THE CHOW–RIMANIĆ COVERING NUMBER $C_k(2)$: EXACT VALUES BY DEGREE WINDOW, AND THE OPTIMAL FAMILIES

KOREA SUPERINTELLIGENCE LABS

Abstract. Chow and Rimanić introduced a function-field analogue of the lonely runner conjecture: for a finite family $\mathcal{F}$ of nonzero polynomials over $\mathbb{F}_q$ one sets $\delta(\mathcal{F}) = \sup_\alpha \min_{f\in\mathcal{F}} \|\alpha f\|$ and lets $C_k(q)$ be the least size of a family with $\delta(\mathcal{F}) < q^{-k}$; they conjectured $C_k(q) = Q_k(q) = 1 + q + \cdots + q^k$. Hu has recently refuted this at $(q,k) = (2,3)$ with thirteen polynomials over $\mathbb{F}_2$ found by a cross-entropy search, giving $9 \le C_3(2) \le 13 < 15$, and asked for the exact value of $C_3(2)$. We settle that question inside the degree window in which the counterexample lives: no twelve nonzero polynomials over $\mathbb{F}_2$ of degree at most four have loneliness less than 2^{-3}, so the thirteen are optimal among all speeds of degree at most four. We then classify the optimal families in that window: there are exactly three, they carry the same covering-multiplicity distribution, and coefficient reversal carries the published family to the second and fixes the third. For $k = 2$ we compute the value restricted to degrees at most four and at most five and find the conjectured 7 in both — past the reach of the Chow–Rimanić small-degree theorem, which at $q = 2$, $k = 2$ reaches only degree two. Every lower bound is an exhaustive search whose branching rule is proved complete rather than assumed, and every result proved here is machine-checked in Lean 4; the two results taken from the literature are marked as such.

1. Introduction

The function-field runner problem. Chow and Rimanić [2] introduced a non-Archimedean, function-field analogue of the lonely runner conjecture, as follows. Fix a prime power q and set $K_\infty = \mathbb{F}_q((T^{-1}))$ and $\mathbb{T} = T^{-1}\mathbb{F}_q[[T^{-1}]]$. Every $\beta \in K_\infty$ splits uniquely as $\beta = [\beta] + \{\beta\}$ with $[\beta] \in \mathbb{F}_q[T]$ and $\{\beta\} \in \mathbb{T}$; put $\|\beta\| = |\{\beta\}|$, the absolute value being normalised by $|\beta| = q^m$ for β of degree m. For a finite family $\mathcal{F} \subset \mathbb{F}_q[T] \setminus \{0\}$ of *speeds* the *loneliness* is

$$\delta(\mathcal{F}) := \sup_{\alpha\in\mathbb{T}} \min_{f\in\mathcal{F}} \|\alpha f\|, \tag{1}$$

and, following Hu [1], we encode the problem by

$$C_k(q) := \min\big\{\, |\mathcal{F}| :\ \mathcal{F} \subset \mathbb{F}_q[T]\setminus\{0\},\ \delta(\mathcal{F}) < q^{-k} \,\big\}, \qquad Q_k(q) := 1 + q + \cdots + q^k. \tag{2}$$

Multiplying a speed by a nonzero scalar changes neither δ nor C_k, so speeds may be taken monic; over $\mathbb{F}_2$ every nonzero polynomial is monic already. The family of all monic polynomials of degree at most k has $Q_k(q)$ members and loneliness at most $q^{-(k+1)}$, while the union-bound theorem [2, Theorem 1.5] shows that no family of size q^k or less has loneliness below q^{-k}; hence

$$q^k + 1 \ \le\ C_k(q) \ \le\ Q_k(q). \tag{3}$$

Hu [1] traces finite-field variants of the runner problem back to Czerwiński and Grytczuk [4]. Chow and Rimanić conjectured that the upper bound is the truth.

Conjecture 1.1 (Chow–Rimanić [2, Conjecture 1.4])**.** $C_k(q) = Q_k(q)$ for every prime power q and every $k \ge 1$; equivalently, every family with $|\mathcal{F}| < Q_k(q)$ satisfies $\delta(\mathcal{F}) \ge q^{-k}$.

They proved the conjecture in a small-degree range and obtained $C_2(q) > q^2 + 0.4877q - C$ for a universal constant C in the first nontrivial case [2, Theorem 1.7], noting there that carefully following their proof would give $C = 6$; a recent survey of the lonely runner conjecture [3] records the function-field conjecture and these results, and no computed value of $C_k(q)$. The small-degree theorem is the benchmark against which everything below should be read; it is the theorem headed

"Small degrees", [2, Theorem 1.6], restated by Hu as [1, Theorem 5.4]. Let $N(m,q)$ be the number of monic irreducible polynomials of degree m over $\mathbb{F}_q$, the notation of [2].

Theorem 1.2 (Chow–Rimanić, "Small degrees" [2, Theorem 1.6])**.** *Let $k>1$ and let q, D satisfy*

$$\frac{N(k+1,q)}{q^k+q^{k-1}+\cdots+q} > \Big\lfloor \frac{D}{k+1} \Big\rfloor .$$

Then every family of nonzero polynomials of degree at most D and of size less than $Q_k(q)$ has $\delta(\mathcal{F}) \geq q^{-k}$.

