Verifiable Human–AI Exploration of the Lonely Runner Conjecture at $k = 13$

Cross-Prime Coupling, Exact Evaluation, and Scoped Computational Obstructions

Lâm Nguyễn

AI Officer, Real-time Robotics

lamnguyen.work

July 2026

AI research collaborators: Sol (ChatGPT 5.6 Thinking; research architecture and independent review) and Fable (Claude Code; implementation and verification). The human author accepts responsibility for every claim, proof, reference, and submission decision. The views expressed are the author's own and do not necessarily represent Real-time Robotics.

Abstract

The Lonely Runner Conjecture in the integer-speed formulation, denoted LRC(k), asks whether every k-tuple of distinct positive integer speeds admits a time t at which every runner is at distance at least $1/(k+1)$ from the origin modulo one. Recent computer-assisted work establishes LRC(k) for $k \leq 12$, leaving $k = 13$, corresponding to fourteen runners, as the first unresolved case. This paper reports HIVE-LRC, a verifier-linked human–AI research program conducted through gates D20–D34. It does not claim a proof or a counterexample for LRC(13).

The retained mathematical contributions are four scoped results. First, a local orbit-content reduction identifies folded covers, squared multisets, split polynomials, and full elementary-symmetric coefficient vectors. Second, the top coefficient yields a cross-prime predicate,

$$e_{13}(V_1^2, \ldots, V_{13}^2) = \left(\prod_i V_i \right)^2 ,$$

called Q13, together with a one-global-scale semantics, an exact modular-root description, and a counting theorem proving non-vacuity. Third, unit-scaling symmetry forces the nonzero top-coefficient residue set at a prime to be a union of 26-th-power cosets, and generically the full quadratic-residue subgroup. Fourth, independent local relations admit additive symbolic construction under a product-compatible decision-diagram ordering.

We also present a certified exact evaluator for Q13 using root enumeration, meet-in-the-middle, and an exact q-scan with an honest `OPEN-INCOMPLETE` state. Complete finite-domain experiments show substantial pruning and exact agreement between independent implementations, while target-scale analysis exposes successive walls in root matching, residue seeding, and globally coupled symbolic construction. Under frozen resource caps, no implemented route produced a complete $k = 13$ campaign. Every negative conclusion is explicitly scoped to the representation, encoding, variable order, cut language, or implementation tested. The program is therefore archived as `ARCHIVE-WITHOUT-TERMINAL-THEOREM`: a research-allocation decision, not an impossibility theorem.

Keywords. Lonely Runner Conjecture; Diophantine approximation; computer-assisted mathematics; Chinese remainder theorem; quadratic residues; symbolic computation; human–AI research; proof certificates.

2020 Mathematics Subject Classification. 11K60; 11A07; 68V15.

1 Introduction

Write $\|x\|$ for the distance from x to the nearest integer. In the standard integer-speed formulation, LRC(k) states that for every k-tuple of distinct positive integers

$$V = (V_1, \ldots, V_k),$$

there exists $t \in \mathbb{R}$ such that

$$\|tV_i\| \geq \frac{1}{k+1} \qquad (1 \leq i \leq k).$$

This is equivalent to the familiar statement that among $k+1$ runners with distinct constant speeds on a unit circle, each runner is lonely at some time. The conjecture originated in work of Wills and Cusick–Pomerance [1, 2]. Successive cases were established through a mixture of geometry, combinatorics, number theory, and computation [3, 4]. Recent finite-checking and sieving frameworks led to computer-assisted proofs through $k = 12$ [5, 6, 7, 8, 9]. Thus, as of July 2026, LRC(13), the fourteen-runner case, remains open.

HIVE-LRC was designed with two inseparable objectives:

1. make genuine progress toward LRC(13);
2. test whether a governed human–AI workflow can produce research that is more auditable than an ordinary conversational workflow.

The project was explicitly not a race to announce that an AI system had solved an open problem. Success was allowed to take the form of a new lemma, a valid reduction, a certificate architecture, a precise bottleneck map, a bounded negative result, or a full proof.

