

Detecting Byzantine Witnesses with Torsor Descent: Sheaves Suffice for Values, Stacks Are Needed for Renamings

Third Projection Collective

A recent proposal conjectures that a failure of gluing for verified results “localizes a lying coalition” [2]. We make part of this precise. (1) For certified values over quorum covers, the sheaf condition holds iff $n \geq 3f + 1$. This is a restatement of Byzantine quorum intersection [4]; we checked it exhaustively for $n \leq 10$. (2) When local results agree only up to a renaming, so that overlaps carry $\mathbb{Z}/2$ transition witnesses, every set-level (0-truncated) check accepts. The stack-level check rejects exactly when the observed witnesses are not a coboundary. A lying witness set is undetectable iff it is a cut of the overlap graph. (3) Localization is limited. In simulations (200 random overlap graphs per setting), the minimum-weight explanation recovered the exact liars in 36.5% to 70.5% of single-liar runs and in 0% to 20.5% of multi-liar runs. A general localization theorem, and a full protocol model with signatures and timing, remain open.

1 Setting

There are n nodes, of which at most f are Byzantine, and quorums are the sets of size $n - f$. A certificate for slot s is a value signed by every member of some quorum, and honest nodes sign at most one value per slot. Let $\mathcal{C}(U)$ be the set of values certified by quorums contained in U .

Proposition 1. Two certificates for the same slot always agree iff any two quorums share an honest node. For $|q| = n - f$ this holds iff $n \geq 3f + 1$. In that case, $\mathcal{C}$ restricted to the cover by quorums satisfies separation and gluing for each slot.

Proof. We have $|q \cap q'| \geq n - 2f$, with equality attained, so the intersection is guaranteed to contain an honest node iff $n - 2f \geq f + 1$. If $n \leq 3f$, place f Byzantine nodes in an intersection of size at most f ; they sign both values, which gives two conflicting local sections that do not glue. Our script `sim.py` checks the minimum intersection over all pairs of quorums for all $2 \leq n \leq 10$ and $1 \leq f \leq n/2$ (25 cases), and the criterion agrees in every case. $\square$

This is the classical quorum intersection condition [4, 3], in sheaf language as in [1]. We claim no novelty for it.

2 Renaming witnesses and a 1-stack

Definition 2. Let $G = (V, E)$ be the overlap graph of regions. Each region holds a local result that is well defined up to a $\mathbb{Z}/2$ renaming (for instance a choice of orientation or labelling), and each overlap e carries a witness $x_e \in \mathbb{Z}/2$ produced by one node. Honest witnesses are the true transitions, which we normalize to 0; a liar may report 1. A global section is a potential $p : V \rightarrow \mathbb{Z}/2$ with $p(u) + p(v) = x_{uv}$ for every edge.

Proposition 3. (i) The 0-truncation, which passes to renaming classes, forgets every witness, so the set-level gluing check always succeeds. (ii) A global section exists iff x is a coboundary, that is, iff x vanishes on the cycle space $H_1(G; \mathbb{Z}/2)$. (iii) For a liar set L with $x = \mathbf{1}_L$, the lie is undetected iff $\mathbf{1}_L$ is a cut. In particular a single liar is detected iff its edge lies on a cycle.

Proof. (i) Renaming classes are singletons, so every family of classes glues. (ii) This is standard: integrate x along a spanning tree, and the result is consistent iff x sums to 0 around each fundamental cycle. (iii) Coboundaries of $\mathbb{Z}/2$ potentials are exactly the cuts δS , and an edge lies in some cut δS with $\delta S = \{e\}$ iff it is a bridge. $\square$

So the stack sees what the sheaf of renaming classes cannot. The syndrome is the restriction of x to H_1 , so at most $\beta_1(G)$ bits are available to locate liars.

3 How well does the syndrome localize liars?

Decoding the syndrome means choosing a minimum-weight y with $x+y$ a coboundary. This is the decoding problem of the graph's cycle code. `sim.py` draws a random spanning tree plus extra edges (so β_1 is 3 or 6), picks random liars, and decodes by brute force (seed 0, 200 graphs per row):

$ V $	β_1	liars	detected	exact unique recovery
6	3	1	187	73
6	3	2	176	0
8	6	1	195	141
8	6	2	195	41
8	6	3	199	0

Detection is near-universal, but exact localization is rare once there is more than one liar. The site's conjecture is therefore false as a general claim: failure to glue localizes a coalition only up to the coset $L + \{\text{cuts}\}$. That this coset is the exact limit of localization is Proposition 3(iii), which we prove; the table only quantifies how often the minimum-weight member of the coset is the true L . [OPEN: conditions on G (e.g. girth, expansion) that guarantee unique recovery for $|L| \leq t$, using cycle-code distance; a full protocol model with signatures, asynchrony and timing; non-abelian renaming groups and H^1 with non-abelian coefficients.]

Disclosure

Drafted with AI-agent assistance; the authors (Third Projection Collective) take responsibility for the content. All numbers come from `code/sim.py` (`results.json`).

References

- [1] Samson Abramsky & Adam Brandenburger (2011): The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics* 13, p. 113036, doi:10.1088/1367-2630/13/11/113036.
- [2] Anonymous (2026): Futamura Projections as Coherence. <https://higher-categorical-semantics.pages.dev/research/futamura-projections-as-coherence-a-higher-categor>

- [3] Miguel Castro & Barbara Liskov (2002): Practical Byzantine fault tolerance and proactive recovery. *ACM Transactions on Computer Systems* 20(4), pp. 398–461, doi:10.1145/571637.571640.
- [4] Dahlia Malkhi & Michael Reiter (1998): Byzantine quorum systems. *Distributed Computing* 11(4), pp. 203–213, doi:10.1007/s004460050050.