

Binary Order Thinning of Finite Posets: Homology, Collapses, and Sharp Limits

Brian Theory

27 September 2026

Abstract

Let P be a finite poset, let $b : P \rightarrow \{0, 1\}$, and thin the order by retaining $x \leq_P y$ exactly when $b(x) \leq b(y)$. We study the inclusion of order complexes $\Delta(Q) \hookrightarrow \Delta(P)$. When $\text{ht}(P) \leq 2$, the one-bit relative boundary $C_2(\Delta P, \Delta Q; \mathbb{Z}) \rightarrow C_1(\Delta P, \Delta Q; \mathbb{Z})$ is a reduced oriented graph-incidence matrix. This gives the complete integral relative homology, proves torsion-freeness, and yields an exact recognition theorem: vanishing relative homology over $\mathbb{F}_2$, over every field, or over $\mathbb{Z}$ is equivalent to a rooted-forest condition, to a perfect acyclic matching of the relative cells, and to an actual simplicial collapse $\Delta(P) \searrow \Delta(Q)$. The rigidity is sharp in two independent directions. At height two, two simultaneous Boolean coordinates already admit a seven-element example whose relative integral boundary has determinant -1 and whose endpoints are contractible, but for which no relative elementary collapse can even begin; an exhaustive enumeration verifies that no such two-bit example exists on at most six vertices. With unrestricted height, a one-bit cone-whisker construction realizes the homotopy type of an arbitrary nonempty finite poset inside a contractible coarse space and produces relative torsion already in height three. The graph-incidence and total-unimodularity ingredients are standard, as are rooted-forest/fitting-orientation methods and the general discrete-Morse mechanism linking appropriate acyclic matchings to collapse. The contribution isolated here is the exact incidence recognition forced by Boolean order thinning at height two, its coefficient-independent forest criterion, and the identification of the two sharp boundaries.

Keywords. finite posets; order complexes; relative homology; simplicial collapse; graph incidence; rooted forests; Boolean observations; discrete Morse theory.

1 Introduction

A simple operation on a finite poset already exhibits a useful boundary between rigid low-dimensional topology and more general cancellation phenomena. Let P be a finite poset and let

$$b : P \longrightarrow \{0, 1\}.$$

Define a same-vertex suborder Q by

$$x \leq_Q y \iff x \leq_P y \text{ and } b(x) \leq b(y). \tag{1}$$

Thus precisely the $1 \rightarrow 0$ comparisons are deleted. The central question of this paper is when the inclusion

$$\Delta(Q) \hookrightarrow \Delta(P)$$

is topologically harmless, and when homological harmlessness is already strong enough to force an explicit collapse.

Finite posets and their order complexes are classical, as are the links between finite spaces and posets due to McCord and subsequent developments [16, 1, 2]. Relative homology, graph incidence, total unimodularity, elementary collapse, and discrete Morse theory are also standard tools [13, 12, 15, 10]. The point here is not to repackage these theories. It is that the Boolean rule (1), together with the height-two restriction, forces a very special relative boundary which can be recognized exactly.

Two closer antecedents deserve explicit separation from the contribution claimed here. Mukherjee developed the topology of rooted forests and, in particular, collapse and relative-homology results for triangulated closed manifolds [17, 18], and Contreras–Tawfeek proved for rooted forests of arbitrary simplicial complexes that correspondence to a discrete gradient is equivalent to simplicial collapse onto the root [6]. Accordingly, the matching-to-collapse mechanism used below is not claimed as new; a short direct proof is retained because in the present two-layer setting it is elementary and makes the relative-collapse order explicit. A different nearby use of Boolean functions and order complexes is due to Björner–Goresky–MacPherson [5]: their complex is the induced order complex on the true states of a Boolean function, whereas (1) keeps the carrier and deletes order relations. The claim here is therefore narrower: the thinning rule at height two forces a reduced graph-incidence boundary, and that forced form yields the exact equivalences and sharp failure regimes proved below.

Our main positive result is the following. Height is counted by strict inequalities, so $\text{ht}(P) \leq 2$ means that no chain has four elements, equivalently $\dim \Delta(P) \leq 2$. Let $K = \Delta(P)$ and $L = \Delta(Q)$. All vertices survive the thinning. If E and T are the deleted edges and deleted triangles, then the only nonzero part of the integral relative chain complex is

$$B : C_2(K, L; \mathbb{Z}) = \mathbb{Z}^T \longrightarrow C_1(K, L; \mathbb{Z}) = \mathbb{Z}^E.$$

For one bit, every column of B has either one nonzero entry or two opposite nonzero entries. After adjoining one formal root for each coarse connected component, B is exactly an oriented graph-incidence matrix with the root rows removed.

This observation has three consequences. First, the entire integral relative homology is explicit and torsion-free. Second, coefficient choice becomes irrelevant. Third, homological neutrality is equivalent to a rooted forest, and leaf pruning of that forest is literally an elementary simplicial collapse. In particular, for a single Boolean coordinate at height at most two,

$$\boxed{H_*(K, L; \mathbb{F}_2) = 0 \iff H_*(K, L; \mathbb{Z}) = 0 \iff K \searrow L.}$$

The theorem also includes homotopy and simple-homotopy equivalence, a perfect acyclic matching of the relative cells, and homology over an arbitrary field.

The equivalence is sharp on two independent axes. The first is bit count. With two coordinates, the relative boundary need no longer be an ordinary graph incidence matrix. We give a seven-element height-two poset and a map $\alpha : P \rightarrow \{0, 1\}^2$ for which the relative 7×7 integral boundary has determinant -1 , yet no deleted edge is initially free. The support graph has three perfect matchings whose signed determinant contributions cancel to -1 . Both endpoint posets beat-reduce to a point, so the inclusion of order complexes is nevertheless a homotopy equivalence. A separate exhaustive computation over all height-two naturally labelled posets and all two-bit labels through six vertices finds no counterexample. Thus the seven-element example is vertex-minimal subject to that exhaustive certificate.