Hu's counterexample and the question it leaves. Hu [1, Theorem 1.2] refuted Conjecture 1.1 at $(q,k)=(2,3)$. His family

$$\begin{aligned}\mathcal{F}_* = \big\{ &1,\ T^2{+}T,\ T^2{+}T{+}1,\ T^3,\ T^3{+}T^2{+}1,\ T^3{+}T{+}1,\ T^3{+}T^2{+}T{+}1,\\ &T^4{+}T^2,\ T^4{+}T,\ T^4{+}T^3{+}1,\ T^4{+}T^2{+}1,\ T^4{+}T{+}1,\ T^4{+}T^3{+}T^2{+}T{+}1 \big\}\end{aligned} \tag{4}$$

of thirteen polynomials of degree at most four satisfies $\delta(\mathcal{F}_*) = 2^{-4}$, so

$$9 \ \leq\ C_3(2) \ \leq\ 13 \ <\ 15 \ =\ Q_3(2), \tag{5}$$

the lower bound being (3). The family was found computationally by a cross-entropy search in the spirit of Wagner's method [5] and is verified afterwards by checking the 2^7 vectors of the relevant coefficient space; no minimality is claimed for it, and (5) leaves a gap of five possible values. Hu asks, verbatim [1, Problem 8.2]:

> **Problem 8.2.** Determine $C_3(2)$ exactly and classify all failures of the projective threshold for small pairs (q,k).

and adds: "The finite problem may be formulated as set cover on the collection of partial-circulant kernels." That is the problem this paper attacks, one degree window at a time.

Degree windows. For $D \geq 0$ put

$$C_k^{(D)}(q) := \min\big\{\, |\mathcal{F}| :\ \mathcal{F} \subset \mathbb{F}_q[T]\setminus\{0\},\ \deg f \leq D\ \ (f\in\mathcal{F}),\ \delta(\mathcal{F}) < q^{-k} \,\big\} \in \mathbb{N}\cup\{\infty\}, \tag{6}$$

with $\min\emptyset = \infty$. The sequence $D \mapsto C_k^{(D)}(q)$ is non-increasing, it is finite and at most $Q_k(q)$ as soon as $D \geq k$, and $C_k(q) = \min_D C_k^{(D)}(q)$, the minimum being attained because the sequence is non-increasing and bounded below by q^k+1. So each window value is an upper bound for $C_k(q)$, and a lower bound for it only inside the window. Theorem 1.2 is exactly a statement about windows: at $q=2$ its hypothesis holds for $D\leq 2$ when $k=2$, for $D\leq 3$ when $k=3$, and for $D\leq 4$ when $k=4$. Hu's counterexample lives at $k=3$, $D=4$, the first window at $k=3$ that Theorem 1.2 does not reach.

Results. Everything below is at $q=2$. Our first result determines the value in the window the counterexample lives in, and so shows that the thirteen polynomials found by cross-entropy search cannot be improved without leaving degree four.

Theorem 1.3 (Theorem 5.1)**.** $C_3^{(4)}(2) = 13$. *That is: no family of at most twelve nonzero polynomials over $\mathbb{F}_2$ of degree at most four has loneliness less than 2^{-3}, while Hu's $\mathcal{F}_*$ does.*

The second result classifies the optimal families in that window. Write ρ for coefficient reversal $f(T) \mapsto T^4 f(1/T)$ on polynomials of degree at most four.

Theorem 1.4 (Theorems 6.1 and 6.2)**.** *There are exactly three families of thirteen nonzero polynomials of degree at most four over $\mathbb{F}_2$ with loneliness less than 2^{-3}: Hu's $\mathcal{F}_*$, its reversal $\mathcal{F}_{\mathrm{rev}} = \rho(\mathcal{F}_*)$, and one further family $\mathcal{F}_\circ$ with $\rho(\mathcal{F}_\circ) = \mathcal{F}_\circ$. All three have the covering-multiplicity distribution that Hu computes for $\mathcal{F}_*$ alone.*

The third result computes the value in the neighbouring case $k=2$, in two windows that Theorem 1.2 does not reach; here the Chow–Rimanić value survives.

Theorem 1.5 (Theorem 7.1)**.** $C_2^{(4)}(2) = C_2^{(5)}(2) = 7 = Q_2(2)$.