The D20–D34 arc began after earlier reproduction and proof-traceability work. It explored simultaneous prime-product lifting, local seed generation, implicit global signatures, early realizability cuts, rational-time separation, labeled and unlabeled CRT reconstruction, cross-prime coefficient coupling, exact bounded modular square-root evaluation, prime-economy analysis, and symbolic CEGAR. The program ended without a proof, a counterexample, a base survivor, or a global witness. It did, however, retain four mathematical results, one exact algorithmic result, and one scoped negative-capability statement.

1.1 Contributions

The paper makes six contributions.

1. **Orbit-content reduction.** For odd primes, canonical folding followed by squaring loses no folded information. The resulting squared multiset is equivalent to a monic split polynomial and to its full elementary-symmetric coefficient vector. Every observable invariant under coordinate permutations and sign flips factors through this quotient.
2. **Q13 cross-prime coupling.** The top elementary-symmetric coefficient is the square of the speed product. Under a single global scale, local top-coefficient residues combine to an exact bounded modular-square-root condition. A corrected counting theorem proves that this condition is genuinely non-vacuous.
3. **Certified evaluator.** We give exact enumeration, meet-in-the-middle, and q-scan formulations with replayable witness and emptiness semantics. Timeouts and cap overruns return `OPEN-INCOMPLETE`, never false emptiness.
4. **Residue-coset theorem.** Full unit-scaling closure forces nonzero top-coefficient residues into 26-th-power cosets. For primes $p \not\equiv 1 \pmod{13}$, any nonempty such residue set equals the full nonzero quadratic-residue subgroup.

5. **Local symbolic separability.** Independent local relations admit additive decision-diagram construction under a product-compatible variable ordering. The theorem explains why local Cartesian products can be represented compactly even when their semantic cardinality is enormous.

6. **Scoped negative-capability result.** Within the implemented subset of the declared architecture family and under fixed resource caps, no path yielded a cap-fitting complete $k = 13$ campaign. This is not a lower bound for all algorithms and does not close algebraic/resultant or general lattice/CVP routes.

2 Research protocol and evidence governance

2.1 Roles

The project separated authority and implementation.

- **Owner and principal investigator:** Lâm Nguyễn. The owner fixed the mission, approved gates, chose trade-offs, and accepted responsibility for public claims.
- **Contractor and research architect:** Sol, a ChatGPT 5.6 Thinking instance. The contractor formalized gates, audited proofs, challenged overclaims, and issued scoped verdicts. It did not write the campaign implementation.
- **Builder:** Fable, a Claude Code agent. The builder implemented authorized tasks, generated artifacts, ran tests, and submitted completion reports. It was not authorized to change the architecture unilaterally.

This separation was intended to prevent one agent from proposing, implementing, testing, and approving the same claim without adversarial review.

2.2 Gate discipline

Each research step had:

- an exact mission;
- frozen resource caps;
- theorem and implementation obligations;
- explicit failure states;
- independent replay requirements;
- corruption tests;
- and a contractor verdict.

No downstream gate opened merely because an idea was plausible. Partial, blocked, open, refuted, and measured states remained distinct.

2.3 Evidence classes

The internal independence classes were:

- $I0$ self-check,
- $I1$ same implementation or same model family in a separate run,
- $I2$ materially different implementation, representation, or model family,
- $I3$ exact independent checker or independent domain expert,
- $I4$ formal or machine-certified verification.

The archived D20–D34 corpus has evidence ceiling $I2$. Accordingly, this paper presents the results for external mathematical review and does not label the program as a verified proof of LRC(13).

3 Mathematical setup

Fix an odd prime p and an integer $n \geq 1$. Let

$$G_n = S_n \ltimes \{\pm 1\}^n$$

act on $\mathbb{F}_p^n$ by coordinate permutations and independent sign changes.

