with scalar stalks, lie of size 3 on the one uncompared cycle edge: every pairwise check passes, harmonic signal $1.225 \neq 0$. The same lie on a bridge of the path P_6: signal 0.000 [internal, `sheaf_mechanism_proof.py`]. And the witness that separates K_c from G_c: the same C_6 and the same lie of 3 on the same relayed edge $(2,3)$, but with one *other* cycle edge severed — $(0,5)$ or $(4,5)$ — gives $r = 1.1 \times 10^{-14}$ and 1.9×10^{-15}. The lie edge is still on a cycle of G_c and its endpoints are still relayed; it is no longer on a cycle of K_c, and the residual is float zero [internal, `sheaf_consistency_radius.py`]. The structural identity behind the mechanism — local consistency with no global section — is Abramsky–Brandenburger contextuality [verified framework, [5]]; the equivocator manufactures for its neighbors exactly a locally-consistent, globally-unglueable family of views.

The cochain, by hand. Before any general theory, the smallest complete instance, in exact quarters. Four relays $0, 1, 2, 3$ in a ring, scalar stalks, one shared coordinate. The links $0-1$, $1-2$ and $2-3$ are compared and agree; the link $3-0$ is never compared, but both its endpoints' reports reach the analyst, so all four links are visible and K_c is the whole ring. Relay 0 equivocates: the head it shows across $3-0$ differs by $s = 3$ from the one it shows relay 1. The analyst's cochain — the vector of edge disagreements, one number per link, the object the whole apparatus operates on and which the reader has not yet seen — is

$$g \;=\; (g_{01},\, g_{12},\, g_{23},\, g_{30}) \;=\; (0,\ 0,\ 0,\ 3).$$

Ask the only question this paper ever asks: is there one history $x = (x_0, x_1, x_2, x_3)$, one number per relay, whose differences reproduce g? Around the loop a coboundary must sum to zero and g sums to 3, so there is none. Least squares finds the closest one: $x = (0,\ 0.75,\ 1.5,\ 2.25)$, whose coboundary is $\delta x = (-0.75,\ -0.75,\ -0.75,\ +2.25)$ — and which therefore misses by exactly 0.75 on *every* link. That is the mechanism in one sentence: the best honest world spreads the lie evenly around the loop and still misses everywhere. Hence

$$r = \sqrt{4 \times 0.75^2} = 1.5 = \frac{|s|}{\sqrt{n}} = \frac{3}{\sqrt{4}} = |s|\sqrt{1 - R_{\text{eff}}^{K_c}(e)} = 3\sqrt{1 - \tfrac{3}{4}},$$

three routes to the same number, none needing a calculator [verified]. Now break the loop, twice, and watch the same arithmetic go silent. *As a chain:* the same four relays as $0-1-2-3$, the same lie of 3 on the never-compared end link 2–3 — the assignment $x = (0, 0, 0, -3)$ reproduces $g = (0, 0, 3)$ *exactly*, so $r = 0$, not small but zero, and the reader has just built the adversary's alibi by hand. *As a ring with a dark link:* keep the ring and the same lie on 3–0, but let 1–2 be *severed*. Now $g_K = (g_{01}, g_{23}, g_{30}) = (0, 0, 3)$ on three visible links, and $x = (0, 0, 3, 3)$ reproduces it exactly: $r = 0$ again. The loop is still there in G_c; it is gone from K_c, and the residual can only see K_c. This is Theorem 1's visibility clause, checkable in four numbers.

Numbers by hand. On C_6, project the lie onto the cycle: the harmonic component of a single-edge offset s on an n-cycle is $s/\sqrt{n}$, so the signal is $3/\sqrt{6} = 1.2247$ — the harness's 1.225 to printed precision, and, in the closed form of §6, $3\sqrt{1 - 5/6}$ [verified]. On P_6 there is no cycle: the cycle space is zero, the projection is zero, the residual is zero — not small, zero (Figure 3 puts the relayed-cycle, cut, and severed cases side by side on the same six-node graph). *Now you try:* C_8, lie of size 2 on the relayed edge — the residual is $2/\sqrt{8} = 0.7071$; the same lie on any bridge, 0; and the same lie on C_8 with any one of the other seven links severed, also 0 [verified].