Theorem 1.2 reaches only $D \le 2$ at $q = 2$, $k = 2$, and the unconditional Chow–Rimanić bound [2, Theorem 1.7] is vacuous at $q = 2$ — with the constant $C = 6$ they record, its hypothesis is empty there — so Theorem 1.5 is not implied by [2]; it confirms Conjecture 1.1 rather than refuting it anywhere. Table 1 collects the windows we have computed, together with those that a published result already decides and that we recompute as controls (Proposition 7.2).

k	D	$C_k^{(D)}(2)$	$Q_k(2)$	status
1	≤ 4	3	3	control; $C_1(q) = q + 1$ follows from (3)
2	≤ 4	7	7	Theorem 1.5
2	≤ 5	7	7	Theorem 1.5
3	≤ 2	∞	15	no family covers at all (Proposition 7.3)
3	≤ 3	15	15	control; Theorem 1.2
3	≤ 4	**13**	15	Hu's family; optimality is Theorem 1.3
3	≤ 5	≤ 13	15	$= 13$ by an unverified computation (Remark 5.2)
4	≤ 4	31	31	control; Theorem 1.2

TABLE 1. The degree-windowed covering number at $q = 2$. Widening the window can only lower the value, and $C_k(2) = \min_D C_k^{(D)}(2)$.

What is not claimed. The value of $C_3(2)$ itself is *not* determined here. Theorem 1.3 pins the minimum only inside deg ≤ 4; a family of fewer than thirteen speeds could still exist at higher degree, and Hu's question remains open. What Theorem 1.3 does settle is that the counterexample is not improvable by a better search in its own degree range, and Theorem 1.4 says exactly what the alternatives in that range are.

Method. Every upper bound in this paper is an explicit family, checked by evaluating k linear forms at every point of a space of size 2^{D+k}. Every lower bound is an exhaustive search. The searches are large — the one behind Theorem 1.3 rules out all $\sum_{i \le 12} \binom{31}{i}$ subfamilies of the 31 admissible speeds — so the object that carries the mathematics is not the search output but the *completeness of the branching rule*, which we prove (Proposition 3.1) rather than assume; a natural-looking variant of the rule is not complete (Remark 3.3). Section 3 also isolates the device that makes the searches cheap enough to be checked by a proof assistant: an exclusion proved over a well-chosen *subset* of the coefficient space is a stronger statement than one proved over the whole space, and the subset shrinks the search tree by up to two orders of magnitude.

2. The covering model

We recall the covering formulation, which goes back to Chow and Rimanić [2, Lemma 2.1] and which we take in the form stated by Hu [1, Lemma 2.1]. Let $V_M = \mathbb{F}_q[T]_{<M}$ carry the coefficient pairing $\langle g, h \rangle = \sum_j g_j h_j$, let $A_k = \mathbb{F}_q[T]_{<k}$, and for $f \in \mathbb{F}_q[T] \setminus \{0\}$ of degree at most D put

$$U_f := fA_k = \operatorname{span}_{\mathbb{F}_q}\{f, Tf, \dots, T^{k-1}f\} \le V_{D+k}, \qquad K_f := U_f^{\perp} \le V_{D+k}. \tag{7}$$

Multiplication by f is injective, so $\dim U_f = k$, $\operatorname{codim} K_f = k$ and $|K_f| = q^D$; the equations cutting out K_f are the first k rows of a circulant convolution matrix, whence the name *partial-circulant kernel*.

Lemma 2.1 (Covering lemma; [2, Lemma 2.1], in the form of [1, Lemma 2.1])**.** *Let $\mathcal{F}$ be a family of nonzero polynomials of degree at most D. Then*

$$\delta(\mathcal{F}) < q^{-k} \quad \Longleftrightarrow \quad \bigcup_{f \in \mathcal{F}} K_f = V_{D+k}.$$

Lemma 2.1 is quoted, not reproved: it is the bridge between the metric definition (1) and the finite combinatorics, and it is the one ingredient of this paper that is taken from the literature rather than verified (see Section 9). Through it, $C_k^{(D)}(q)$ of (6) is the minimum number of partial-circulant kernels of speeds of degree at most D whose union is V_{D+k}, and all statements below are proved in that covering formulation.

The encoding. From here on $q = 2$. A speed $f = \sum_{j \le D} a_j T^j$ is recorded by the natural number whose j-th binary digit is a_j, and a point $x = \sum_{j<D+k} x_j T^j \in V_{D+k}$ likewise; so the 31 admissible speeds at $D = 4$ are the integers $1, \dots, 31$ and the coefficient space is $\{0, 1, \dots, 127\} \cong \mathbb{F}_2^7$. In this encoding (7) reads

$$x \in K_f \quad \Longleftrightarrow \quad \sum_{j=0}^{D} a_j x_{j+s} = 0 \ \text{ in } \mathbb{F}_2, \qquad 0 \le s < k, \tag{8}$$

which for $D = 4$, $k = 3$ is the displayed system [1, (3.1)] term for term; in the general form [1, (2.5)] the same k equations are written with the shift running over $1 \le s \le k$ and the point coordinates indexed from 1. We say that $\mathcal{F}$ *covers* if the kernels of its members exhaust V_{D+k}.

The definition (7) depends on the ambient space V_{D+k}, and D is a choice. The next proposition says the choice is immaterial, which is what makes the window value $C_k^{(D)}(2)$ a statement about the speeds; on the metric side this is transparent from (6), on the covering side it is not.