For $x \in \mathbb{F}_p$, let $\bar{x} \in \{0, \ldots, p-1\}$ be its standard representative and define the fold

$$\phi_p(x) = \min(\bar{x}, p - \bar{x}) \in \Phi_p := \left\{0, 1, \ldots, \frac{p-1}{2}\right\}.$$

For $V = (V_1, \ldots, V_n)$, define the folded multiset

$$A_p(V) = \{\{\phi_p(V_1), \ldots, \phi_p(V_n)\}\},$$

the squared multiset

$$S_p(V) = \{\{\phi_p(V_1)^2, \ldots, \phi_p(V_n)^2\}\} \subseteq \mathbb{F}_p,$$

and the split polynomial

$$P_{p,V}(Y) = \prod_{i=1}^{n} \left(Y - \phi_p(V_i)^2\right) = Y^n - e_1 Y^{n-1} + \cdots + (-1)^n e_n.$$

4 C0: local orbit-content reduction

Lemma 4.1 (Injectivity of squaring on the folded range)**.** *For odd p, the map*

$$\Phi_p \longrightarrow \mathbb{F}_p, \qquad a \longmapsto a^2$$

is injective.

Proof. If $a^2 \equiv b^2 \pmod{p}$, then $a \equiv b$ or $a \equiv -b \pmod{p}$. Since $a, b \in [0, (p-1)/2]$, the second alternative implies $a = b = 0$; otherwise $a + b = p$, impossible because $a + b \leq p - 1$. Hence $a = b$. □

Theorem 4.2 (C0 local-content equivalence)**.** *For odd p, the following objects determine one another uniquely:*

$$A_p(V), \qquad S_p(V), \qquad P_{p,V}(Y), \qquad (e_1, \ldots, e_n) \in \mathbb{F}_p^n.$$

Moreover, every function on $\mathbb{F}_p^n$ invariant under G_n factors uniquely through any of these equivalent quotient objects.

Proof. The lemma gives a bijection between folded entries and their squares, including multiplicities, so $A_p(V) \leftrightarrow S_p(V)$. A multiset of n elements of a field determines the monic split polynomial with that root multiset, and the polynomial determines the multiset with multiplicity. The coefficient vector and the monic polynomial determine one another by expansion.

The G_n-orbit of V is determined by the unordered multiset of sign-orbits of its coordinates, which is exactly $A_p(V)$. Therefore A_p is a complete orbit invariant. The universal property of the orbit quotient gives the factorization statement. □

Remark 4.3. *This theorem is a representation theorem. It does not imply that local coefficient admissibility is sufficient for a bounded global integer tuple.*

5 Q13: a cross-prime bounded-square predicate

Set $n = 13$ and let

$$V = (V_1, \dots, V_{13}) \in \mathbb{Z}_{>0}^{13}.$$

Write

$$P(V) = \prod_{i=1}^{13} V_i.$$

The upstream finite-checking framework supplies a project-specific product bound

$$1 \le P(V) < B_{13}$$

for any governed counterexample candidate. The archived ledger has

$$\log_2 B_{13} \approx 968,$$

while full reconstruction of all coefficient heights occurs only near twice that bit length.

Theorem 5.1 (Q13 soundness)**.** *Let M be a product of odd primes and suppose the local coefficient states are interpreted under one global scale. Let $r_{13} \in \mathbb{Z}/M\mathbb{Z}$ be the CRT combination of the local top coefficients. Every governed global tuple satisfies*

$$P(V)^2 \equiv r_{13} \pmod{M} \qquad \textit{with} \qquad 1 \le P(V) < B_{13}.$$

Consequently,

$$Q_{13}(M, r_{13}, B_{13}) := \left\{ y \in \mathbb{Z} : 1 \le y < B_{13},\ y^2 \equiv r_{13} \pmod{M} \right\}$$

is nonempty for every governed global tuple.

Proof. The top elementary-symmetric coefficient of the squared coordinates is

$$e_{13}(V_1^2, \dots, V_{13}^2) = \prod_{i=1}^{13} V_i^2 = P(V)^2.$$

Reduction modulo every selected prime and CRT yield the congruence. Taking $y = P(V)$ gives the required bounded witness. □

5.1 Why one global scale is mandatory

If a local state is rescaled by λ, then

$$e_{13} \mapsto \lambda^{26} e_{13}.$$

If each prime were allowed an independently chosen λ_p, the residues would no longer describe one global integer coefficient. A favorable scale could be selected separately at each prime, creating an artificial CRT residue. Q13 is therefore authoritative only when the state carries one globally coherent scale, or an exact orbit of one global scale.