The misread to preempt. "Cohomology sees hidden data" is exactly what the theorem does *not* say. The residual reads only relayed evidence, and squeezes it through the one constraint honesty imposes — cycle-sums vanish. Where no *visible* cycle passes (a bridge, or a cycle broken by a severed link) or no evidence arrives at all (severed edge), there is no constraint to violate and nothing is seen; the harness's cut-edge maximum of 1.5×10^{-13} over 400 trials is float epsilon, not a small detection [internal, `sheaf_harness_v2.py`]. A second misread: the detection statistic is the *data* residual, never $\dim H^1$ — a loopy topology has $\beta_1 > 0$ for purely topological reasons on honest data too, so any counting claim must net β_1 out; the harness prints the structural cokernel separately (e.g. the two-path split-view topology: $\beta_1 = 1$, coker $= 2$) and scores only the residual, which is identically zero for honest worlds.

Honest boundary. Detection requires the reports, not the check: severed links stay dark, cut edges silent, and the v1 "detections" there were D2 artifacts. Vanishing signal over $\mathbb{R}$ is not an all-clear — real coefficients abelianize the obstruction, and contextual models exist whose cohomological invariant vanishes (the Carù gap [9]); §8 prices this. And the mechanism convicts an inconsistency; it does not name its author (§6, §8).

5 The harness: pre-registered gates, verdict COMMIT

Before running, the harness committed to its gates: the completion residual must detect split-view equivocation on relayed cycles where pairwise comparison is blind, and must stay silent on honest worlds, cut edges, and severed arms — it passed, the rival CUT condition was refuted, and both v1 defects are guarded by mutants — verification in the box below.

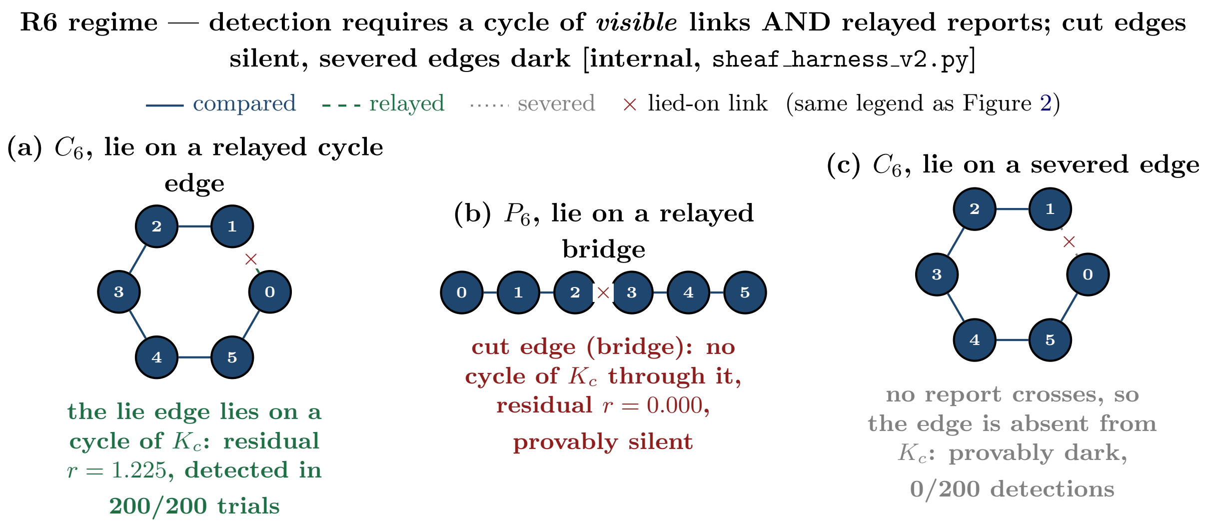

Figure 3: Regime diagram for the mechanism, in one legend with Figure 2 (harborblue solid compared, seagreen dashed relayed, gray dotted severed; the shipred × marks the lied-on link, independently of its tier). (a) The lie sits on a relayed link that still lies on a cycle of the visible subgraph K_c — residual $r = 1.225$, detected in 200/200 trials, cohomology-only. (b) The same relayed tier but a bridge: no cycle of K_c runs through it, so $r = 0.000$ and the completion absorbs any lie — provably silent. (c) The link is severed, absent from K_c entirely, and equivocation there is provably dark (0/200 detections). Detection needs both clauses — a cycle the analyst can see *and* relayed endpoint reports; deleting either one, as (b) and (c) do, deletes the detection [internal, `sheaf_harness_v2.py`].