Proposition 2.2. *Let $\mathcal{F}$ be a family of nonzero polynomials over $\mathbb{F}_2$ of degree at most D and let $M \ge D$. Then $\mathcal{F}$ covers V_{D+k} if and only if it covers V_{M+k}. Moreover covering is monotone: if $\mathcal{F} \subseteq \mathcal{G}$ and $\mathcal{F}$ covers, then $\mathcal{G}$ covers.*

Proof sketch. The membership test (8) reads only the digits $0, \dots, D$ of the speed and $0, \dots, D + k - 1$ of the point, so truncating a point of V_{M+k} modulo T^{D+k} changes neither side; this is the remark in the proof of [1, Lemma 2.1] that "the coefficients of α beyond $T^{-(D+k)}$ do not affect these k equations", proved here as a statement about the encoded test rather than used informally. Monotonicity is immediate. No computation is involved. □

Remark 2.3. Proposition 2.2 gives the monotonicity of $D \mapsto C_k^{(D)}(2)$ in the covering formulation: a family of degree at most D that covers V_{D+k} also covers V_{D+1+k}, so widening a window can only lower the value. In particular a value computed in a wide window bounds the values in all narrower ones from below.

3. An exact set-cover search with a proved branching rule

Hu's suggestion — "the finite problem may be formulated as set cover on the collection of partial-circulant kernels" — is the algorithm of this section. For a list U of points still to be covered and a budget b of kernels still allowed, define

$$\operatorname{srch}(b, U) = \begin{cases} \text{true} & U = \emptyset, \\ \text{false} & U \neq \emptyset,\ b = 0, \\ \displaystyle\bigvee_{\deg a \le D,\ p \in K_a} \operatorname{srch}\big(b - 1,\ U \setminus K_a\big) & \text{otherwise, where } p \text{ is the first entry of } U. \end{cases} \tag{9}$$

The rule is: take any still-uncovered point and branch over *every* kernel that contains it. Its completeness is the following proposition, stated in the contrapositive form in which it is used.

Proposition 3.1. *Let U be a list of points of V_{D+k} and $b \ge 0$. If $\operatorname{srch}(b, U)$ is false, then for every family $\mathcal{F}$ of at most b nonzero polynomials of degree at most D there is a point of U lying in no K_f, $f \in \mathcal{F}$. Consequently, if $\operatorname{srch}(b, U)$ is false for some $U \subseteq V_{D+k}$, then no family of at most b speeds of degree at most D covers V_{D+k}.*

Proof sketch. Induction on b. A family covering the head p of U must contain some f with $p \in K_f$, and that f is one of the branches; deleting it from the family and K_f from U is the inductive step. The two lookup tables used by the implementation — the incidence lists $\{a : p \in K_a\}$ and the membership table of the kernels — are proved to agree with (8) rather than checked on instances. The second sentence follows because a point of $U \subseteq V_{D+k}$ missed by $\mathcal{F}$ witnesses that $\mathcal{F}$ does not cover. □

Remark 3.2 (Core compression). The second sentence of Proposition 3.1 is what makes the searches affordable: U need not be all of V_{D+k}. Covering a subset is easier than covering the whole space, so a subset that b kernels already fail to cover is a *stronger* hypothesis and a much cheaper search. At $k = 3$, $D \le 4$, $b = 12$ the search over all 128 points of $\mathbb{F}_2^7$ visits 43,193,961 nodes, whereas over the 37-point subset

$$\begin{aligned}\mathcal{C} = \{\, &6, 8, 9, 13, 17, 21, 24, 25, 30, 39, 43, 45, 47, 50, 52, 56, 62, 66, 69, 72, 73, 74, 79, 80,\\ &87, 88, 90, 95, 97, 99, 107, 111, 114, 115, 119, 121, 125 \,\} \subset \mathbb{F}_2^7\end{aligned} \tag{10}$$

it visits 797,161, a 54-fold reduction. (Node counts throughout are as measured by a reference implementation of the same recursion in C; the formal proofs evaluate that recursion inside the proof assistant, which reports no counts.) How $\mathcal{C}$ was found — greedy deletion of points in random order — is irrelevant to the proof, which uses only $\mathcal{C} \subseteq \mathbb{F}_2^7$. The same device takes the $k = 2$, $D \le 5$ exclusion from 79,029,201 to 12,204,241 nodes over a 61-point subset, and the enumeration of Section 6 from 129,582,648 to 2,391,484.

Remark 3.3 (A rule that is not complete). It is tempting to combine the branching rule of (9) with an index rule — to require that the kernels chosen along a branch have increasing indices, as one does when enumerating subsets. The combination is not complete: the branch point is forced by the point-selection rule, so the kernel that a given covering family uses to cover it may carry a smaller index than a kernel already chosen on that branch, and the index rule then discards that family. In our own preliminary computations this variant reported "no cover exists" on an instance of this very problem that has one. Proposition 3.1 is proved for the rule as displayed and for no other, which is what makes each exclusion below a theorem rather than the output of a program.

4. Hu's certificate, re-verified

Before using the search we re-derive the published data about $\mathcal{F}_*$ from the definitions. Write $\Phi(x) = \#\{f \in \mathcal{F} : x \in K_f\}$ for the covering multiplicity.