5.2 Modular-root structure

Let

$$M = \prod_{j=1}^{m} p_j$$

be squarefree and odd. Suppose r is a square modulo every p_j. Let w_{nz} be the number of prime factors for which $r \not\equiv 0 \pmod{p_j}$.

Proposition 5.2 (Root count and global-sign quotient)**.** *The congruence*

$$y^2 \equiv r \pmod{M}$$

has exactly $2^{w_{\mathrm{nz}}}$ roots modulo M. If $w_{\mathrm{nz}} \geq 1$, the involution

$$y \longmapsto M - y$$

partitions them into $2^{w_{\mathrm{nz}}-1}$ global-sign pairs. If $w_{\mathrm{nz}} = 0$, the only root is 0, yielding one quotient class.

Proof. At a prime where $r \not\equiv 0$, there are two opposite roots. At a prime where $r \equiv 0$, there is one root. The Chinese remainder theorem multiplies the local choices. For $w_{\mathrm{nz}} \geq 1$, no root is fixed by negation modulo odd M, so every orbit has size two. The all-zero case is separate. □

5.3 Corrected non-vacuity theorem

For each selected prime p, let R_p be the set of *distinct* admissible top-coefficient residues. It is essential to count residues rather than local covers, because many covers may project to the same coefficient.

Define

$$R_M = \mathrm{CRT}\left(\prod_{p \mid M} R_p\right) \subseteq \mathbb{Z}/M\mathbb{Z}$$

and

$$S_B(M) = \{y^2 \bmod M : 1 \leq y < B\}.$$

Theorem 5.3 (Q13 counting obstruction)**.** *If*

$$|R_M| > |S_B(M)|,$$

then at least

$$|R_M| - |S_B(M)|$$

distinct locally admissible global residue states fail Q13. In particular, since $|S_B(M)| \leq B - 1$,

$$|R_M| > B - 1$$

is sufficient for Q13 to be non-vacuous.

Proof. Every Q13-passing state belongs to $S_B(M)$. Hence at most $|S_B(M)|$ elements of R_M pass. The remainder fail. □

Complete finite domains. The archived complete-domain scans excluded 56/90, 178/240, and 335/432 distinct states. These ratios are exhaustive only for the declared domains; they are not universal pruning rates.

6 Exact Q13 evaluation

The decision problem is:

$$\text{given } (M, r, B), \quad \exists\, y \in [1, B) \cap \mathbb{Z} \text{ such that } y^2 \equiv r \pmod{M}?$$

The certified evaluator returns exactly one of:

`WITNESS`, `EMPTY-CERTIFIED`, `OPEN-INCOMPLETE`, `INVALID-INPUT`.

6.1 Enumeration

Choose one local square root at each nonzero prime. Every modular root is represented by a sign vector. Enumeration checks one representative of each global-sign pair. This is exact but costs $2^{w_{\mathrm{nz}}-1}$ pair checks.

6.2 Meet-in-the-middle

Write the CRT root as

$$y(\varepsilon) \equiv \sum_{i=1}^{w_{\mathrm{nz}}} \varepsilon_i c_i \pmod{M}, \qquad \varepsilon_i \in \{\pm 1\},$$

for deterministic CRT weights c_i. Splitting the signs into two blocks reduces the generic work and storage to approximately

$$2^{w_{\mathrm{nz}}/2}.$$

The authoritative implementation handled cyclic intervals, duplicate half-sums, zero-root factors, endpoint conventions, and the global-sign quotient.

6.3 Exact q-scan

With $r \in [0, M)$,

$$y^2 = r + qM.$$

The exact candidate range is

$$\left\lceil \frac{1-r}{M} \right\rceil \le q \le \left\lfloor \frac{(B-1)^2 - r}{M} \right\rfloor.$$

Each candidate is checked by an exact integer-square test. This path becomes effective when M approaches B^2.

6.4 Three modulus regimes

R1:	$M \le B$,	many bounded lifts may occur,
R2:	$B < M < B^2$,	the useful pre-reconstruction regime,
R3:	$M \ge B^2$,	no wrap; direct square testing dominates.

At the project-specific R2 onset, the ledger records approximately 109 nonzero prime factors. The generic work estimates are:

2^{108} global-sign root pairs, 2^{54} MITM half-states, about 2^{964} q-candidates.