The design. Two hundred trials per arm at seed 20260816, stalks $\mathbb{R}^5$, shared-prefix sizes 2–4. The detector is the completion residual of §3; the score is *cohomology-only* detections — trials where the residual fires and every pairwise comparison passes. The pre-registered Stage-1 gates from the rescued prototyping plan: (i) the split-view arm must be detected cohomology-only; (ii) the control arms must be silent.

Harness results (all arms, seed 20260816). Partition-on-a-cycle split view (`two_path`): 200/200 cohomology-only — pairwise blind in every trial. Severed-cycle arm: 0/200 — the honest boundary, measured. Cut edge (`single_bridge`): max residual 1.5×10^{-13} over 400 trials — silence at float epsilon. Full visibility: 0 cohomology-only detections and 200/200 redundant (pairwise catches everything first, as it must). Expander with 3 severed and 3 relayed edges: partition-straddling equivocator 113/200 cohomology-only, 87 dark — both clauses of the contract live in one topology, and the dark trials split cleanly along them. Classifying all 87: in 79 the equivocator has no relayed incident link at all, so every lie lands on a severed edge and Theorem 1's relayed-reports clause fails — darkness predicted; in the remaining 8 the lie edge *is* relayed and does lie on a cycle of G_c, but not on a cycle of K_c, because the rest of that cycle is severed. Those 8 are the arm's only trials that distinguish the two readings of the cycle clause, and they are dark, which is what the K_c form of Theorem 1 predicts and the G_c form does not [internal, `sheaf_harness_v2.py`]. Localization: the max-residual edge lies on a known-graph cycle 200/200 (upgraded to "through the equivocator" as a theorem in §6). Pre-registered gates (i) and (ii) **PASS**; verdict **COMMIT**. The rival CUT condition — "residual detections are just pairwise detections in disguise" — is refuted: 313 of 1101 residual detections across arms were pairwise-blind. Mutation suite: D1 reintroduction refused by the structural self-check (coker $= 0$ against $\beta_1 = 7$); D2 reintroduction flagged (mutant "partition" residuals bit-identical to full visibility); no-equivocator control 0/200 [internal, `sheaf_harness_v2.py`].

What it buys. The W7 mechanism proof was a two-topology existence demonstration; the harness upgrades it to a validated statistical detector with its false-positive behavior measured (honest and control arms identically zero), its blindness measured (severed arm 0/200, and the expander's 87 dark trials accounted for one by one — 79 by the relayed-reports clause, 8 by the visible-cycle clause), and its guards regression-tested by reintroducing the exact defects that fooled v1. The three-tier contract stopped being prose and became the experiment's arm structure.

Honest boundary. The harness validates the detector at the sharpened scope only: synthetic gossip, prefix restrictions, single equivocator per trial, $\mathbb{R}^5$ stalks. The 113/200 expander number is a property of that topology's severed/relayed mix, not a universal rate. And a statistical PASS is not yet a theorem — which is exactly why §6 exists.

6 The consistency radius: what $r > 0$ certifies, where it points, what it costs

The residual r is the exact minimum lie any adversary must have injected to produce the visible data, it localizes to visible *cycles through the equivocator, and it computes as at most L scalar graph-Laplacian solves — with the single-equivocator condition proved to be the honest scope, because coalition lies on a cycle can cancel r to zero — formal statements in the box below.*

Intuition. Robinson's consistency radius asks how far local data sits from the nearest global section [8]; our r is its least-squares form under the completion over severed blocks. The new content is that r is not merely a distance — it is a *certificate about the adversary*: every offset pattern that could have produced the visible data is at least this large, and one of exactly this size exists. The topology enters through electrical network theory: a lie on edge e hides in proportion to how well the rest of the *visible* graph can explain it away, which is the effective resistance $R_{\text{eff}}^{K_c}(e)$ measured in the visible coordinate subgraph K_c of §3 — on a bridge $R_{\text{eff}} = 1$ and the lie hides completely; on a long cycle $R_{\text{eff}} \to 1$ and the certified fraction $\sqrt{1 - R_{\text{eff}}}$ decays like $1/\sqrt{n}$ (Figure 4). Severing a link raises $R_{\text{eff}}^{K_c}$ of everything that was routed around through it, which is the same fact as Theorem 1's visibility clause, written as a price.