The second sharpness axis is height. For any nonempty finite poset R , adding a new greatest element and choosing one Boolean label in a particular way makes the coarse order complex a cone while the thinned order complex is $\Delta(R)$ with one collapsible whisker. Consequently unrestricted one-bit thinning contains arbitrary finite-poset homotopy types and arbitrary corresponding reduced homology shifts. In particular, one bit already produces relative torsion in height three. The height-two hypothesis therefore does real mathematical work; it is not merely a convenient dimension cutoff.

The paper is deliberately narrower than the broader logical-topology program from which the problem arose. We use general finite-poset language first. Section 9 returns to the observation-refinement interpretation only after the structural results are complete. We do not develop a separate observation-dimension theory, generic persistence constructions, formula-level exact sequences, or other branches that do not support the theorem above.

2 Order thinning, deleted cells, and relative chains

2.1 Finite-poset conventions

For a finite poset P , write $\Delta(P)$ for its order complex: its vertices are the elements of P , and a k -simplex is a strict chain

$$x_0 < x_1 < \cdots < x_k.$$

The height $\text{ht}(P)$ is the largest number of strict inequalities in a chain. Thus $\text{ht}(P) = \dim \Delta(P)$. Connected components of P mean connected components of its undirected comparability graph, equivalently of $\Delta(P)$.

Given $b : P \rightarrow \{0, 1\}$, define $Q = Q_b$ by (1). Since the Boolean order is transitive, Q is a partial order on the same ground set. We set

$$K = \Delta(P), \quad L = \Delta(Q).$$

Because Q and P have the same vertices, L is a spanning subcomplex of K . A simplex $x_0 < \cdots < x_k$ of K lies in L exactly when

$$b(x_0) \leq b(x_1) \leq \cdots \leq b(x_k).$$

Equivalently, it is deleted exactly when its Boolean word contains a $1 \rightarrow 0$ descent.

More generally, for a block of r Boolean coordinates

$$\alpha : P \longrightarrow \{0, 1\}^r,$$

we write Q_α for the coordinatewise thinning

$$x \leq_{Q_\alpha} y \iff x \leq_P y \text{ and } \alpha(x) \leq \alpha(y) \tag{2}$$

where the last order is coordinatewise. The one-bit case is $r = 1$.

2.2 The relative complex

The inclusion $L \subseteq K$ gives the usual relative chain complex $C_*(K, L; R)$ over a coefficient ring R [13]. Since L contains all vertices, $C_0(K, L; R) = 0$. If $\text{ht}(P) \leq 2$, there are no simplices above dimension two, so with E the deleted edges and T the deleted triangles,

$$0 \longrightarrow R^T \xrightarrow{B_R} R^E \longrightarrow 0 \tag{3}$$

is the complete relative complex. In particular,

$$H_2(K, L; R) = \ker B_R, \quad H_1(K, L; R) = \operatorname{coker} B_R, \quad H_0(K, L; R) = 0. \quad (4)$$

This is a homology calculation, not by itself a homotopy test. The one-bit theorem below explains why the special Boolean incidence structure supplies the missing homotopical information at height two.

3 One bit at height two: an incidence theorem

Assume throughout this section that $\operatorname{ht}(P) \leq 2$ and $b : P \rightarrow \{0, 1\}$. Orient every simplex by the poset order. For $x < y < z$,

$$\partial[x, y, z] = [y, z] - [x, z] + [x, y]. \quad (5)$$

There are eight Boolean words on a three-chain. Four are nondecreasing and hence survive. The four deleted words have the following relative boundaries.

word $b(x)b(y)b(z)$	deleted edges	integral coefficients in $\partial^{\operatorname{rel}}$
010	yz	+1
101	xy	+1
100	xy, xz	+1, -1
110	xz, yz	-1, +1

In particular, a deleted triangle contains exactly one or two deleted edges, never zero and never three.

Definition 3.1 (rooted incidence graph). Let E be the deleted edges. For every connected component C of P , introduce a formal root $*_C$. The multigraph $G_b(P)$ has vertex set

$$V(G_b(P)) = E \sqcup \{*_C : C \in \pi_0(P)\}.$$

Each deleted triangle τ contributes one graph edge. If τ has two deleted edges e, f , join the vertices e and f . If it has one deleted edge e , join e to the root of the coarse component containing τ . Orient each graph edge so that, after deleting the root rows, its incidence column agrees with the signs in (5).

Parallel graph edges are allowed. Every connected component of $G_b(P)$ contains at most one formal root: all of its deleted-edge vertices and triangle edges lie inside a single connected component of P , for which only one formal root was introduced.

Proposition 3.2 (integral incidence recognition). *With the ordered deleted edges as rows and deleted triangles as columns, the integral relative boundary*

$$B : C_2(K, L; \mathbb{Z}) \longrightarrow C_1(K, L; \mathbb{Z})$$

is the oriented vertex-edge incidence matrix of $G_b(P)$ with all formal-root rows deleted.

Proof. A 010 or 101 column has one coefficient +1. It is the reduced incidence column of an edge directed from the component root to the unique deleted-edge vertex. A 100 or 110 column has two coefficients +1 and -1, so it is the ordinary oriented incidence column of an edge between the two deleted-edge vertices. These are all deleted three-chain patterns. $\square$




Figure 1: The two possible incidence types for a deleted triangle in the one-bit, height-two regime. A triangle with one deleted edge is represented by an edge to the formal component root; a triangle with two deleted edges is represented by an ordinary edge between their vertices. Deleting the root rows recovers the relative boundary columns.

The proposition turns the complete relative homology into elementary graph topology. For a finite graph G , let $c(G)$ be its number of connected components and put

$$\beta_1(G) = |E(G)| - |V(G)| + c(G).$$

Let $c_0(G_b(P))$ be the number of connected components containing no formal root. Isolated roots are permitted and contribute neither a relative chain generator nor a homology generator.