Proposition 4.1. *The thirteen kernels of $\mathcal{F}_*$ cover $\mathbb{F}_2^7$. The number of $x \in \mathbb{F}_2^7$ with $\Phi(x) = m$, for $m = 0, 1, \ldots, 13$ in order, is*

$$0,\ 83,\ 26,\ 15,\ 2,\ 0,\ 0,\ 1,\ 0,\ 0,\ 0,\ 0,\ 0,\ 1,$$

and the total incidence is $\sum_{x \in \mathbb{F}_2^7} \Phi(x) = 208 = 13 \cdot 2^4$.

This is Hu's certificate proposition [1, Proposition 3.1], verified rather than discovered; the only strengthening is that the multiplicity vector is asserted at every $m \le 13$, so that the absent entries are part of the statement, whereas [1] prints the six nonzero entries. The proof is the evaluation of (8) at all 13×128 speed-point pairs, done inside the proof assistant's kernel; it was reproduced beforehand by two independent programs (Section 9).

Proposition 4.2. *Every member of $\mathcal{F}_*$ is needed: for each $f \in \mathcal{F}_*$, the twelve kernels of $\mathcal{F}_* \setminus \{f\}$ do not cover $\mathbb{F}_2^7$.*

Proposition 4.2 says that $\mathcal{F}_*$ is minimal for inclusion. That is a weaker statement than Theorem 1.3 — it does not exclude a different family of twelve — but it is proved by a completely independent route, thirteen evaluations of the covering test with no search at all, and it is the control that would fail first if the encoding (8) were wrong.

5. The value in the window $D \le 4$

Theorem 5.1. *No family of at most twelve nonzero polynomials over $\mathbb{F}_2$ of degree at most four covers $\mathbb{F}_2^7$ at $k = 3$; equivalently, no such family has $\delta(\mathcal{F}) < 2^{-3}$. Together with Proposition 4.1,*

$$C_3^{(4)}(2) = 13,$$

and Hu's thirteen polynomials are a minimum covering family among all speeds of degree at most four.

Proof sketch. The upper bound is Proposition 4.1. The lower bound rests on an exhaustive computation, as follows. There are 31 admissible speeds (the nonzero polynomials of degree at most four over $\mathbb{F}_2$), so the families to be excluded are the $\sum_{i \le 12} \binom{31}{i}$ subsets of size at most twelve. We run the search (9) with $D = 4$, $k = 3$, budget $b = 12$ over the 37-point set $\mathcal{C}$ of (10); it terminates with the value false after 797,161 nodes. Since $\mathcal{C} \subseteq \mathbb{F}_2^7$, Proposition 3.1 converts that single Boolean into the assertion that no twelve of the 31 kernels cover $\mathbb{F}_2^7$. The corresponding search over all 128 points, which is the same statement proved the expensive way, closes after 43,193,961 nodes; that larger run is a cross-check outside the formal development. Finally $\delta(\mathcal{F}) < 2^{-3}$ is equivalent to covering by Lemma 2.1, and by Proposition 2.2 neither side depends on the ambient space. □

Remark 5.2 (The window $D \le 5$). Hu's family also lies in the window $D \le 5$, and by Proposition 2.2 it covers there, so $C_3^{(5)}(2) \le 13$. For the matching lower bound we have only an unverified computation, which we record as such: two independently written C programs report that the search (9) with $D = 5$, $k = 3$, $b = 12$ over the full space returns false after 54,045,122,343 nodes, agreeing digit for digit, which would give $C_3^{(5)}(2) = 13$. That search is four to five orders of magnitude larger than the one behind Theorem 5.1 and is not part of the formal development; core compression does not reduce it, since finding a subset as in Remark 3.2 costs one full search per deleted point. We therefore do not state $C_3^{(5)}(2) = 13$ as a theorem anywhere in this paper. The window $D \le 6$ is a further factor of roughly 4000 by the observed growth and would need integer programming or a SAT solver rather than depth-first search.

6. The optimal families

Theorem 5.1 leaves open how many families attain the minimum. Hu exhibits one and claims nothing about the set of optima. There are exactly three. In the encoding of Section 2, and then in polynomial notation,

$$\begin{aligned}
\mathcal{F}_* &= \{1, 6, 7, 8, 11, 13, 15, 18, 19, 20, 21, 25, 31\},\\
\mathcal{F}_{\mathrm{rev}} &= \{2, 5, 9, 12, 16, 19, 21, 22, 25, 26, 28, 30, 31\},\\
\mathcal{F}_\circ &= \{3, 4, 10, 14, 17, 19, 21, 23, 24, 25, 27, 29, 31\},
\end{aligned}$$

that is, $\mathcal{F}_*$ is (4) and

$$\begin{aligned}
\mathcal{F}_{\mathrm{rev}} = \{\, & T,\ T^2{+}1,\ T^3{+}1,\ T^3{+}T^2,\ T^4,\ T^4{+}T{+}1,\ T^4{+}T^2{+}1,\ T^4{+}T^2{+}T,\\
& T^4{+}T^3{+}1,\ T^4{+}T^3{+}T,\ T^4{+}T^3{+}T^2,\ T^4{+}T^3{+}T^2{+}T,\ T^4{+}T^3{+}T^2{+}T{+}1\,\},\\
\mathcal{F}_\circ = \{\, & T{+}1,\ T^2,\ T^3{+}T,\ T^3{+}T^2{+}T,\ T^4{+}1,\ T^4{+}T{+}1,\ T^4{+}T^2{+}1,\ T^4{+}T^2{+}T{+}1,\\
& T^4{+}T^3,\ T^4{+}T^3{+}1,\ T^4{+}T^3{+}T{+}1,\ T^4{+}T^3{+}T^2{+}1,\ T^4{+}T^3{+}T^2{+}T{+}1\,\}.
\end{aligned}$$