Theorem CR-1 (soundness and achievability). For the completion-residual detector with prefix restrictions and three-tier visibility: $r > 0$ iff no global section explains the compared-plus-relayed data, and honest executions give $r = 0$ exactly, for every visibility pattern. Every offset pattern ε consistent with the visible data satisfies $\|\varepsilon_K\| \geq r$, with equality at the completion — r is the exact minimum lie. For a single equivocator q with offset vector o and visible-lie operator A_q: $\sigma_{\min}^{+}(\Pi_K A_q)\,\|o_\perp\| \leq r \leq \|o_\perp\|$, where $o_\perp$ is the component of o outside $\ker A_q$. For a single-edge lie of size s in coordinate c,

$$r \;=\; |s|\sqrt{1 - R_{\text{eff}}^{K_c}(e)} \qquad \text{— on } C_n:\; |s|/\sqrt{n};\ \text{the harness's 1.2247 } \textit{is } 3\sqrt{1 - 5/6},$$

where the effective resistance is taken in the *visible* coordinate subgraph K_c, not in G_c: on the theta graph (C_6 plus a chord) with the chord severed and a lie of 3 on a relayed cycle edge, the measured residual is $1.2247 = 3\sqrt{1 - 5/6}$, the K_c value; the G_c reading would predict 1.7928 [internal, `sheaf_consistency_radius.py`].

Theorem CR-2 (localization). Noise-free, single equivocator: the support of the edge residual lies on cycles of the visible coordinate subgraphs K_c through q — by the electrical-flow support argument, the harness's 200/200 localization is now a theorem. With bounded noise η, localization is guaranteed while $2\|\eta_K\| < \max_f \|(\Pi_K \varepsilon_K)_f\|$.

Theorem CR-3 (complexity). Because δ_K splits per shared coordinate into graph incidence operators, r is one sparse least-squares problem that decomposes into at most

L scalar graph-Laplacian systems: $\widetilde{O}(|E| \cdot L\log(1/\epsilon))$ with SDD solvers [11]. (Bach's #P-hardness of sheaf cohomology [12] concerns coherent sheaves on projective space — the wrong category; finite cellular sheaves over $\mathbb{R}$ on graphs are linear algebra.)

Verification. Soundness 0/600 violations; achievability 600/600 (a consistent pattern of norm exactly r constructed every time). The R_{eff} closed form is exact on C_4–C_{12} and 200 random graphs; the $\sigma^+_{\min}$ lower bound is tight under the adversarial offset ($r/\|o_\perp\| = \sigma^+_{\min} = 0.4082$); a uniform kernel lie of norm 5 gives $r = 2\times10^{-14}$ — a consistent alternative state rather than a split view, invisible in principle and correctly not flagged; and on a cycle of K_c the uniform offsets are the *only* such directions, every kernel direction of $\Pi_K A_q$ sitting in the uniform span to 5×10^{-16}. Localization 200/200 (two-cycles topology) plus 128/128 (severed expanders); noise boundary measured: $\geq 95\%$ localization up to $\sigma_\eta/\sigma_{\text{eq}} \approx 0.1$, 76.5% at 0.2; sufficient-condition violations 0. Timing slope versus $|E|$ (wall-clock, machine-dependent — only the qualitative split reproduces across runs): conjugate gradient stays well under the 1.6 gate (≈ 0.6–0.7 across repeated runs) while dense solve is consistently super-linear and several times steeper ($\gtrsim 3.0$); it is the CR-3 decomposition, not the two digits, that the gate certifies. Two coordinated liars on a common cycle: each alone certified at $r \geq |s|/\sqrt{8}$, the coalition cancels to $r = 6\times10^{-15}$. The D2-unmasked mutant fires on a severed lie at 0.839 — exactly the full-visibility value, caught [internal, `sheaf_consistency_radius.py`].