Theorem 3.3 (complete integral relative homology). *For one-bit thinning of a finite poset of height at most two, the relative boundary matrix B is totally unimodular and*

$$H_2(K, L; \mathbb{Z}) \cong \mathbb{Z}^{\beta_1(G_b(P))}, \quad (6)$$

$$H_1(K, L; \mathbb{Z}) \cong \mathbb{Z}^{c_0(G_b(P))}, \quad (7)$$

$$H_0(K, L; \mathbb{Z}) = 0. \quad (8)$$

Consequently all integral relative homology groups are torsion-free, and for every field $\mathbb{k}$,

$$\dim_{\mathbb{k}} H_2(K, L; \mathbb{k}) = \beta_1(G_b(P)), \quad \dim_{\mathbb{k}} H_1(K, L; \mathbb{k}) = c_0(G_b(P)).$$

Proof. An oriented graph-incidence matrix is totally unimodular. One elementary proof is by induction on a square submatrix: a column with at most one nonzero reduces the determinant to a smaller minor, while if every column has two nonzeros then the opposite signs make the row sum zero. Deleting rows preserves total unimodularity. Thus Proposition 3.2 gives the first claim. This is the ordinary graph-incidence phenomenon; more general relations between total unimodularity and relative torsion are well established [10].

Consider one connected component of $G_b(P)$. If it contains a root, deleting the root row from its incidence matrix gives full row rank. More is needed over $\mathbb{Z}$ to identify the cokernel: a spanning tree of the component gives a square maximal minor of the reduced incidence matrix with determinant ± 1 . Hence the image is all of $\mathbb{Z}^{V-1}$, so the cokernel is zero; its kernel is the integral cycle space. If the component contains no root, the matrix is the ordinary incidence matrix of a connected graph, whose cokernel is $\mathbb{Z}$ and whose kernel is again its cycle space. Taking direct sums over graph components gives (6)–(8). The universal coefficient theorem for homology of the pair gives, for every field $\mathbb{k}$, a short exact sequence

$$0 \longrightarrow H_k(K, L; \mathbb{Z}) \otimes_{\mathbb{Z}} \mathbb{k} \longrightarrow H_k(K, L; \mathbb{k}) \longrightarrow \text{Tor}_1^{\mathbb{Z}}(H_{k-1}(K, L; \mathbb{Z}), \mathbb{k}) \longrightarrow 0.$$

The integral relative groups above are free, so the Tor term vanishes and the stated field dimensions follow. $\square$

Remark 3.4 (novelty boundary). The graph-incidence matrix, its total unimodularity, and the formulas for graph cycle and component spaces are standard. The candidate contribution is the recognition that the particular Boolean thinning rule forces the relative simplicial boundary into precisely this reduced-incidence form. No new general matrix-tree theorem is claimed.

4 The exact forest/collapse theorem

Call $G_b(P)$ a *rooted forest* if it is acyclic and every connected component contains exactly one formal root. This includes isolated roots.

For later use, let $L \subseteq K$ be any same-vertex simplicial inclusion with $\dim K \leq 2$. Its *relative support graph* $\Gamma(K, L)$ is the bipartite graph whose left vertices are the deleted edges $E = K_1 \setminus L_1$, whose right vertices are the deleted triangles $T = K_2 \setminus L_2$, and in which $e \in E$ is adjacent to $\tau \in T$ exactly when $e \subset \tau$. A perfect matching $M = \{(e_i, \tau_i)\}$ is *acyclic* if the directed graph on the matched pairs, with an arrow $i \rightarrow j$ whenever $i \neq j$ and $e_i \subset \tau_j$, is acyclic. This is the usual acyclicity condition for the corresponding matching of the relative face poset in this two-layer setting.

Lemma 4.1 (two-layer matching–collapse lemma). *Let $L \subseteq K$ be as above. This is the present two-layer specialization of the standard acyclic-matching/collapse principle when the unmatched cells form the target subcomplex; compare Forman [12], Kozlov [15], and the rooted-forest formulation of Contreras–Tawfeek [6]. The following are equivalent:*

- (i) K collapses simplicially onto L using only deleted edge–triangle pairs;
- (ii) $\Gamma(K, L)$ has an acyclic perfect matching;
- (iii) the poset of relative cells $K \setminus L$, ordered by inclusion, has a perfect acyclic matching.

Proof. A relative collapse sequence pairs every deleted triangle with the deleted edge removed with it. If the pair (e_i, τ_i) is removed at step i and $e_i \subset \tau_j$ with $j \neq i$, then τ_j must already have been removed before step i ; otherwise e_i would not be free. Thus every arrow in the matched-pair digraph points strictly backward in collapse order, so the resulting perfect matching is acyclic.

Conversely, let $M = \{(e_i, \tau_i)\}$ be an acyclic perfect matching. Its matched-pair digraph has a sink i . By definition of sink, e_i lies in no remaining deleted triangle except τ_i . No retained triangle can contain a deleted edge, because L is a subcomplex, and there are no simplices above dimension two. Hence e_i is a free face of τ_i . Collapse this pair and restrict M to the remaining relative cells. The restricted matched-pair digraph is still acyclic, so iteration removes all relative cells. This proves (ii) $\Rightarrow$ (i); condition (iii) is the same matching condition expressed in the relative face poset. $\square$

Theorem 4.2 (one-bit height-two recognition theorem). *Let P be a finite poset with $\text{ht}(P) \leq 2$, let $b : P \rightarrow \{0, 1\}$, define Q by (1), and put $K = \Delta(P)$ and $L = \Delta(Q)$. The following are equivalent.*

- (i) $H_*(K, L; \mathbb{F}_2) = 0$.
- (ii) $H_*(K, L; \mathbb{k}) = 0$ for every field $\mathbb{k}$.
- (iii) $H_*(K, L; \mathbb{Z}) = 0$.
- (iv) $G_b(P)$ is a rooted forest.
- (v) the poset of relative cells $K \setminus L$ admits a perfect acyclic matching (equivalently, $\Gamma(K, L)$ has an acyclic perfect matching);
- (vi) K collapses simplicially onto L by deleted edge–triangle pairs;
- (vii) the inclusion $L \hookrightarrow K$ is a simple-homotopy equivalence;
- (viii) the inclusion $L \hookrightarrow K$ is a homotopy equivalence.

When these conditions hold and deleted cells are nonempty, B is square and $\det B = \pm 1$ after compatible row and column orderings.

Proof. Theorem 3.3 makes (i)–(iv) equivalent: vanishing is exactly $\beta_1(G_b(P)) = 0$ and $c_0(G_b(P)) = 0$, which says that every graph component is a tree containing its unique possible root.