Theorem 6.1. *Let $\mathcal{F}$ be a family of at most thirteen nonzero polynomials over $\mathbb{F}_2$ of degree at most four whose kernels cover $\mathbb{F}_2^7$ at $k = 3$ — equivalently, with $\delta(\mathcal{F}) < 2^{-3}$. Then $\mathcal{F}$ is one of $\mathcal{F}_*$, $\mathcal{F}_{\mathrm{rev}}$, $\mathcal{F}_\circ$. Each of the three does cover, and each has the covering-multiplicity vector of Proposition 4.1.*

Proof sketch. Let $\mathrm{enum}(b, U)$ be the variant of (9) that returns the list of the families it builds instead of a Boolean; the same induction as Proposition 3.1 proves it complete, in the form: every family of at most b admissible speeds covering U contains one of the returned families. We run it with $b = 13$

over the 37-point set $\mathcal{C}$ of (10), an exhaustive computation returning 5,907 leaves; that they collapse to 1,274 distinct sets is a count made outside the formal development and is not used below. Each returned family is checked to have exactly thirteen members and no repetitions — thirteen and not fewer because the search behind Theorem 5.1 already excludes twelve kernels on $\mathcal{C}$ alone. Now let $\mathcal{F}$ cover $\mathbb{F}_2^7$ with $|\mathcal{F}| \le 13$. Then $\mathcal{F}$ covers $\mathcal{C}$, so it contains one of the returned families, which has thirteen distinct members; a thirteen-element subset of a family of at most thirteen elements is the whole family. It remains to test the returned families for covering all of $\mathbb{F}_2^7$, and exactly three pass. The multiplicity vectors of $\mathcal{F}_{\mathrm{rev}}$ and $\mathcal{F}_\circ$ are then computed as in Proposition 4.1. Run over the whole of $\mathbb{F}_2^7$ instead of $\mathcal{C}$ the enumeration visits 129,582,648 nodes and returns the same three families; that larger run is a cross-check outside the formal development. □

The three are not unrelated. The degree-4 window carries the involution $\rho\colon f(T) \mapsto T^4 f(1/T)$, which reverses the coefficient window $0, \dots, 4$; let ρ' denote the analogous reversal of the seven coordinates of $\mathbb{F}_2^7$.

Theorem 6.2. *For every f of degree at most four and every $x \in \mathbb{F}_2^7$,*

$$x \in K_{\rho f} \quad \Longleftrightarrow \quad \rho' x \in K_f.$$

Moreover $\rho(\mathcal{F}_*) = \mathcal{F}_{\mathrm{rev}}$ and $\rho(\mathcal{F}_\circ) = \mathcal{F}_\circ$.

Proof sketch. Reversing the speed window and the point window matches the s-th equation of (8) with the $(k-1-s)$-th. The displayed equivalence is verified at $(D,k) = (4,3)$ by evaluating both sides at all 32×128 pairs, and the two images are computed directly from the lists above. □

Since ρ is an involution of the 31 admissible speeds, the displayed equivalence gives at once that $\mathcal{F}$ covers iff $\rho(\mathcal{F})$ does, so ρ permutes the covering families of each size in this window; that last deduction is made here in prose, not in Lean.

Remark 6.3 (Outside the formal development). The substitution $\sigma\colon f(T) \mapsto f(T+1)$ is a ring automorphism of $\mathbb{F}_2[T]$ preserving degrees, and it extends to an isometry of $\mathbb{F}_2((T^{-1}))$ fixing $\mathbb{F}_2[T]$, so it preserves δ and the degree window. An auxiliary computation, made outside the formal development and stated here as an observation only, gives

$$\sigma(\mathcal{F}_*) = \mathcal{F}_*, \qquad \sigma(\mathcal{F}_{\mathrm{rev}}) = \mathcal{F}_\circ.$$

With Theorem 6.2, ρ acts on the three optimal families as the transposition $(\mathcal{F}_*\ \mathcal{F}_{\mathrm{rev}})$ and σ as $(\mathcal{F}_{\mathrm{rev}}\ \mathcal{F}_\circ)$; two transpositions generating the symmetric group on three letters, so the three optimal families form a single orbit. Note also, directly from the lists, that all three contain $T^4{+}T{+}1$, $T^4{+}T^2{+}1 = (T^2{+}T{+}1)^2$ and $T^4{+}T^3{+}T^2{+}T{+}1$.