Numbers by hand. C_6 with every link visible, so $K_c = G_c = C_6$: R_{eff} of an edge is $\frac{n-1}{n} = \frac{5}{6}$ (the rest of the cycle is a 5-edge series path in parallel with the edge), so a lie of 3 certifies $r = 3\sqrt{1/6} = 1.2247$ [verified] — the number the mechanism script printed a week before the closed form existed. The $\sigma^+_{\min}$ tightness value 0.4082 is $1/\sqrt{6}$: the same geometry seen from the operator side [verified]. *Now you try:* a federation ring of 12 relays with one never-compared link and a head-offset lie of 0.5 — the certified minimum lie is $0.5/\sqrt{12} = 0.144$; double the ring and the certificate shrinks by $\sqrt{2}$: long cycles dilute conviction, which is an argument for short reconciliation loops in federation design.

What it buys. The harness's statistic is upgraded to a certificate with an interpretation a federation operator can act on: r is the smallest lie consistent with what you saw, the hot edges lie on cycles through the culprit's neighborhood, and the whole computation is a handful of Laplacian solves — cheap enough to run on every gossip round. The closed form also converts topology into policy: effective resistance is the price list for where a federation should force direct reconciliation (bridges: mandatory — the residual cannot help there) and where relayed gossip suffices (short cycles: cheap certified detection).

7 Related work: imported, adjacent, and new

Imported. The structural identity — local consistency without a global section — is Abramsky–Brandenburger contextuality [5]; the equivocator's split view is an empirical model that is locally consistent and globally unglueable, and we use the sheaf language exactly as that program fixed it. Cellular sheaves on graphs and their coboundary calculus are Curry [6]; the sheaf Laplacian and its harmonic/spectral reading are Hansen–Ghrist [7]. The consistency radius — distance from local data to the nearest global section, with its sup-norm corollaries — is Robinson [8]; our r is its least-squares completion form. The nearly-linear Laplacian solvers behind CR-3 are the Spielman–Teng line [11].

Adjacent, and honestly positioned. Herlihy and Shavit [1] and, independently, Saks and Zaharoglou [2] founded the topological line in distributed computing — co-recipients of the 2004 Gödel Prize — by characterizing which tasks have wait-free protocols via a structure-preserving map between the protocol's and the task's simplicial complexes of process views; Herlihy, Kozlov, and Rajsbaum [3] is the standard text. We inherit their conviction that topol-

CR-1 closed form: the lie the residual certifies, by topology
[verified; `sheaf_consistency_radius.py`]

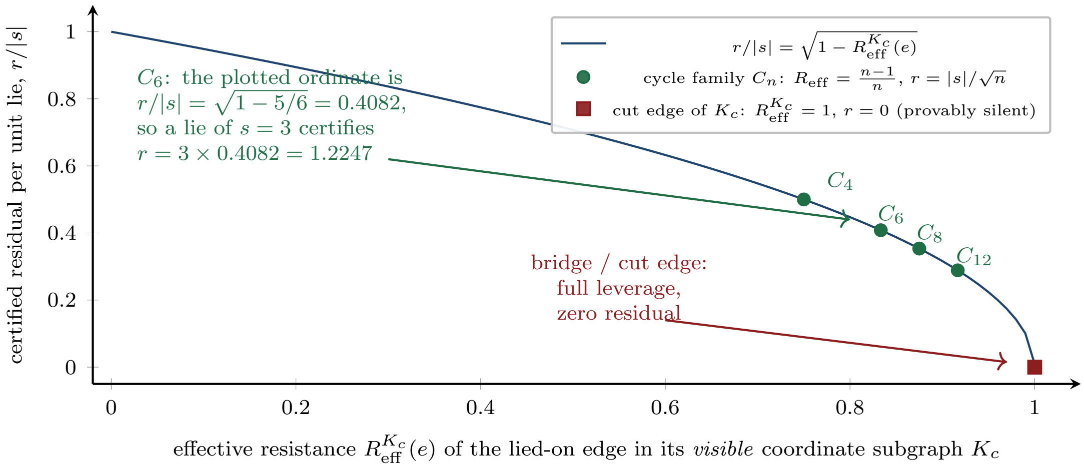

Figure 4: Conviction strength is set by topology alone. The CR-1 closed form plots certified residual *per unit lie*, $r/|s| = \sqrt{1 - R_{\text{eff}}^{K_c}(e)}$, against the effective resistance of the lied-on edge in its *visible* coordinate subgraph K_c — compared-plus-relayed links only, so severing a link on the edge's cycle moves the point right, toward silence. The cycle family C_n sits at $R_{\text{eff}}^{K_c} = \frac{n-1}{n}$; the marked C_6 point is the plotted ordinate 0.4082, which an absolute lie of $s = 3$ scales to the harness's $r = 1.2247$. The cut-edge corner $R_{\text{eff}}^{K_c} = 1$ is provable silence [verified; closed form exact on C_4–C_{12} and 200 random graphs, `sheaf_consistency_radius.py`].