Assume (iv). Any nontrivial rooted tree has a nonroot leaf. The corresponding deleted edge e lies in exactly one currently remaining deleted triangle τ . No retained triangle can contain a deleted edge, because L is a subcomplex, and there are no cells above dimension two. Thus e is a free face of τ in the current complex. Delete the pair (e, τ) . On $G_b(P)$ this is exactly pruning a nonroot leaf and its incident graph edge, leaving a rooted forest. Iteration removes all deleted cells and proves (vi). Recording the pairs in the pruning order yields an acyclic perfect matching of the relative cells, proving (v) as well.

Lemma 4.1 gives the converse (v) $\Rightarrow$ (vi) directly in this two-layer situation. This step is a transparent specialization of standard discrete-Morse collapse theory, not a new matching-to-collapse theorem. A simplicial collapse gives a simple-homotopy equivalence and therefore a homotopy equivalence, so (vi) $\Rightarrow$ (vii) $\Rightarrow$ (viii). Finally a homotopy equivalence induces an integral homology equivalence, giving (viii) $\Rightarrow$ (iii) and closing the cycle.

In the neutral nonempty case, (3) has zero kernel and cokernel, so $|E| = |T|$ and B is an integral isomorphism. Total unimodularity then forces determinant ± 1 . $\square$

Corollary 4.3 (greedy collapse algorithm). *After the deleted edges and triangles have been constructed, homotopy neutrality of a one-bit height-two thinning can be decided in time linear in the size of the defect incidence data: repeatedly remove a deleted edge incident to exactly one remaining deleted triangle together with that triangle. The procedure succeeds exactly when all deleted cells are exhausted.*

Proof. This is nonroot-leaf pruning in $G_b(P)$. Equivalently, one may test directly that each graph component is a tree with exactly one root. $\square$

Remark 4.4. The theorem is stronger than a matrix-rank test for a particular coefficient field. The extra force comes from the one-bit incidence structure, not from homology alone. Section 7 gives a unimodular relative boundary for which direct collapse fails as soon as a second coordinate is allowed.

5 Structural sufficient conditions and logical examples

The exact theorem above is specific to height two, but constructive sufficient conditions are available in arbitrary height. We record only the forms that will be useful for interpreting Boolean observations.

Proposition 5.1 (repair maps). *Let P be any finite poset and $b : P \rightarrow \{0, 1\}$.*

- (a) *If there is an order-preserving $\rho : P \rightarrow P$ such that $\rho(p) \leq p$ and $b(\rho(p)) = 0$ for every p , then $\Delta(Q_b) \hookrightarrow \Delta(P)$ is a homotopy equivalence.*
- (b) *Dually, if there is an order-preserving $\lambda : P \rightarrow P$ such that $p \leq \lambda(p)$ and $b(\lambda(p)) = 1$ for every p , the same conclusion holds.*

Proof. In (a), ρ is also monotone as a map $P \rightarrow Q_b$, and $i\rho \leq \text{id}_P$ and $\rho i \leq \text{id}_{Q_b}$. Pointwise comparable monotone maps induce homotopic maps on order complexes, so i and ρ are homotopy inverses; see, for example, the comparable-map machinery in Barmak [1]. Part (b) is dual. $\square$

Let $F = b^{-1}(0)$ and $T = b^{-1}(1)$. Proposition 5.1 immediately gives the following closure-system criteria.

Corollary 5.2 (floors, ceilings, and semilattice closure). *Each of the following is sufficient for homotopy preservation.*

- (a) *For every $p \in P$, $F \cap \downarrow p$ has a greatest element.*
- (b) *For every $p \in P$, $T \cap \uparrow p$ has a least element.*
- (c) *P is a finite join-semilattice, F is join-closed, and $F \cap \downarrow p \neq \emptyset$ for every p .*
- (d) *P is a finite meet-semilattice, T is meet-closed, and $T \cap \uparrow p \neq \emptyset$ for every p .*

In (c) the repair is $\rho(p) = \bigvee(F \cap \downarrow p)$; in (d) it is $\lambda(p) = \bigwedge(T \cap \uparrow p)$.

Corollary 5.3 (Horn-type specialization). *Suppose a realized-signature poset is a meet-subsemilattice of a Boolean cube and the truth states of b are the models of a Horn CNF in the old coordinates. If every realized signature lies below at least one such model, then the thinning preserves homotopy type. Dually, on a join-subsemilattice, a union-closed false region that is lower cofinal gives a down-repair.*

Remark 5.4. The closure properties of Horn and dual-Horn relations are classical [14, 19]. The corollary is only their use as an explicit repair certificate; no blanket claim that every Horn-defined observation is homotopy-neutral is intended. The cofinality and semilattice hypotheses are essential parts of the statement.

Example 5.5 (the Boolean diamond). Let P be the Boolean diamond on $D = \emptyset$, $A = \{X\}$, $C = \{Y\}$, and $B = \{X, Y\}$. Thus $D < A < B$ and $D < C < B$. The following four one-bit profiles illustrate the criteria and the exact graph formula.

observation	profile on (D, A, C, B)	certificate	effect
$X \vee Y$	0, 1, 1, 1	isotone	$Q = P$
$X \text{ XOR } Y$	0, 1, 1, 0	rooted tree	proper thinning; homotopy-neutral
$X \Leftrightarrow Y$	1, 0, 0, 1	rooted tree / true ceiling	proper thinning; homotopy-neutral
$\neg(X \vee Y)$	1, 0, 0, 0	one rootless component	topology changes

For XOR the two deleted edges $A < B$ and $C < B$ are each anchored by a one-deletion triangle, giving a rooted two-leaf tree. XNOR is dual. For NOR the root is isolated and the three deleted-edge vertices form an unrooted path; Theorem 3.3 gives $H_1(K, L; \mathbb{Z}) \cong \mathbb{Z}$.

