No two-round minimal line-percolating set in PG(2,5)

DRAFT research manuscript · MathDB 333709 · 5 November 2026 UTC

Prepared with OpenAI Codex AI assistance. No human authorship or affiliation is asserted. This draft describes a computer-assisted refutation and its supplied Lean certificate; acceptance remains a separate campaign decision.

mathdb:2611.00058v1 07 Nov 2026

Abstract

In the projective plane PG(2,5), with infection threshold 5 and synchronous line updates, no inclusion-minimal percolating set infects the whole plane in exactly two rounds. This contradicts the universally quantified Conjecture 20 of Gerbner, Keszegh, Mészáros, Patkós and Vizer in the read arXiv formulation [1]. The counterinstance lies at the permitted upper boundary r = q. A minimal two-round set exists at (q,r) = (4,4), so that boundary alone does not explain the obstruction. A minimal three-round set also exists in PG(2,5) at threshold 5.

The argument reduces minimality to closed deletion separators, encodes the resulting finite system, normalizes an infected noncollinear triple by collineations, and refutes the normalized formula by a checked reverse-unit-propagation certificate. The final Lean formula has 1,085 Boolean variables and 14,698 clauses. Its trimmed certificate contains 129,671 additions and 3,340,122 hint uses. The supplied formal endpoint is unconditional. This dossier independently recomputes the geometry, controls, formula bytes and RUP checks, without invoking Lean or a solver.

The journal version of record was not read. Positive author-side retention evidence and its limits are stated below; the result is asserted against the read arXiv statement, without a priority claim.

1. Definitions and source question

A finite projective plane of order q is an incidence structure with $q^2 + q + 1$ points and as many lines, q + 1 points on each line and q + 1 lines through each point, a unique joining line for distinct points, a unique intersection point for distinct lines, and a quadrangle. Write P for its point set. For an integer threshold r, define

$S_r(A) = A \cup \bigcup \{ l : |l \cap A| \geq r \}$, $A^0 = A$, $A^{n+1} = S_r(A^n)$.

All line thresholds in a round are evaluated in the old state A^n. Thus this is a synchronous process. The increasing sequence stabilizes; its stable value is $cl_r(A)$. A set percolates if $cl_r(A) = P$. For such a set, $t_r(A)$ is the least n with $A^n = P$. A percolating set is inclusion-minimal if no proper subset percolates. This is distinct from having minimum cardinality among percolating sets.

The source studies the smallest percolation time among inclusion-minimal sets. Its Proposition 17 excludes one-round minimal sets. Proposition 18 constructs three-round minimal sets when $4 \leq r \leq (q + 7)/3$. The ensuing question is whether two rounds can always be achieved throughout the larger stated threshold range.

2. Statement, scope and positive controls

Conjecture 20 in [1, p. 14] reads, verbatim apart from mathematical typesetting:

“For each r with $4 \le r \le q$ there exists a minimal percolating set A_r with $t_r(A_r) = 2$.”

The surrounding convention is an arbitrary finite projective plane of order q. There is no sufficiently-large-order condition, restriction to Desarguesian planes, or exclusion of r = q. The formal universal statement quantifies over q, every plane of order q, and every r satisfying the displayed inequalities. A single plane with no such set therefore refutes it.

Theorem. In PG(2,5), no inclusion-minimal percolating point set at threshold 5 has percolation time 2. Consequently the read formulation of Conjecture 20 is false.

Here $4 \le 5 \le 5$, so this is the allowed upper boundary, not an interior parameter. At q = 5 the hypothesis of Proposition 18 permits $r \le 4$; it does not cover the theorem’s threshold. Its three-round construction therefore does not already imply this obstruction. The theorem excludes every inclusion-minimal two-round set, of any cardinality, rather than only small or sampled configurations.

Positive controls use the coordinate ordering specified in Section 3. A mask m means that point i is initially infected exactly when bit i of m is set. Each listed set percolates at the stated time, and every one-point deletion has a nonfull stable closure. All five controls were freshly recomputed for this dossier.

q	r	Initial mask	Initial size	Successive infected counts	Time
4	4	1299211	11	11, 15, 21	2
5	4	70588740	10	10, 20, 31	2
7	4	6831815575732480	10	10, 26, 57	2
7	5	19422875086709251	15	15, 33, 57	2
5	5	241579737	16	16, 22, 25, 31	3

Full deletion closures are in recomputation.json. A proper subset is contained in a one-point deletion, so monotonicity upgrades these deletion checks to exact inclusion minimality. At order 4, field arithmetic is in $F_4 = F_2[X]/(X^2 + X + 1)$, rather than arithmetic modulo 4.