ogy is the right language for what a distributed system can and cannot know, but we ask a different question of a different object: not which tasks are solvable before execution, but what an analyst can certify *after* one execution from partial, already-relayed evidence — cellular sheaves with real coefficients on the gossip graph itself, not simplicial complexes of possible process states. PeerReview [4] is the systems-level ancestor of the exact scenario: witnesses cross-check signed logs, and a node that forks its log is caught because every authenticator it issues is eventually forwarded to a witness set that can compare them directly. But that guarantee rests on PeerReview's witness-set assumption — that some correct node eventually obtains both sides of any comparison — which is precisely the assumption our three-tier contract drops for the relayed tier: we characterize detection when no such direct comparison point exists and only cycle-relayed reports remain. Carù [9] shows the Čech-cohomological obstruction with abelian coefficients can vanish on contextual models — imported here as the reason vanishing r over $\mathbb{R}$ is never an all-clear. Sheng et al.'s BFT forensics [10] works the attribution side: with signed protocol messages, equivocation yields a transferable cryptographic proof *naming* the culprit, and their impossibility results bound when. We sit deliberately on the other side of that line: cohomology reads unsigned *structure* — it can convict "someone on these cycles lied" from evidence no signature covers (relayed, never-compared links), and can never name the liar (§8); signatures attribute, cohomology localizes, and a deployed federation wants both. Bach [12] is cited only to preempt a category error: #P-hardness of coherent-sheaf cohomology does not touch finite cellular sheaves on graphs; the claim rests on a Semantic Scholar TLDR of the primary source rather than the source itself, which the primary text (institutional access blocked, see `deep-dives/BLOCKED-SOURCES-REQUESTS.md`) would close outright.

New, honestly. No component of the mathematics is exotic: incidence matrices, cycle spaces, effective resistance, least squares. The contribution is the assembled result at its exact scope: (1) the mechanism boundary — for a single equivocator, detection beyond pairwise iff the missing edge closes a cycle the analyst can see, one built entirely from compared and relayed links, with cut-edge silence and severed darkness proved rather than observed; (2) the three-tier observability contract as the object the theorems quantify over, extracted from two named defects that each faked a different clause of it; (3) the certificate reading of the completion residual — r as the exact minimum injected lie, with the $R_{\mathrm{eff}}^{K_c}$ closed form tying conviction strength to the federation topology *the analyst can see*; and (4) the coalition-cancellation construction that fixes the honest scope at single-equivocator, turning what would be a reviewer's counterexample into the paper's own boundary statement.

8 The boundary is critical

This paper's claims are narrow by construction, and the narrowness is the result. Each clause below is a theorem-grade limitation, stated at the same prominence as the theorems, because a federation that deploys this detector believing any one of them false will be harmed by the belief.

Honest boundary — the four clauses.