6 A controlled chain corollary

The complete one-bit classification for a coarse chain is included only as a compact order-theoretic corollary. If $P = C_n$, two positions $i < j$ are incomparable in the thinned order exactly when $b_i = 1$ and $b_j = 0$. Hence the incomparability graph of Q is a bipartite Ferrers graph (the neighborhoods of the 1-vertices are nested), Writing $\text{Inc}(Q)$ for this incomparability graph, pairwise comparable sets in a poset are chains, so

$$\Delta(Q) = \text{Ind}(\text{Inc}(Q)),$$

the independence complex of that Ferrers graph. Dochtermann–Engström give the corresponding contractible-versus-two-point homotopy dichotomy for independence complexes of Ferrers graphs [9, Proposition 3.7]; related Ferrers-graph topology is also developed by Claesson–Kitaev–Ragnarsson–Tenner [8]. We therefore retain the result for its explicit Boolean-word formulation and elementary beat-point proof, not as a novelty claim.

Theorem 6.1 (binary chain trichotomy). *Let $n \geq 1$, let $P = C_n$ be the n -element chain, and write $b_i = b(i)$. Exactly one of the following occurs.*

- (a) *The word $b_1 \cdots b_n$ is nondecreasing. Then $Q = P$.*
- (b) *The word is nonconstant and nonincreasing, hence $1^a 0^{n-a}$ for some $1 \leq a < n$. Then $Q \cong C_a \sqcup C_{n-a}$; in particular the thinning changes H_0 .*
- (c) *In every other case, Q is beat-contractible and $\Delta(Q) \hookrightarrow \Delta(P)$ is a homotopy equivalence.*

Consequently there are $n + 1$ profiles with $Q = P$, $n - 1$ profiles for which the thinning changes H_0 , and $2^n - 2n$ remaining profiles with $Q \neq P$ for which $\Delta(Q) \hookrightarrow \Delta(P)$ is nevertheless a homotopy equivalence.

Proof. Case (a) is immediate from (1). In case (b), all comparisons inside the initial 1-block and final 0-block survive, while every comparison from the first block to the second is deleted, giving the stated disjoint union.

For case (c), if $b_1 = 0$ then the first chain element remains a global minimum of Q ; if $b_n = 1$ the last remains a global maximum. Either condition makes $\Delta(Q)$ a cone. The only remaining endpoint pattern is $b_1 = 1, b_n = 0$. Since the word is not nonincreasing, some 0 occurs before a later 1. Remove the initial 1-vertices one at a time: for each such vertex, the first later 1 is the least element of its strict upper set in the current poset, so it is an up-beat point. After these removals the first remaining vertex has label 0 and is a global minimum. Thus Q beat-reduces to a point. The coarse chain is also contractible, and any map between nonempty contractible spaces is a homotopy equivalence. Counting binary words gives the final formula. $\square$

7 Sharpness I: two bits already fail

For a general block thinning at height two, the relative complex still has the form (3), but its columns need not be graph-incidence columns. The relative support graph $\Gamma(K, L)$ introduced before Lemma 4.1 is therefore the natural combinatorial object. Perfect matchings are the *fitting orientations* of higher-dimensional rooted-forest theory in this special dimension [4, 11].

Proposition 7.1 (direct collapse and unique fitting orientation). *For any same-vertex block thinning $Q \subseteq P$ with $\text{ht}(P) \leq 2$, the following are equivalent:*

- (i) *$\Delta(P)$ collapses directly to $\Delta(Q)$ by deleted edge-triangle pairs;*
- (ii) *the relative support graph has an acyclic perfect matching;*
- (iii) *the relative support graph has exactly one perfect matching.*

The empty support graph has its unique empty matching and represents the identity collapse.

Proof. Lemma 4.1 gives (i) $\iff$ (ii). For a perfect matching, a directed cycle in the matched-pair digraph is exactly an alternating cycle of the bipartite support graph. Flipping an alternating cycle gives a second perfect matching. Conversely, the symmetric difference of two distinct perfect matchings is a nonempty disjoint union of alternating cycles. Thus a perfect matching is acyclic exactly when it is unique, proving (ii) $\iff$ (iii). The matching/fitting-orientation viewpoint is standard [12, 4, 6]; the uniqueness argument above is only the bipartite two-layer form needed here. $\square$

We now give the sharp counterexample. It is stated completely so that the existence claim does not depend on the exhaustive search used later for minimality.

Theorem 7.2 (seven-element two-bit counterexample). *There is a seven-element poset P of height two and a map $\alpha : P \rightarrow \{0,1\}^2$ such that, for the coordinatewise thinned suborder Q ,*

$$\Delta(Q) \hookrightarrow \Delta(P)$$

is an integral homology equivalence and a homotopy equivalence, but $\Delta(P) \not\searrow \Delta(Q)$ by a relative elementary-collapse sequence.

Proof. Let P be the transitive closure of the cover relations

$$\begin{aligned} 0 < 2, \ 0 < 3, \ 0 < 4, \ 1 < 2, \ 1 < 4, \\ 2 < 5, \ 2 < 6, \ 3 < 5, \ 3 < 6, \ 4 < 5. \end{aligned} \tag{9}$$

Every maximal chain has three elements, so $\text{ht}(P) = 2$. Define

x	0	1	2	3	4	5	6
$\alpha(x)$	10	01	01	00	11	00	01

The retained strict comparisons in Q are

$$(0, 4), (1, 2), (1, 4), (1, 6), (2, 6), (3, 5), (3, 6). \tag{10}$$

Order the deleted edges as

$$\begin{aligned} e_1 &= (0, 2), \ e_2 = (0, 3), \ e_3 = (0, 5), \ e_4 = (0, 6), \\ e_5 &= (1, 5), \ e_6 = (2, 5), \ e_7 = (4, 5), \end{aligned}$$

and the deleted triangles as

$$\begin{aligned} t_1 &= (0, 2, 5), \ t_2 = (0, 2, 6), \ t_3 = (0, 3, 5), \ t_4 = (0, 3, 6), \\ t_5 &= (0, 4, 5), \ t_6 = (1, 2, 5), \ t_7 = (1, 4, 5). \end{aligned}$$

With the order orientation, the relative integral boundary is

$$B = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ -1 & 0 & -1 & 0 & -1 & 0 & 0 \\ 0 & -1 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & -1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{bmatrix}, \quad \det B = -1. \tag{11}$$

Thus B is an integral isomorphism and the relative integral homology vanishes in all degrees. Hence the inclusion is an integral homology equivalence (and is a homology equivalence over every field).