These controls distinguish the obstruction from a general failure at r = q, from an inability to percolate at threshold 5, and from a failure of the round counter to detect time 2. The three-round control at (5,5) is a concrete certificate, not a claim that Proposition 18 extends beyond its hypotheses. No theorem for other odd orders is asserted.

mathdb:2611.00058v1 07 Nov 2026

3. Geometry, stabilization and exact finite reduction

Represent points and line normals of PG(2,5) by the normalized vectors

(1,a,b) for a,b ∈ F_5, followed by (0,1,a) for a ∈ F_5, followed by (0,0,1).

Within the first block a varies before b. Indices begin at 0. Incidence is vanishing dot product modulo 5. This gives 31 points, 31 lines and degree 6 on either side. Recalculation verifies all 465 distinct point-pair joins and all 465 distinct line-pair meets. The formal plane checks the literal incidence table against the dot-product rule and verifies the plane axioms, including a quadrangle. Coordinate data are not accepted merely because a generator emitted them.

The map S_r is inflationary and monotone. On a finite set of N points, every nonfixed update strictly increases cardinality, and a fixed state remains fixed. Hence A^N is stable. In the present plane N = 31. This supplies a uniform finite closure bound for arbitrary initial sets, including deletion sets whose behavior must not be truncated to two rounds.

Separator lemma. If A is inclusion-minimal percolating and v ∈ A, then $B_v = cl_r(A \setminus \{v\})$ contains $A \setminus \{v\}$, is closed, and omits v.

Indeed, minimality makes B_v nonfull. If v belonged to B_v, it would contain A. Monotonicity, closedness and percolation of A would force $B_v = P$, a contradiction. Conversely, if A percolates and for every v ∈ A a closed B_v contains $A \setminus \{v\}$ but omits v, no deletion can percolate. Every proper subset lies in some deletion; hence A is inclusion-minimal. This converse permits arbitrary closed separators, not only canonical closures.

Introduce Boolean arrays A_p and X_p for the initial and first-round states, I_l and J_l for line activations, and B_{vp} for separators. The finite system requires:

- $I_l \leftrightarrow |l \cap A| \geq 5$, and $J_l \leftrightarrow |l \cap X| \geq 5$.
- $X = A \cup \bigcup\{l : I_l\}$; every point belongs to X or a J-line; $X \neq P$.
- If A_v holds, then B_v contains $A \setminus \{v\}$, is closed, and excludes v.

A target set gives a model by its actual activations and canonical deletion closures. A model conversely gives percolation by round 2 and failure at round 1. Inflationarity rules out time 0. Its separators prove inclusion minimality by the lemma. Thus existence of a model is equivalent to existence of the forbidden set. The Lean theorem finite5_iff_target proves this equivalence with the original unbounded-time and proper-subset definitions.

mathdb:2611.00058v1 07 Nov 2026

4. Truth-table CNF and collineation normalization

The formal encoding uses one clause for each forbidden local truth-table row. For support $z_1,\dots,z_k$ and a rejected row $b_1,\dots,b_k$, the blocking clause is $\bigvee_i (z_i \neq b_i)$. Such a clause excludes exactly that row. Activation equivalences, spread equivalences, final coverage and separator closedness are therefore exact local constraints. A six-point line is closed at threshold 5 precisely when its infected count is not exactly 5.

There are $4 \cdot 31 + 31^2 = 1{,}085$ variables. Their DIMACS numbering is $A_p \mapsto p + 1$, $X_p \mapsto p + 32$, $I_l \mapsto l + 63$, $J_l \mapsto l + 94$, $B_{vp} \mapsto 125 + 31v + p$. A forbidden bit true denotes a negative literal. The final clause families are:

Family	Clauses
Activation from A	1,984
Activation from X	1,984
First-round spread	3,968
Second-round coverage	31
Closed separators	5,766
Nonfull first round	1
Separator omission	31
Separator containment	930
Normalization units	3
Total	14,698

Without the units the total is 14,695. Fresh Python enumeration reproduces every DIMACS byte in the final normalized.cnf. The archived round-4 receipt separately records equality with the actual Lean generator’s export. The final packet retains the semantic generator and bridge; this dossier does not rerun that Lean export. The DIMACS SHA-256 is recorded in the reproduction materials.

Normalization is justified on witnesses. If A were contained in a line l, every other line would meet its closure in at most one point, below threshold 5, so infection could never leave l. Thus a percolating set contains a noncollinear triple. Its representative vectors form a basis. An invertible change of coordinates sends that triple to indices 0, 25 and 30; inverse-transpose action on line normals preserves incidence.