- *Single equivocator is the scope, and provably so.* Detection and localization are single-equivocator statements. A coalition whose combined lie is a coboundary is a consistent counterfactual world: two coordinated liars on a common cycle, each alone certified at $r \geq |s|/\sqrt{8}$, jointly cancel to $r = 6\times10^{-15}$ [internal, `sheaf_consistency_radius.py`]. r never overstates the lie, but it can understate a coalition's to zero — soundness is unconditional, detection is not.
- *No attribution, ever.* A $+o$ lie by q toward w and a $-o$ lie by w toward q produce identical observations — the identical-data twin construction. The residual localizes to cycles through the equivocator; it cannot distinguish which endpoint of the offending relationship lied. Naming culprits is the signatures-and-forensics business [10], which this detector complements and must never be marketed as replacing.
- *Vanishing r is not an all-clear.* Three independent reasons: uniform (kernel) lies move every neighbor's view identically and are invisible in principle ($\|o\| = 5$, $r = 2 \times 10^{-14}$ [internal]) — such a lie is a consistent alternative state rather than a split view, so it falls outside what Theorem 1 quantifies over rather than contradicting it, and the exclusion is exact: on a cycle of K_c the uniform offsets are the *only* invisible ones, every kernel direction lying in the uniform span to 5×10^{-16} [internal, `sheaf_consistency_radius.py`]; severed edges admit consistent counterfactuals by construction; and over $\mathbb{R}$ the obstruction is abelianized, so a vanishing real-coefficient signal does not certify gluability in richer coefficient systems (the Carù gap [9]). $r = 0$ means "no conviction from this evidence," never "honest."
- *Severed darkness, and the noise floor.* Where no report crosses, the theorem is that nothing can be seen — and the damage is not confined to the severed link itself: severing one link of a loop darkens the whole loop, because the cycle Theorem 1 needs must live in K_c and no longer does. Budget accordingly (force comparison on bridges; keep reconciliation loops short; and count a loop as a loop only when every link on it reports). Under measurement noise, soundness needs a threshold above the honest floor $\|\Pi_K \eta\|$: the measured localization boundary is $\geq 95\%$ at noise-to-lie ratio 0.1 falling to 76.5% at 0.2 [internal], and below threshold the certificate reverts to "consistent with noise."

The scope discipline pays twice. Inside the boundary the claims are sharp enough to certify (a lower bound on the lie, a support statement, a solver cost); outside it the paper states its own counterexamples with witnesses, so the failure modes are part of the artifact rather than awaiting discovery by an adversary who read the appendix more carefully than the authors.

Reproducibility

Every [internal] number regenerates deterministically at seed 20260816 from three self-contained scripts shipped with the program's results estate: `sheaf_mechanism_proof.py` (the C_6/P_6 minimal pair: harmonic signal 1.225 versus 0.000), `sheaf_harness_v2.py` (all harness arms of §5, the pre-registered gates, the D1/D2 reintroduction mutants, and the structural β_1/cokernel printout), and `sheaf_consistency_radius.py` (CR-1/2/3: the 0/600–600/600 soundness/achievability sweep, the R_{eff} closed-form checks on C_4–C_{12} and 200 random graphs, the localization and noise-boundary measurements, the timing slopes, and the coalition-cancellation witness). Each script exits nonzero on any failed check. Every [verified] number is recomputable by hand from the closed forms restated in the boxes. Figures regenerate from `figures/src/` at the same seed. The one exception to bit-exact reproduction: CR-3's timing slopes are wall-clock and machine-dependent, so the script asserts and re-asserts only the qualitative gate (conjugate-gradient sub-1.6, dense solve several times steeper), not the printed digits.

References

[1] M. Herlihy and N. Shavit. The topological structure of asynchronous computability. *Journal of the ACM*, 46(6):858–923, 1999.

[2] M. Saks and F. Zaharoglou. Wait-free k-set agreement is impossible: the topology of public knowledge. *SIAM Journal on Computing*, 29(5):1449–1483, 2000.

[3] M. Herlihy, D. Kozlov, and S. Rajsbaum. *Distributed Computing Through Combinatorial Topology*. Morgan Kaufmann, 2013.

[4] A. Haeberlen, P. Kouznetsov, and P. Druschel. PeerReview: practical accountability for distributed systems. In *Proc. 21st ACM Symposium on Operating Systems Principles (SOSP)*, pp. 175–188, 2007.

[5] S. Abramsky and A. Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13:113036, 2011.

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

[7] J. Hansen and R. Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3:315–358, 2019.

[8] M. Robinson. Sheaves are the canonical data structure for sensor integration. *Information Fusion*, 36:208–224, 2017.

[9] G. Carù. On the cohomology of contextuality. In *Proc. 13th International Conference on Quantum Physics and Logic (QPL 2016)*, EPTCS 236, pp. 21–39, 2017.

[10] P. Sheng, G. Wang, K. Nayak, S. Kannan, and P. Viswanath. BFT protocol forensics. In *Proc. ACM Conference on Computer and Communications Security (CCS)*, pp. 1722–1743, 2021.

[11] D. A. Spielman and S.-H. Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proc. 36th ACM Symposium on Theory of Computing (STOC)*, pp. 81–90, 2004.

[12] E. Bach. Sheaf cohomology is #P-hard. *Journal of Symbolic Computation*, 27(4):429–433, 1999.