The support graph has exactly three perfect matchings. As permutations assigning to row e_i the column $t_{p(i)}$, they are

$$\begin{aligned} (1, 4, 3, 2, 7, 6, 5), \\ (2, 3, 1, 4, 7, 6, 5), \\ (2, 3, 5, 4, 6, 1, 7), \end{aligned}$$




Figure 2: Relative support graph of the seven-state two-bit example. Every deleted edge has degree at least two, so no relative elementary collapse can start. The graph nevertheless has three perfect matchings, listed in the text, whose signed determinant contributions sum to -1 .

where the displayed indices are one-based. Their determinant contributions are, respectively, $-1, +1, -1$, summing to -1 . This is precisely the kind of cancellation among fitting orientations already present in the general higher-dimensional rooted-forest framework [4]; the point here is that two Boolean coordinates can realize it inside this order-thinning problem.

More directly, the seven deleted-edge degrees into deleted triangles are

$$(2, 2, 3, 2, 2, 2, 2).$$

No deleted edge is initially free. Since every vertex is retained and there are no simplices above dimension two, a relative elementary collapse would have to start by pairing a deleted edge with its unique incident deleted triangle. No such first move exists. Therefore $K \not\searrow L$. Proposition 7.1 gives the same conclusion from the existence of three perfect matchings.

It remains to check homotopy equivalence. Both endpoint posets beat-reduce to a point. One valid reduction for P is

$$3 \downarrow 0, \quad 4 \uparrow 5, \quad 0 \uparrow 2, \quad 1 \uparrow 2, \quad 5 \downarrow 2, \quad 2 \uparrow 6,$$

and for Q it is

$$0 \uparrow 4, \quad 2 \downarrow 1, \quad 4 \downarrow 1, \quad 1 \uparrow 6, \quad 5 \downarrow 3, \quad 3 \uparrow 6.$$

Here $x \downarrow y$ means that x is removed as a down-beat point with the indicated extremal lower neighbor, and similarly for $\uparrow$. Thus both order complexes are nonempty and contractible. Their inclusion is therefore a homotopy equivalence. $\square$

The preceding theorem is a finite explicit proof. Vertex minimality is a separate, computer-assisted statement.

Proposition 7.3 (computer-assisted vertex minimality). *No height-two two-bit counterexample of the type in Theorem 7.2 exists on at most six vertices. Consequently the seven-element example is vertex-minimal, conditional only on the correctness of the exhaustive checker and its stated enumeration model.*

Exhaustive certificate. For each $n \leq 6$, the checker enumerates every naturally labelled transitive relation on $\{0, \dots, n-1\}$ of height at most two and every one of the 4^n maps to $\{0, 1\}^2$. Every finite

poset admits a linear extension, so natural labellings cover all isomorphism types, with harmless duplication. For each pair it constructs the coordinatewise thinned order, the deleted edges and triangles, and the relative matrix over $\mathbb{F}_2$. Relative acyclicity requires a square full-rank matrix. For each such neutral case, the program tests uniqueness of a perfect matching of the support graph and uses Proposition 7.1: a neutral example fails direct collapse exactly when the perfect matching is not unique. The search intentionally tests $\mathbb{F}_2$ -neutrality rather than the stronger integral-homology and endpoint-homotopy conditions in Theorem 7.2. This enlarges the candidate class: integral relative acyclicity implies $\mathbb{F}_2$ -relative acyclicity, so any smaller counterexample of the theorem's type would necessarily be detected by this weaker filter. Therefore failure to find a noncollapsible case in the larger $\mathbb{F}_2$ -neutral class suffices for the stated lower bound; no separate homotopy test is needed for exclusion.

A fresh independent rerun on 27 September 2026 produced

n	posets	two-bit profiles	square cases	$\mathbb{F}_2$ -neutral	max. matrix dim.
1	1	4	4	4	0
2	2	32	25	25	0
3	7	448	256	256	1
4	39	9,984	4,174	4,174	2
5	330	337,920	104,114	103,562	4
6	4,117	16,863,232	3,865,228	3,767,044	8

No neutral case through $n = 6$ had more than one perfect matching. The independent bounded-check script, its recorded output, and the fixed seven-state rerun are available in the website certificate package; they are not embedded in this PDF. Section 10 records the distinction between this exhaustive evidence and the analytic proof of Theorem 7.2. $\square$

Theorem 7.4 (sharp two-bit failure with certified minimality). *There exists a seven-element poset P of height two and a map $\alpha : P \rightarrow \{0, 1\}^2$ such that, for the coordinatewise thinned suborder Q , the inclusion $\Delta(Q) \hookrightarrow \Delta(P)$ is an integral homology equivalence and a homotopy equivalence but $\Delta(P) \not\searrow \Delta(Q)$ by a relative elementary-collapse sequence. No such two-bit example exists on at most six vertices.*

The existence assertion is the explicit mathematical construction of Theorem 7.2. The vertex-minimality sentence is computer-assisted and is exactly Proposition 7.3; it is not being presented as a separate noncomputational proof.

Corollary 7.5 (bit-count sharpness). *At height two, one bit is the maximal universal regime in which homology neutrality forces direct collapse. For every block size $r \geq 2$, there are height-two r -coordinate thinnings that are homology-neutral but do not collapse directly.*

Proof. The one-bit assertion is Theorem 4.2, and the two-bit failure is Theorem 7.2. For any $r > 2$, append $r - 2$ constant coordinates to the seven-state two-bit labeling. Coordinatewise comparability is unchanged, so the thinned suborder and the same noncollapse certificate persist. $\square$

8 Sharpness II: greater height destroys one-bit rigidity

The height hypothesis is independently sharp. The following construction is elementary but strong enough to show why no unrestricted one-bit analogue of Theorem 4.2 can hold.