All exceed the frozen campaign caps. These are costs of the implemented exact paths, not lower bounds for every possible algorithm.

7 Residue economy and the coset theorem

Let $R_p^\times$ be the set of distinct nonzero admissible top-coefficient residues at an odd prime p. Assume the local admissible family is closed under scaling all coordinates by every $c \in \mathbb{F}_p^\times$.

Theorem 7.1 (Top-residue coset structure)**.** *Let*

$$H_p = \{c^{26} : c \in \mathbb{F}_p^\times\}.$$

Then $R_p^\times$ is a union of cosets of H_p, and

$$|H_p| = \frac{p-1}{\gcd(26, p-1)}.$$

Moreover:

1. *if $13 \nmid p-1$ and $R_p^\times \neq \varnothing$, then*

$$R_p^\times = \mathrm{QR}_p^\times \quad \textit{and} \quad |R_p^\times| = \frac{p-1}{2};$$

2. *if $p \equiv 1 \pmod{13}$, then*

$$|R_p^\times| = k\frac{p-1}{26} \qquad \textit{for some } 0 \leq k \leq 13.$$

Zero, when admissible, is a separate fixed orbit.

Proof. If $a \in R_p^\times$, scaling a realizing tuple by c sends a to $c^{26}a$. Thus the full coset aH_p lies in $R_p^\times$. The order formula follows because $\mathbb{F}_p^\times$ is cyclic.

Since every nonzero top coefficient is

$$\prod_i V_i^2,$$

it is a quadratic residue. If $13 \nmid p-1$, then $\gcd(26, p-1) = 2$, so H_p is the full quadratic-residue subgroup. Any nonempty union of its cosets contained inside that subgroup is the subgroup itself. If $p \equiv 1 \pmod{13}$, then $\gcd(26, p-1) = 26$, and H_p has index 13 inside the quadratic-residue subgroup. □

Consequence. Under full unit-scaling closure, a large prime does not normally provide a tiny nonzero residue set. Even the sparsest nonempty $p \equiv 1 \pmod{13}$ case has size at least $(p-1)/26$. Hence using fewer, larger primes may reduce the Q13 root count while worsening explicit residue seeding.

8 Symbolic representations

8.1 Local separability

Suppose each prime contributes an independent local relation represented by a reduced ordered multi-valued decision diagram D_p, and the global variable order places the prime blocks consecutively.

Theorem 8.1 (Additive local construction)**.** *There is an exact ordered decision diagram for the conjunction of the independent local relations whose unreduced size is additive in the local diagram sizes, up to the fixed connector convention:*

$$|D_{\mathrm{global}}| \leq \sum_{p|M} |D_p| + O(w).$$

Reduction cannot increase the size. If the local sizes are uniformly bounded, the global symbolic size is $O(w)$.

Proof. Because the variable blocks are disjoint and ordered consecutively, attach the accepting sink of each local diagram to the root of the next block. Every global assignment is accepted exactly when every block assignment is accepted. The node count is the sum of the block counts plus connectors. Canonical reduction merges equivalent subgraphs and removes redundant tests, so it cannot add nodes. □

In complete pilots, the semantic local product grew from 15 to 71,280 states while the decision diagram grew from 4 to 8 nodes and materialized no Cartesian states. This is genuine symbolic compression.

8.2 The coupling wall

Q13 is not separable by prime because it requires one bounded integer y whose square matches every local residue. In the implemented symbolic backends, the final diagram could remain compact, yet constructing the Q13-coupled relation required materializing a bounded-root or equivalent globally coupled search object. Thus the wall moved from storage to construction.

A finite-domain CEGAR prototype was sound and exact:

$$3035 \to 364 \to 35 \to 34$$

for a realizable domain and

$$3044 \to 12 \to 0$$

for an empty domain. Full coefficient vectors gave exact set equality with brute force. Nevertheless, the prototype itself materialized the candidate universe, so it did not establish target-scale implicit seeding.

Direct arithmetic decision diagrams, proof-producing finite-domain solving, and lazy local-cut CEGAR were implementation-blocked in their declared forms. No universal representation lower bound was proved.