The Lean proof uses 86 explicitly checked collineations: 31 in family F, 30 in family G and 25 in family H. In succession these position the first point, the second while fixing the first, and the third while fixing the first two. Each table has proved inverse laws and incidence preservation. Transport commutes with step, hence with every round, and preserves fullness, percolation, exact time and minimality. The latter follows by transporting any hypothetical proper percolating subset back through the inverse permutation.

It is therefore sufficient to refute the formula with A_0, A_{25} and A_{30} true, the positive units 1, 26 and 31. This is a normalization theorem about geometric witnesses; adding units to an arbitrary formula is not presumed satisfiability-preserving.

mathdb:2611.00058v1 07 Nov 2026

5. Certificate and formal proof architecture

The normalized trimmed trace contains 129,671 additions, 3,340,122 hint uses and 1,868,145 literal occurrences in added clauses. It refers to 11,294 input clauses and ends with the empty clause at identifier 184281. The normalized-trimmed.json representation occupies 34,919,674 bytes. These are fresh counts from the frozen trace, rather than estimates from solver logs.

For a proposed clause C, reverse unit propagation (RUP) begins by making every literal in C false. Each positive hint names an input clause or an earlier proved clause. After removing currently false literals, a singleton forces its remaining literal; an empty remainder is a contradiction. A sequence reaching contradiction establishes C. Retaining earlier proved clauses when a trace requests deletion is sound, since all retained clauses are consequences of the original inputs.

The independent review [3] re-derived the finite model with its own encoding and checker. Its normalized formula has 8,312 clauses and its own certificate has 141,005 additions. It also checked the discovery encoding’s trimmed certificate. Those are independent computational cross-checks, not the final Lean trace. The discovery formula had 7,940 normalized clauses and 82,646 trimmed additions; its line activation constraints were one-way. The final truth-table formula uses full activation equivalences and different clause enumeration. The certificates are not interchangeable across these formulas.

In the final Lean implementation [4], clauses and falsifying assignments are pairs of natural-number bitmasks. A proved checker validates coverage and propagation choices. Its induction maintains that satisfied hint clauses cannot all be covered by a falsifying assignment. Literal-list semantics and mask semantics are proved equivalent. Every used input clause is proved from the actual generated CNF, and its numeric literal list is matched to the frozen DIMACS data.

For each addition, a data preparer supplies masks and propagation choices. The expression builder assembles ordinary proof terms: a reflexivity equality must reduce the actual Boolean RUP check in the kernel, and named earlier clause theorems discharge all hint premises. Shared reducible mask definitions keep imports within fixed memory limits. Neither Python, a solver response nor compiled native evaluation supplies a proof oracle.

The empty-clause theorem refutes normalizedCNF5. The geometric normalization bridge then gives noMinimalTwoRound5_unconditional and the delivered endpoint

Percolation.not_conjecture20_unconditional : ¬ Percolation.Conjecture20.

The supplied final replay covers 830 modules. Its endpoint axiom closure is exactly propext, Classical.choice and Quot.sound. The metadata audit includes the expression builder and emitted declarations, and the source scan checks forbidden proof shortcuts. Historical conditional endpoints from earlier rounds are preserved as development records; they are not presented as completed refutations. The final endpoint discharges the CNF-unsatisfiability assumption.

mathdb:2611.00058v1 07 Nov 2026

6. Replays, reproducibility and trust limits

The developer’s fresh archive replay records successful compilation and separate explicit-module checking of every ordered module: compiler time 3,213.01 s plus checker time 2,019.08 s, totaling 5,232.09 s, with peak RSS 2,738,672 KiB. Fresh aggregation of module-times.json reproduces these totals. They are summed successful process wall times, not an elapsed end-to-end measurement including driver overhead or failed development attempts.

The development record discloses earlier overlapping replay chunks and an import-memory failure in an unshared-mask implementation. Affected checks were repeated serially; the delivered shared-mask replay is the successful run above. These incidents are retained rather than erased. They must not be combined with successful-run timing or represented as a raised resource limit.

At this manuscript’s final evidence check, percolation-parent-final-r1/replay-result.json and registry-refresh.json were absent. The root replay and registry-refresh gates are therefore pending. A directory, source copy or running-progress log is not a completed root replay receipt. No Lean, lake, leanchecker or SAT/ILP solver was run in this manuscript lane.