Theorem 8.1 (cone-whisker realization). *Let R be any nonempty finite poset and choose a minimal element $z \in R$. Adjoin a new greatest element t to form $P = R \cup \{t\}$. Define*

$$b(z) = b(t) = 0, \quad b(x) = 1 \quad (x \in R \setminus \{z\}).$$

Then the one-bit thinning Q retains all comparisons of R and precisely one new comparison involving t , namely $z < t$. Therefore

$$\Delta(P) = \text{cone}(\Delta R), \quad \Delta(Q) = \Delta(R) \cup_{\{z\}} [z, t], \quad (12)$$

and $\Delta(Q)$ collapses onto $\Delta(R)$. In particular,

$$\Delta(Q) \hookrightarrow \Delta(P) \text{ is a homotopy equivalence} \iff \Delta(R) \text{ is contractible}, \quad (13)$$

and for $k \geq 1$,

$$H_k(\Delta P, \Delta Q; \mathbb{Z}) \cong \tilde{H}_{k-1}(\Delta R; \mathbb{Z}). \quad (14)$$

Proof. Every comparison internal to R survives: it is either $0 \rightarrow 1$ (when it starts at z) or $1 \rightarrow 1$, apart from identities. For $x \in R \setminus \{z\}$, the new comparison $x < t$ is $1 \rightarrow 0$ and is deleted, while $z < t$ is $0 \rightarrow 0$ and survives. Since t is greatest in P , $\Delta(P)$ is a cone. In $\Delta(Q)$ the edge $[z, t]$ is a whisker and t is a free vertex, so deleting it leaves $\Delta(R)$. Equation (13) follows because the coarse complex is contractible, and (14) follows from the long exact sequence of the pair together with (12). $\square$

Corollary 8.2 (arbitrary refined homotopy type). *Within a contractible coarse order complex, one Boolean coordinate of unrestricted height can realize, up to a single collapsible whisker, the order-complex homotopy type of any nonempty finite poset.*

Corollary 8.3 (torsion in height three). *There are one-bit height-three thinnings with relative integral torsion. In particular, taking a height-two finite-poset model of $\mathbb{RP}^2$ as R in Theorem 8.1 gives*

$$H_2(\Delta P, \Delta Q; \mathbb{Z}) \cong \mathbb{Z}/2.$$

A 13-point projective-plane finite model is standard prior art [7], so this yields a 14-state example. More generally, choosing R with torsion in reduced homology transfers that torsion one degree upward to the relative pair.

Remark 8.4. Theorem 8.1 also shows that integral homology neutrality need not imply homotopy neutrality beyond height two: choose an acyclic but noncontractible finite complex and take its face poset for R . This is exactly the separation that the rooted incidence structure rules out in Theorem 4.2.

9 Positioning and the logical-observation interpretation

9.1 What is standard and what is specific here

The mathematical ingredients used in the proofs have substantial prior art. Finite-poset/order-complex methods and finite-space realization are classical [16, 1, 2]. Elementary collapse and acyclic matchings belong to standard simple-homotopy and discrete-Morse theory [12, 15]. Higher-dimensional rooted forests, fitting orientations, determinant weights, and cancellation among several

fitting orientations are established in the work of Bernardi–Klivans and related higher-dimensional tree theory [4, 11]. Mukherjee developed topological and collapse/homology criteria for rooted forests, including triangulated closed manifolds [17, 18], and Contreras–Tawfeek characterized, for rooted forests in arbitrary simplicial complexes, when the forest corresponds to a discrete gradient in terms of simplicial collapse to its root [6]. Thus neither the generic rooted-forest language nor the matching-to-collapse mechanism is a novelty claim of this paper. Total-unimodularity and torsion questions for simplicial boundary matrices likewise have a broader established setting [10]. The Horn closure statements in Section 5 rely on classical closure properties rather than a new classification of Horn logic [14, 19]. Finally, Björner–Goresky–MacPherson attach order complexes to the true-state subposet of a Boolean function in their study of circuit complexity [5]. That construction is topically close but mathematically different from the same-carrier relation thinning studied here.

The narrower contribution claimed by this paper is the exact recognition theorem for the particular Boolean order-thinning rule: one coordinate plus height at most two forces the defect into an ordinary reduced graph-incidence matrix. Combining this forced incidence form with the standard collapse machinery makes coefficient independence, rooted forests, perfect acyclic matchings of the relative cells, and actual collapse equivalent for this thinning problem. The seven-element two-bit example and the cone–whisker construction then identify clean limits of that recognition theorem. We make no claim that the underlying graph, matching, or rooted-forest theories are new.

9.2 Propositional logical topology interpretation

In an observation-generated finite logical topology, one may pass to the Kolmogorov quotient and regard the specialization structure as a finite poset of realized observation signatures. The same-carrier thinning model applies literally when the new Boolean observation *factors through the current T_0 quotient*: equivalently, it is constant on every current observational-equivalence class and therefore does not split a quotient point into several newly distinguishable states. Under this hypothesis, the new observation retains a coarse comparison exactly when its truth value is nondecreasing along that comparison. A single new observation is then precisely the thinning (1), and a simultaneous block that factors through the current quotient is (2).

If a new observation separates worlds that were identified in the current quotient, the refined T_0 carrier itself gains points; that operation is not merely a same-carrier order thinning and is outside the literal scope of the theorems proved here.

This statement should not be confused with McCord’s theorem. McCord’s canonical map from the order-complex realization to a finite T_0 space is a weak homotopy equivalence [16]; it is not an assertion that every such realization map is an ordinary homotopy equivalence. Our main theorems are statements about the order complexes $\Delta(Q) \subseteq \Delta(P)$. Beat-point reductions are used only where they explicitly contract the relevant finite posets/order complexes, and relative homology computations are never presented as homotopy tests without an additional theorem such as Theorem 4.2.

For the logical program, the operational message is correspondingly narrow. A single Boolean observation on a height-two current signature poset admits an exact, integral, combinatorial safety test and an explicit collapse certificate. Simultaneous observation blocks do not inherit that equivalence automatically, and higher-height models require genuinely stronger hypotheses such as the repair maps of Section 5.