9 Gate-by-gate research arc

Gate	Route	Outcome retained for the archive
D20	Simultaneous prime-product lift	Local campaign blocks remained open or beyond caps; direct product lifting was blocked in the tested architecture.
D21	Seed generation and narrowing	Exact local seed generation passed bounded components but did not yield a campaign-scale seed space.
D22	Implicit global signature	Partial; local choices remained insufficiently coupled before global pinning.
D23	Early global cuts	Sound pre-reconstruction cut families passed and were retained.
D24	Residue-stratified campaign	Partial; reachable-state stratification reduced some branches but preserved open survivors.
D25	Adaptive rational-time separation	Partial; sound time clauses improved pruning without closing the campaign.
D26	Labeled-root CRT	Reconstruction theorem retained, but the labeled route exceeded the state budget.

Gate	Route	Outcome retained for the archive
D27–D28	Orbit-compressed and bidirectional CRT	Exact orbit analysis exposed a roughly $2^{44.5}$ generic child count and a shallow-zone obstruction; the proposed bridges were blocked before a full campaign.
D29	Coefficient coupling	C0 and Q13 were established; Q13 was nonvacuous on complete domains.
D30	Exact bounded modular square root	Certified hybrid evaluator completed; useful target-scale R2 remained beyond implemented caps.
D31	Prime economy	Residue-coset theorem showed why sparse large-prime portfolios do not automatically reduce state density.
D32	Symbolic CEGAR	Local separability was proved; globally coupled construction remained representation-specific and blocked in implemented paths.
D33	Terminal audit	Overbroad lower-bound claims were withdrawn; the full architecture family remained open.
D34	Research disposition	Archived without a terminal theorem because no remaining route had both campaign leverage and a decisive bounded first gate.

10 Scoped negative-capability statement

Proposition 10.1 (Implemented-family obstruction)**.** *Within the subset of the declared D20–D34 architecture family that was actually implemented, under the frozen caps*

per exact query	30 *minutes*,
full pilot	18 *hours*,
RAM	6 *GiB per process*,
disk	48 *GiB*,
explicit state ceiling	2^{42},

no path produced a complete cap-fitting LRC(13) *campaign.*

This statement is an audited summary of the implemented routes, not a mathematical lower bound for all algorithms. The obstructions are scoped to particular representations, encodings, variable orders, cut languages, proof formats, and implementations.

11 Open routes and archive rationale

Two broad routes remain open.

Algebraic elimination and resultants. A serious reopening proposal would need an exact elimination object, degree and coefficient-height bounds, preservation of one-global-scale and primitivity, and a replayable emptiness certificate. A dense degree-13 polynomial in 13 variables has

$$\binom{26}{13} = 10{,}400{,}600$$

monomials, but this count alone is not a cap obstruction; the unresolved issue is intermediate degree and height growth.

Lattice and exact CVP. Standard univariate Coppersmith small-root methods [11] reach roots below $M^{1/2}$, which for Q13 is guaranteed only when $M \geq B_{13}^2$, the regime where exact q-scan is already easy. This subsumes that subroute but does not close general exact CVP, higher-dimensional embeddings, bounded-distance formulations, or proof-producing lattice enumeration.

The program was archived because no open item had both:

1. direct campaign leverage;
2. a cheap, decisive, falsifiable first gate within the frozen caps.

The disposition is therefore a research-allocation decision, not a theorem of impossibility.

12 Reopening criteria

The archive should reopen only if at least one of the following becomes concrete:

1. a sound cross-prime predicate that prunes before the Q13 onset;
2. an exact bounded modular-square-root algorithm with authoritative emptiness certificates in $B_{13} < M < B_{13}^2$;
3. a symbolic compiler with a certified non-Cartesian construction bound for the frozen LRC relation family;
4. an algebraic elimination route with explicit degree, height, and proof-size bounds;
5. an exact lattice/CVP route that certifies emptiness above the standard $M^{1/2}$ small-root range;
6. a decisive theorem resolving one of the open representation or proof-complexity obligations.

