A Density-One Sequence with Distinct Consecutive-Block Products

Przemyslaw Chojecki

January 19, 2026

Abstract

We construct a density-1 increasing sequence $1 \leq d_1 < d_2 < \cdots$ such that all consecutive-block products $\prod_{u \leq i \leq v} d_i$ are distinct. We implement a dyadic-stage deletion scheme in which collisions are encoded by separated-variable curve families of the form $\prod(x - a_i) = \prod(y + b_j)$. We introduce a capacity constraint enforced via the Moser–Tardos algorithmic Lovász local lemma, which yields a small deletion set at each dyadic scale hitting all "logarithmic-scale" collisions. We further show how to verify the requisite subpolynomial dependency degree in this logarithmic regime using uniform determinant-method bounds of Castryck–Cluckers–Dittmann–Nguyen and a degeneracy analysis via Bilu–Tichy and Hajdu–Tijdeman (Pell-type families).

1 Introduction

Let $d_1 < d_2 < \cdots$ be an increasing sequence of positive integers. We say that (d_i) has the *distinct consecutive-block product property* if

$$\prod_{u \leq i \leq v} d_i \neq \prod_{u' \leq i \leq v'} d_i \tag{1}$$

for all distinct pairs of index intervals $[u, v] \neq [u', v']$.

A convenient equivalent formulation is in terms of *disjoint block collisions*. If (1) fails, then after cancelling any overlap of the two index-intervals, one obtains a nontrivial equality of products of *disjoint* consecutive blocks:

$$d_{M-\ell+1} \cdots d_M = d_N \cdots d_{N+k}, \qquad M < N, \tag{2}$$

with $\ell \geq 1$, $k \geq 0$. Thus the problem is to build a density-1 sequence with *no* solutions to (2).

We seek such a sequence with asymptotic density 1:

$$\lim_{n\to\infty} \frac{|\{d_i \le n\}|}{n} = 1.$$

This would settle positively an open problem of Erdős:

Problem 1.1 (Problem 421, [1])**.** *Is there a sequence* $1 \le d_1 < d_2 < \cdots$ *with density* 1 *such that all products* $\prod_{u\le i\le v} d_i$ *are distinct?*

1.1 Main result

Theorem 1.1 (Density-1 construction)**.** *There exists a density-*1 *increasing sequence* (d_i) *satisfying* (1) *(equivalently, admitting no solutions to* (2)*).*

Acknowledgments. I am grateful to Terence Tao, uona and natso26 for useful discussions on erdosproblems.com forum regarding this problem.

1.2 Collision-length regimes

A collision (2) is specified by its two block lengths ℓ and $k+1$ and by its maximal term $m := d_M$. A useful meta-parameter is

$$d := \max(\ell, k+1),$$

which is also the algebraic degree of the associated split-product curve once we encode the collision (Module III). We will repeatedly choose the dyadic scale X with $m \in [X, 2X)$.

The argument naturally splits into several regimes for (ℓ, k) as functions of m (or X). The point of spelling these out early is that *only one* regime requires serious Diophantine inputs with uniformity in degree/height; the others are handled by much softer counting.

The long cutoff. Define the polylogarithmic threshold

$$L(X) \;:=\; 10\big(\log(2X)\big)^3. \tag{3}$$

Note that if $m \in [X, 2X)$ then $10(\log m)^3 \le L(X)$.

Long regime (killed at the sprinkling stage). We call a collision *long* if

$$d = \max(\ell, k+1) \;\ge\; 10(\log m)^3.$$

Module I constructs a density-1 set $A_0 \subset \mathbb{N}$ such that, for all sufficiently large m, no long collision with maximal term m occurs inside A_0 (Corollary 2.1). In particular, after passing to A_0, every collision at scale X must satisfy

$$d \le L(X).$$

Thus there is no remaining length regime above the polylog threshold.

Logarithmic regime (the only regime remaining after sprinkling). Fix dyadic X. A collision is in the *logarithmic regime at scale* X if its maximal term m lies in $[X, 2X)$ and

$$d = \max(\ell, k+1) \le L(X).$$

The rest of the paper eliminates *all* logarithmic-regime collisions at each dyadic scale X by deleting a small set $D_X \subset [X, 2X) \cap A_0$, while preserving global density 1.