10 Reproducibility and proof status

The manuscript is accompanied by a `reproducibility/` directory. The principal files are:

- `verify_two_bit_height_two_counterexample.py`, an exact checker for Theorem 7.2;
- `two_bit_height2_exhaustive.cpp`, the exhaustive checker used for Proposition 7.3, with supplied and rerun outputs through $n = 6$;
- retained one-bit verifiers for the height-two matrix criterion, repair-map sufficient conditions, chain classification, and repair-forest equivalence;
- a reproducibility README and SHA-256 manifest.

There are two logically different kinds of claims. The one-bit incidence theorem, Theorem 4.2, the explicit seven-element construction, and the cone-whisker theorem have mathematical proofs in the text. The statement that seven vertices are minimal for the two-bit failure is computer-assisted: the program exhausts the finite search space through six vertices, but the paper does not disguise that enumeration as a noncomputational proof. For the exclusion step it is enough to enumerate the larger class of $\mathbb{F}_2$ -neutral cases: every integral-homology-neutral candidate of the target type is $\mathbb{F}_2$ -neutral, so a smaller target counterexample could not be filtered out by omitting an endpoint-homotopy test. The exact seven-state verifier independently reconstructs the orders, the relative matrix, its determinant and $\mathbb{F}_2$ rank, the matching count, the deleted-edge degrees, and the two beat reductions.

The retained one-bit checks are supplementary regression tests rather than foundations for the proof. In the post-audit rerun, the repair-forest verifier checked all height-at-most-two naturally labelled posets through six vertices (4,117 posets at $n = 6$, 263,488 one-bit cases there) and reported all assertions passed. The independent height/homology and chain scripts likewise passed their stated ranges. These finite checks help guard transcription and implementation errors but are not substituted for the proofs in Sections 3–4.

11 Conclusion

A one-bit observation on a height-two finite poset has an exactly classifiable incidence defect. The relative integral boundary is a reduced graph-incidence matrix, so relative homology is torsion-free and determined by graph cycles and rootless components. In this regime, homology neutrality is not merely necessary for topological neutrality: it forces a rooted forest and therefore an explicit simplicial collapse.

Both hypotheses are sharp. Keeping height two but allowing two Boolean coordinates already permits determinant cancellation among several fitting orientations; the seven-element example is homology- and homotopy-neutral but admits no direct relative collapse, and exhaustive computation certifies that six vertices do not suffice. Keeping one bit but dropping the height bound allows arbitrary finite-poset homotopy types inside a contractible coarse cone and relative torsion in height three. These two failures isolate the precise scope in which the one-bit forest theorem should be used.

References

- [1] J. A. Barmak, “On Quillen’s Theorem A for posets,” *Journal of Combinatorial Theory, Series A* 118 (2011), 2445–2453. <https://arxiv.org/abs/1005.0538>.

- [2] J. A. Barmak, *Algebraic Topology of Finite Topological Spaces and Applications*, Lecture Notes in Mathematics 2032, Springer, 2011.
- [3] J. A. Barmak and E. G. Minian, “Simple homotopy types and finite spaces,” *Advances in Mathematics* 218 (2008), 87–104.
- [4] O. Bernardi and C. J. Klivans, “Directed rooted forests in higher dimension,” *Electronic Journal of Combinatorics* 23(4) (2016), P4.35. DOI: 10.37236/5819.
- [5] A. Björner, M. Goresky, and R. MacPherson, “Topological aspects of Boolean functions,” *Pure and Applied Mathematics Quarterly* 20(3) (2024), 1029–1063. DOI: 10.4310/PAMQ.2024.v20.n3.a1.
- [6] I. Contreras and A. R. Tawfeek, “On discrete gradient vector fields and Laplacians of simplicial complexes,” *Annals of Combinatorics* 28(1) (2024), 67–91. DOI: 10.1007/s00026-023-00655-1.
- [7] N. Cianci and M. Ottina, “Poset splitting and minimality of finite models,” *Journal of Combinatorial Theory, Series A* 157 (2018), 120–161. DOI: 10.1016/j.jcta.2018.02.010.
- [8] A. Claesson, S. Kitaev, K. Ragnarsson, and B. E. Tenner, “Boolean complexes for Ferrers graphs,” *Australasian Journal of Combinatorics* 48 (2010), 159–173.
- [9] A. Dochtermann and A. Engström, “Algebraic properties of edge ideals via combinatorial topology,” *Electronic Journal of Combinatorics* 16(2) (2009), R2. DOI: 10.37236/68.
- [10] T. K. Dey, A. N. Hirani, and B. Krishnamoorthy, “Optimal homologous cycles, total unimodularity, and linear programming,” *SIAM Journal on Computing* 40(4) (2011), 1026–1044. <https://arxiv.org/abs/1001.0338>.
- [11] A. M. Duval, C. J. Klivans, and J. L. Martin, “Simplicial and cellular trees,” in *Recent Trends in Combinatorics*, IMA Volumes in Mathematics and its Applications 159, Springer, 2016, 713–752.
- [12] R. Forman, “Morse theory for cell complexes,” *Advances in Mathematics* 134 (1998), 90–145. DOI: 10.1006/aima.1997.1650.
- [13] A. Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
- [14] P. Jeavons, D. Cohen, and M. Gyssens, “Closure properties of constraints,” *Journal of the ACM* 44(4) (1997), 527–548. DOI: 10.1145/263867.263489.
- [15] D. Kozlov, *Combinatorial Algebraic Topology*, Algorithms and Computation in Mathematics 21, Springer, 2008.
- [16] M. C. McCord, “Singular homology groups and homotopy groups of finite topological spaces,” *Duke Mathematical Journal* 33 (1966), 465–474. DOI: 10.1215/S0012-7094-66-03352-7.
- [17] S. K. Mukherjee, “On the topology of rooted forests in higher dimensions,” *Topology and its Applications* 247 (2018), 50–56. DOI: 10.1016/j.topol.2018.07.013.
- [18] S. K. Mukherjee, “On the rooted forests in triangulated closed manifolds,” *Linear and Multilinear Algebra* 68(10) (2020), 2034–2043. DOI: 10.1080/03081087.2019.1570066.
- [19] M. Wild, “The joy of implications, aka pure Horn formulas: Mainly a survey,” *Theoretical Computer Science* 658 (2017), 264–292.