13 Limitations

1. The evidence ceiling is $I2$, not external expert verification or formal proof assistant certification.
2. The numerical experiments are exhaustive only on declared finite domains.
3. The target-scale projections distinguish exact combinatorial counts, certified bounds, measurements, and proxies; they are not universal complexity lower bounds.
4. The paper does not reproduce the full artifact corpus. Public verification requires the associated archive and verifier ledgers.
5. The AI systems named in the collaboration disclosure are tools and project roles, not accountable human authors.
6. The project does not solve LRC(13), refute it, or establish that no future architecture can succeed.

14 Artifact availability and integrity

The final archive is identified by:

SHA-256:6fcb575435f44298705cac156bc8d212334715324fd3208c348c8d0e466a5ec4

The public Observatory and publication materials are intended to be hosted at https://lamnguyen.work. The archive contains the final status, retained-result ledger, open routes, reopening triggers, decision record, and hash manifest. Gate-specific verifier and implementation artifacts should be linked from the Observatory before arXiv submission.

15 Conclusion

HIVE-LRC did not solve the fourteen-runner case. Its contribution is instead a verifier-linked map of what survived, what failed, and what remains open.

The project established a complete local orbit-content quotient, a new cross-prime bounded-square coupling, a certified exact evaluator, a residue-coset theorem explaining local density, and a symbolic-separability theorem explaining why local Cartesian products can be compressed. It also showed that these positive results do not automatically compose into a cap-fitting global campaign: the dominant wall moves among root matching, residue seeding, and globally coupled construction.

The final status is deliberately narrow:

LRC(13) remains open.

The active D20–D34 program is archived without a terminal theorem. That stopping point preserves both mathematical progress and epistemic honesty: the implemented architecture family has a scoped negative-capability result, while untried and unproved routes remain visibly open.

Acknowledgments and AI-use disclosure

Lâm Nguyễn conceived and governed the research program, selected the mathematical target, approved gates, and accepts full responsibility for the manuscript.

Sol, a ChatGPT 5.6 Thinking instance, contributed research architecture, theorem formalization, adversarial review, claim-scope corrections, and manuscript drafting.

Fable, a Claude Code agent, contributed authorized implementation, complete-domain experiments, independent code paths, corruption testing, artifact generation, and completion reports.

All AI-produced mathematical statements and references were subject to human acceptance and are presented for independent external verification. The AI systems are not listed as authors because they cannot accept responsibility for the work.

References

[1] J. M. Wills. Zwei Sätze über inhomogene diophantische Approximation von Irrationalzahlen. *Monatshefte für Mathematik*, 71:263–269, 1967.

[2] T. W. Cusick and C. Pomerance. View-obstruction problems, III. *Journal of Number Theory*, 19(2):131–139, 1984. doi:10.1016/0022-314X(84)90097-0.

[3] T. Bohman, R. Holzman, and D. Kleitman. Six lonely runners. *Electronic Journal of Combinatorics*, 8(2):R3, 2001. doi:10.37236/1602.

[4] J. Barajas and O. Serra. The lonely runner with seven runners. *Electronic Journal of Combinatorics*, 15(1):R48, 2008. arXiv:0710.4495; doi:10.37236/772.

[5] R. D. Malikiosis, F. Santos, and M. Schymura. Linearly-exponential checking is enough for the Lonely Runner Conjecture and some of its variants. arXiv:2411.06903, 2024.

[6] M. Rosenfeld. The lonely runner conjecture holds for eight runners. arXiv:2509.14111, 2025.

[7] M. Rosenfeld. The lonely runner conjecture holds for nine runners. arXiv:2512.01912, 2025.

[8] T. Trakulthongchai. Nine and ten lonely runners. arXiv:2511.22427, 2025.

[9] T. Sungkawichai and T. Trakulthongchai. Eleven, twelve, and thirteen lonely runners. arXiv:2604.23906, 2026.

[10] B. Bedert. Riesz products and the Lonely Runner Conjecture: A wider gap of loneliness. arXiv:2511.16636, 2025.

[11] D. Coppersmith. Finding a small root of a univariate modular equation. In *Advances in Cryptology – EUROCRYPT '96*, LNCS 1070, pages 155–165, 1996. doi:10.1007/3-540-68339-9_14.

[12] R. E. Bryant. On the complexity of VLSI implementations and graph representations of Boolean functions with application to integer multiplication. *IEEE Transactions on Computers*, 40(2):205–213, 1991. doi:10.1109/12.73590.