Short vs. long logarithmic (inside the logarithmic regime). Within the logarithmic regime we further split by degree. Define

$$d_0(X) \ := \ \left\lceil (\log X)^{1/3} \log\log X \right\rceil . \tag{4}$$

- *Short logarithmic:* $d < d_0(X)$. These collisions are already sparse in $[X, 2X)^2$ by Bombieri–Pila type bounds (small fixed degree), and we remove them by a direct $o(X)$ "sparse cleanup" step (delete one endpoint from each witness).
- *Long logarithmic:* $d_0(X) \le d \le L(X)$. Here the degree grows (slowly) with X, and we need uniform determinant-method bounds (Castryck–Cluckers–Dittmann–Nguyen) together with a degeneracy analysis (Bilu–Tichy and Hajdu–Tijdeman, yielding only Pell-type exceptional families, which are dyadically $\ll \log X$). These bounds control the dependency degree in an algorithmic Lovász local lemma, and a Moser–Tardos construction then produces a deletion set D_X that simultaneously respects the capacity constraint and hits every long-logarithmic collision event.

Why this covers all collisions. Given any collision (2) with maximal term m and dyadic X such that $m \in [X, 2X)$, either it is long (hence ruled out already inside A_0), or it lies in the logarithmic regime $d \le L(X)$. In the latter case it is either short logarithmic (removed by sparse cleanup) or long logarithmic (removed by the capacity/encoding/Diophantine/LLL stage). Thus no collision survives in the final set.

1.3 Modular outline of the proof

After Module I produces the base set A_0 and eliminates the *long* regime, it suffices to destroy the remaining *logarithmic* collisions scale by scale. For each large dyadic X we will delete a small set

$$D_X \subset [X, 2X) \cap A_0$$

and define the final set

$$A := A_0 \setminus \bigcup_{X \text{ dyadic}} D_X. \tag{5}$$

The key requirements are:

(i) *(Density)* $\sum_{X \le N} |D_X| = o(N)$, so A has density 1.

(ii) *(Local regularity)* D_X satisfies a capacity constraint, ensuring that consecutive elements of $A \cap [X, 2X)$ can be recovered from bounded local windows.

(iii) *(Hitting)* every logarithmic-regime collision with maximal term in $[X, 2X)$ is hit by D_X.

Items (ii)–(iii) are achieved by a Moser–Tardos algorithmic Lovász local lemma, once we verify that the dependency degree is subpolynomial; this is where the uniform Diophantine bounds (Castryck–Cluckers–Dittmann–Nguyen) and the degeneracy analysis (Bilu–Tichy; Hajdu–Tijdeman; Pell-type families) enter.

1.4 Module interfaces (statements used as black boxes)

We now state the main proposition supplied by each module. The remainder of the paper is devoted to proving these propositions, and Theorem 1.1 follows immediately from them.

Module I (sprinkling): eliminate long collisions and produce density-1 A_0.

Proposition 1.2 (Module I)**.** *There exists a set $A_0 \subset \mathbb{N}$ of asymptotic density* 1 *such that the following hold.*

(i) ***(Polylog kill).*** *If* $(d_i^{(0)})$ *is the increasing enumeration of* A_0*, then for all sufficiently large maximal terms* $m = d_M^{(0)}$ *there is no collision* (2) *with maximal term* m *and*

$$\max(\ell, k+1) \;\geq\; 10(\log m)^3.$$

Equivalently, for all sufficiently large dyadic X*, every collision inside* A_0 *with maximal term in* $[X, 2X)$ *must satisfy* $d = \max(\ell, k+1) \leq L(X)$*.*

(ii) ***(Local regularity of*** A_0***).*** *Define*

$$U(X) := \lceil 100 \log(2X) \rceil .$$

Then for all sufficiently large dyadic X*, every integer interval* $I \subset [X/2, 2X)$ *with*

$$|I| \;\leq\; 2L(X)$$

satisfies

$$|I \setminus A_0| \;\leq\; U(X).$$

Sketch of Proposition 1.2. We delete each $n \geq 3$ independently with probability $\delta(n) = 1/(\log n)^2$ and take A_0 to be the complement. Density-1 follows from a standard strong-law/second-moment estimate. For a fixed maximal term m, there are at most m^4 candidate ordered pairs of consecutive blocks among integers $\leq m$. Any collision with $\max(\ell, k+1) \geq 10(\log m)^3$ requires at least that many distinct integers to survive, which happens with probability $\leq m^{-10}$; a union bound gives $\ll m^{-6}$, summable in m, so Borel–Cantelli applies.