How special is "exactly three"? The same enumeration in the window one degree lower returns exactly one optimal family, and there the optimum is the family Chow and Rimanić's construction names.

Proposition 6.4. *Let $\mathcal{F}$ be a family of at most fifteen nonzero polynomials over $\mathbb{F}_2$ of degree at most three whose kernels cover $\mathbb{F}_2^6$ at $k = 3$. Then $\mathcal{F}$ is the set of all fifteen nonzero polynomials of degree at most three.*

The proof is the enumeration used for Theorem 6.1, run at $D = 3$, $k = 3$ and budget 15 over the whole of $\mathbb{F}_2^6$.

Remark 6.5 (Outside the formal development). Auxiliary computations with the same enumeration, not part of the formal development, give the number of minimum covering families in further windows at $q = 2$: one in each of the tight windows $(k, D) = (2,2)$ and $(4,4)$ — always the standard family of all nonzero polynomials of degree at most k, as in the formally proved tight case $(3,3)$ of Proposition 6.4 — but 155 at $(k,D) = (1,4)$ and 281 at $(2,4)$. So the enumeration is not a procedure that returns a handful of families by construction: one degree past tight, at $k = 1$ and

$k = 2$, the optimum is very far from unique, and it becomes rigid again exactly in the window where Conjecture 1.1 fails.

7. Neighbouring windows

Theorem 7.1. $C_2^{(4)}(2) = C_2^{(5)}(2) = 7 = Q_2(2)$*: in both windows the minimum covering family has seven members, and the seven nonzero polynomials of degree at most two are one.*

Proof sketch. The upper bound in both windows is the standard family, checked by evaluation. The lower bounds are exhaustive. At $D \le 4$ there are 31 admissible speeds and 64 points; the search (9) with budget 6 over the whole of V_6 returns false after 966,377 nodes. At $D \le 5$ there are 63 admissible speeds and 128 points; the search with budget 6 over the whole space needs 79,029,201 nodes, and it is the 61-point subset of Remark 3.2 that is used in the formal proof, at 12,204,241 nodes. Proposition 3.1 turns each Boolean into the exclusion of every family of at most six speeds in the corresponding window. □

At $q = 2$, $k = 2$ the hypothesis of Theorem 1.2 holds only for $D \le 2$, and the unconditional bound [2, Theorem 1.7] is vacuous at $q = 2$; so Theorem 7.1 is a new determination in each of the two windows. The window between them follows without a further search: the standard family lies in it, and Remark 2.3 bounds $C_2^{(3)}(2)$ below by $C_2^{(4)}(2) = 7$, so $C_2^{(3)}(2) = 7$ as well — a corollary of Theorem 7.1, not a separate machine-checked statement. We stress the direction of the result: it *confirms* Conjecture 1.1 in these windows. The next proposition collects the windows where a published theorem predicts the answer and our computation reproduces it; they are controls on the instrument, not new values.

Proposition 7.2. $C_1^{(4)}(2) = 3$, $C_3^{(3)}(2) = 15$ *and* $C_4^{(4)}(2) = 31$ *— in each case the value* $Q_k(2)$*, attained by the family of all nonzero polynomials of degree at most* k*.*

For $k = 1$ the value follows from (3) alone, since $q^k + 1 = Q_1(q) = 3$ at $q = 2$; the searches involved are small (497 nodes at budget 2). For $k = 3$, $D \le 3$ and $k = 4$, $D \le 4$ the values are given by Theorem 1.2, whose hypothesis holds because $N(4,2)/(2^3 + 2^2 + 2) = 3/14 > \lfloor 3/4 \rfloor = 0$ and $N(5,2)/(2^4 + 2^3 + 2^2 + 2) = 6/30 > \lfloor 4/5 \rfloor = 0$; the exhaustive exclusions here take 1,707 and 63,579 nodes at budgets 14 and 30. These two are the controls that matter most, because $k = 3$, $D \le 4$ is the first window at $k = 3$ that Theorem 1.2 does not reach and is exactly where the value drops from 15 to 13.

Finally, two checks that the covering condition is a genuine constraint in the window where Theorem 5.1 lives.

Proposition 7.3. *At* $k = 3$ *no family of speeds of degree at most two covers* V_5 *— not even the family of all seven such speeds — so* $C_3^{(2)}(2) = \infty$*. At* $k = 3$*,* $D \le 4$ *the family* $\{1, T\}$ *is admissible and does not cover.*

The first is forced by density — seven kernels of density 2^{-3} cannot exhaust the space — and the computation confirms it; the second rules out the opposite degeneracy, that admissibility might imply covering. Together with Proposition 4.2 they are the negative controls of the development: a formalisation in which the covering predicate were accidentally vacuous, or accidentally universal, would fail one of them.

8. Further windows

Hu's problem has a second half — "classify all failures of the projective threshold for small pairs (q, k)" — and everything above is at $q = 2$. We close by pricing the three directions that seem nearest, in the units of the searches used here.