The accompanying fresh isolated Python script reconstructs finite-field incidence, verifies degree/join/meet axioms, computes the five controls to stability, tests every deletion, regenerates the final DIMACS exactly, counts certificate statistics and checks every RUP addition. It also aggregates supplied module measurements and counts normalization modules. Mathematical calculations are recomputed; bibliographic identifiers, historical dates and reported external audit statuses are provenance claims bound to source files, not newly measured facts.

The final packet is reused externally at percolation-formal-r8/source-only.zip, byte-identically identified by SHA-256:

b2e78af5a5191a37f364f39a89272af5898ca6a7ecc59058a7661649837ea9fb.

It is not duplicated into this dossier. The archive is 64,586,499 bytes compressed and 323,843,462 bytes expanded. A lane-local safe-extraction test rejects absolute paths, traversal, symlinks, duplicate names and oversized members; it verifies the source manifest and the 830-entry module order, then removes the test tree. Only its documentation and module-order.txt are retained from that extraction. Exact raw MathDB bytes, source copies, frozen-input hashes, recomputation output and packaging receipts accompany the paper.

The trust boundary has several layers. The human-readable reduction must match the source question; the formal definitions and bridge must express that reduction; the delivered proof must be checked by the claimed kernel and toolchain. This lane verifies data and certificate arithmetic, while relying on supplied receipts for Lean replay and audits. Source scans alone prove neither semantics nor correctness. The independent review predates final formal completion and thus cannot be described as an independent full replay of the final endpoint.

mathdb:2611.00058v1 07 Nov 2026

7. Related work, limitations and significance

Gerbner, Keszegh, Mészáros, Patkós and Vizer [1] study line percolation on general finite projective planes, including percolating-set sizes, nonpercolating sets and time parameters. Their Section 5 supplies the precise question refuted here. The result separates minimum achievable time for minimal sets from the existence of short percolating configurations without minimality. It is a finite nonexistence theorem over all initial sets in one plane, not a statistical inference from a greedy search.

The later paper Triangle Percolation on the Grid [2] concerns lattice-triangle update families and cites the journal version of [1] as background. The cached review did not locate an exact refutation in its bounded literature and public-proof-registry searches. These searches are incomplete: repository descriptions, citation indexes and selected registry files cannot rule out unindexed or uninspected work. No worldwide priority claim is made, and this offline manuscript preparation performs no new live search.

The version of record, SIAM Journal on Discrete Mathematics 32(2) (2018), 864–881, DOI 10.1137/16M1087709, remains unread because the recorded publisher full-text request returned HTTP 403. Root ruling r2 [5] records positive author-side evidence: the authors' Hungarian Academy REAL repository deposit is the arXiv v1 PDF stating Conjecture 20; arXiv lists only v1; and the later citing paper cites the journal version without reporting a refutation. This supports retention cautiously, but does not establish that the typeset journal statement is unchanged. The refutation here is of the read arXiv formulation. The unread version of record remains a qualified limitation.

No short human proof of the nonexistence is known from the inspected materials. Complement-set observations and the exclusion of an initially full line do not alone force a third round. An obstruction for odd orders is conceivable, but is neither proved nor claimed. The theorem makes no assertion of a smallest counterexample, no classification at other thresholds, and no separate credit claim for the three-round control.

This AI-assisted campaign draft carries zero accepted credit; root alone decides acceptance. Formal completion, source fidelity, novelty evidence and acceptance remain distinct gates.

References and evidence

[1] Dániel Gerbner, Balázs Keszegh, Gábor Mészáros, Balázs Patkós and Máté Vizer. Line Percolation in Finite Projective Planes. arXiv:1608.00531v1. Journal metadata as above; DOI verified in cached publisher metadata and [2]. Read text: definitions p. 2; Propositions 17–18 pp. 11–13; Conjecture 20 p. 14.

[2] Igor Araujo, Bryce Frederickson, Robert A. Krueger, Bernard Lidický, Tyrrell B. McAllister, Florian Pfender, Sam Spiro and Eric Nathan Stucky. Triangle Percolation on the Grid. arXiv:2303.15402v2. The cached full text supplies its scope and reference to [1].

[3] Campaign independent adversarial review, percolation-review-r1, REVIEW.md and associated encoding, certificate and provenance receipts. Internal evidence, not an external publication.

[4] Campaign formalization rounds r1–r8; final source-only packet and replay/audit receipts, percolation-formal-r8. Lean 4.34.1; Std/Lean and local modules. No mathlib dependency is asserted.

[5] Root rulings percolation-root-ruling-r1.json and percolation-root-ruling-r2.json. Internal scope and retention adjudications; both record zero credit.

mathdb:2611.00058v1 07 Nov 2026