Module II (capacity and encoding): from collisions to split-product curves. Fix dyadic X and write

$$A_X := ([X, 2X) \cap A_0) \setminus D_X.$$

We impose a *capacity constraint* at scale X: there exists a parameter $T(X) \ll L(X)$ such that

$$\forall \text{ integer intervals } I \subset [X, 2X) \text{ with } |I| = L(X), \qquad |D_X \cap I| \leq T(X). \quad (6)$$

Proposition 1.3 (Module II)**.** *Assume* D_X *satisfies* (6)*. Then any logarithmic-regime collision* (2) *inside* A_X *with maximal term in* $[X, 2X)$ *and degree*

$d = \max(\ell, k+1) \le L(X)$ *yields an integer solution* $(x, y) \in [X, 2X)^2$ *to an equation of the form*

$$\prod_{i=0}^{\ell-1}(x - a_i) = \prod_{j=0}^{k}(y + b_j), \tag{7}$$

where $a_0 = b_0 = 0$*,* $\ell, k+1 \le L(X)$*, and the shifts satisfy*

$$0 \le a_i, b_j \le C\big(L(X) + T(X) + U(X)\big)$$

for an absolute constant C *and* $U(X) := \lceil 100 \log(2X) \rceil$*.*

Sketch of Proposition 1.3. Given a collision with maximal term $x \in [X, 2X)$, define backward shifts by $a_i := x - d_{M-i}$ and forward shifts by $b_j := d_{N+j} - d_N$. This transforms (2) into (8). Capacity ensures that in any length-$L(X)$ window, only $T(X)$ points are missing, so the predecessor and successor gaps needed to define the $\le L(X)$ factors are bounded by $O(L(X) + T(X))$. Thus the shift vectors are uniformly bounded in terms of $\mathrm{polylog}(X)$, which in turn controls the degree and coefficient height of the defining polynomial.

We emphasize that we only use the forward implication

$$\text{collision in } A_X \implies \text{integer solution of (8)},$$

and we do *not* claim the converse for arbitrary bounded shifts.

Module III (Diophantine bounds): uniform point bounds for the collision curve family. Write $d = \max(\ell, k+1)$ for the degree of (8), and split the logarithmic regime at $d_0(X)$ (defined in (4)).

Proposition 1.4 (Module III)**.** *There is a function* $\Delta(X) = X^{o(1)}$ *such that for all sufficiently large dyadic* X *the following holds.*

Consider any split-product equation

$$\prod_{i=0}^{\ell-1}(x - a_i) \;=\; \prod_{j=0}^{k}(y + b_j), \tag{8}$$

with $a_0 = b_0 = 0$*,* $\ell, k+1 \le L(X)$*, and shifts satisfying*

$$0 \le a_i, b_j \;\le\; C\big(L(X) + T(X) + U(X)\big)$$

for an absolute constant C*. Write* $d = \max(\ell, k+1)$ *and split at* $d_0(X)$*.*

(i) ***(Short logarithmic: Bombieri–Pila sparsity.)*** *If $d < d_0(X)$, then the set of endpoints $x \in [X, 2X)$ for which there exist admissible parameters and some $y \in [X, 2X)$ satisfying (8) has size $o(X)$.*

(ii) ***(Long logarithmic: CCDN + degeneracy filter.)*** *If $d_0(X) \le d \le L(X)$, then for each fixed admissible equation (8) and each absolutely irreducible component C of the curve defined by (8) one has*

$$|C(\mathbb{Z}) \cap [X, 2X)^2| \ \le \ \Delta(X).$$

Moreover, any exceptional (degenerate) components not controlled directly by determinant-method bounds are shown, via Bilu–Tichy and Hajdu–Tijdeman, to belong to Pell/Dickson-type families; such families contribute $O(\log X)$ solutions in $[X, 2X)^2$ and hence are absorbed into $\Delta(X)$.

All bounds are uniform over admissible parameters at scale X.

Sketch of Proposition 1.4. All collision curves are of the form $f(x) = g(y)$ with

$$f(t) = \prod_{i=0}^{\ell-1}(t - a_i), \qquad g(t) = \prod_{j=0}^{k}(t + b_j),$$

where $d = \max(\ell, k+1) \le L(X)$ and all shifts satisfy $a_i, b_j \ll \mathrm{polylog}(X)$. Consequently the coefficient height of the defining polynomial is bounded in terms of d and $\log X$ (and remains subpolynomial in X throughout the range $d \le L(X)$).