The rest of $C_3(2)$. Theorem 5.1 settles $D \le 4$ and Remark 5.2 reports an unverified computation at $D \le 5$; the exclusion at $D \le 6$ is a further factor of roughly 4000 by the growth observed between

consecutive windows, and depth-first search of the shape (9) is the wrong instrument at that size. What would change the picture is not a larger budget but a *proved* pruning rule usable inside the completeness argument of Proposition 3.1 — for instance a cardinality bound, discarding a branch as soon as the uncovered set is larger than the remaining budget times the kernel size. In the $k = 3$, $D \le 4$ window such a rule cuts the full-space search from 43,193,961 nodes to about 1.3×10^7 in our C implementation, a factor of 3.3, and combined with the subset device of Remark 3.2 it is the plausible route to $D \le 5$.

Odd characteristic. The smallest untouched cell is $q = 3$, $k = 2$, $D \le 3$: the ambient space is V_5 with $3^5 = 243$ points, there are $(3^4 - 1)/2 = 40$ monic speeds, and the target is $Q_2(3) = 13$. That is the size class of the $k = 4$ cell computed here, so the search itself is not the obstacle; what is missing is the encoding, since the bit-level representation of Section 2 is specific to $q = 2$.

A second counterexample? At $q = 2$, $k = 5$ the hypothesis of Theorem 1.2 holds for $D \le 5$, since $N(6, 2) = 9$ and $9/62 > \lfloor D/6 \rfloor$ exactly when $D \le 5$; so the first open window is $D = 6$, with $2^{11} = 2048$ points, 127 admissible speeds and target $Q_5(2) = 63$. A covering family of fewer than 63 speeds there would be a second failure of Conjecture 1.1, and the first at $k \ne 3$. Against that, Proposition 7.2 records a negative datum: at $k = 4$, $D \le 4$ — the window immediately after Hu's — the value is exactly $Q_4(2) = 31$, so whatever makes $k = 3$ special does not simply persist.

9. Verification

Every proposition and theorem of this paper except the two quoted results, Lemma 2.1 and Theorem 1.2, has been formally verified in Lean 4 [6] (toolchain `v4.32.0`) against *Mathlib* [7], in a development that is free of `sorry`; repository reference to be added at submission. Lemma 2.1 is quoted from [2, 1] and is the one step not verified here: the formal statements are about the covering formulation, and the passage from $\delta(\mathcal{F}) < q^{-k}$ to the covering property is that lemma. The reasoning layer — the completeness of the branching rule and of its enumerating variant (Proposition 3.1 and the argument of Theorem 6.1), the ambient independence of Proposition 2.2, the correctness of the two lookup tables, and the assembly of each window value from a witness and an exclusion — uses no axioms beyond `propext`, `Classical.choice` and `Quot.sound`, and so does every published datum re-derived in Section 4: Propositions 4.1, 4.2, 7.3 and Theorem 6.2 are checked by the kernel of the proof assistant, with no compiled evaluation. What is delegated to compiled evaluation, through Lean's `native_decide`, is exactly the exhaustive exclusions and the enumerations: one Boolean equation for Theorem 5.1 (797,161 nodes over the 37-point set), two for Theorem 7.1 (966,377 and 12,204,241 nodes), three for Proposition 7.2, and three each for Theorem 6.1 and Proposition 6.4, with one more recording the leaf count 5,907; each such equation carries one additional axiom, and each is given its meaning by the kernel-checked soundness statement above it. Every number attributed to [1] was reproduced from that paper's own definitions by two independent implementations, one in C and one in Python, before any formal proof was written, and the 37-point exclusion was reproduced by a third implementation, which reports the same 797,161 nodes and the same verdict. Finally, for the record: the arXiv comment of [1] announces that "Work-in-progress formalization and related code will be developed and maintained at" `github.com/hxypqr/lonely-runners-function-fields`. Checked on 28 August 2026, that address returns HTTP 404, as does the matching GitHub API entry; the author's public repository listing is readable and holds no formalization of [1]. The development reported here was therefore written independently of any formalization of [1].

References

[1] X. Hu, *Lonely Runners over Function Fields: Quantized Phase–Riesz product*, arXiv:2608.18818 (v1, 19 Aug 2026; v2, 21 Aug 2026). Numbers cited above are v2's; v2 differs from v1 only in a historical-attribution paragraph and in the acknowledgements.

[2] S. Chow and L. Rimanić, *Lonely runners in function fields*, Mathematika **65** (2019), no. 3, 677–701; arXiv:1711.01207.

[3] G. Perarnau and O. Serra, *The Lonely Runner Conjecture turns 60*, Comput. Sci. Rev. **58** (2025), Paper 100798; arXiv:2409.20160.

[4] S. Czerwiński and J. Grytczuk, *Invisible runners in finite fields*, Inform. Process. Lett. **108** (2008), no. 2, 64–67.
[5] A. Z. Wagner, *Constructions in combinatorics via neural networks*, arXiv:2104.14516 (2021).
[6] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE-28, Lecture Notes in Computer Science 12699, Springer, 2021.
[7] The mathlib Community, *The Lean mathematical library*, in: Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020.