For *short logarithmic* degree $d < d_0(X)$, Bombieri–Pila bounds the number of integer points per irreducible component in $[X, 2X)^2$ by $X^{1/d+o(1)} \le X^{1/2+o(1)}$. Summing over all admissible parameter choices in this range (whose total number is $X^{o(1)}$ because $d_0(X) = o(\log X)$ and shifts are $\mathrm{polylog}(X)$) yields a total witness set occupying only $o(X)$ distinct endpoints x, hence removable by a sparse cleanup.

For *long logarithmic* degree $d_0(X) \le d \le L(X)$, we apply determinant-method bounds with polynomial dependence on d and controlled dependence on coefficient height, in particular the uniform results of Castryck–Cluckers–Dittmann–Nguyen. Since $d \ge d_0(X)$ one has $X^{1/d} = X^{o(1)}$, and since $d \le L(X) = \mathrm{polylog}(X)$ all degree- and height-dependent losses remain subpolynomial in X. This gives a uniform $X^{o(1)}$ bound on $|C(\mathbb{Z}) \cap [X, 2X)^2|$ for each nondegenerate component C.

The remaining risk is that $f(x) = g(y)$ has a *degenerate* component supporting atypically many integer points. Here we use Bilu–Tichy as a structural filter: infinite families force (f, g) into a finite list of standard pairs. In our split-root setting, refinements such as Hajdu–Tijdeman (for polynomials with simple rational roots) rule out all but Pell/Dickson-type exceptional families. Such Pell-type parametrizations have exponential growth, so they contribute $\ll \log X$ solutions in a dyadic box, which are absorbed into $\Delta(X)$.

Finally, in Module IV the Lovász local lemma needs a *per-variable* dependency bound. This is obtained by combining the per-component dyadic bound $|C(\mathbb{Z}) \cap [X, 2X)^2| \le \Delta(X)$ with the fact that the number of admissible parameter choices at scale X is $X^{o(1)}$ (shifts are polylog(X) and $d \le L(X)$), so each fixed integer participates in at most $X^{o(1)}$ long-logarithmic collision events.

Module IV (dyadic stage deletion): sparse cleanup + Moser–Tardos with capacity.

Proposition 1.5 (Module IV)**.** *For all sufficiently large dyadic X, there exists a deletion set $D_X \subset [X, 2X) \cap A_0$ such that:*

(i) D_X satisfies the capacity constraint (6)*;*

(ii) in $A_X := ([X, 2X) \cap A_0) \setminus D_X$ there is no logarithmic-regime collision with maximal term in $[X, 2X)$, i.e. no collision (2) *with maximal term in $[X, 2X)$ and $d \le L(X)$;*

(iii) $|D_X| = o(X)$.

Sketch of Proposition 1.5. At scale X, first remove short-logarithmic witnesses by deleting one endpoint from each (using Proposition 1.4(i)); this costs $o(X)$ deletions and can be arranged without breaking capacity because the witness endpoints are sparse. For the remaining long-logarithmic constraints $d_0(X) \le d \le L(X)$, we set up an LLL instance with variables indexed by integers in $[X, 2X) \cap A_0$ and bad events consisting of (a) capacity violations in length-$L(X)$ windows and (b) survival of a collision pattern (equivalently, an admissible curve point surviving with no deletion among its participants). Proposition 1.3 reduces (b) to integer points on admissible split-product curves.

Proposition 1.4(ii) bounds the number of dyadic solutions on each absolutely irreducible component of any fixed admissible curve by $\Delta(X) = X^{o(1)}$ (with degenerate Pell/Dickson families contributing only $O(\log X)$).

Together with the parameter count for admissible shift patterns at scale X (which is $X^{o(1)}$ because $d \le L(X) = \text{polylog}(X)$ and shifts are polylog(X)), this yields a subpolynomial dependency bound for the collision events. Indeed, fixing a variable corresponding to a specific integer $n \in [X, 2X) \cap A_0$, any collision event that depends on this variable must come from an admissible curve point $(x, y) \in [X, 2X)^2$ and an admissible local shift pattern in which n appears among the at most $O(d)$ participating integers (on either side of the collision). For each fixed admissible pattern, Proposition 1.4(ii) bounds the total number of dyadic solutions (x, y) on each irreducible component by $\Delta(X) = X^{o(1)}$ (degenerate Pell/Dickson families contribute only $O(\log X)$ and are absorbed). Multiplying by the number of admissible patterns and the $O(d)$ ways for n to occur within a collision, we obtain that n is involved in at most $X^{o(1)}$ collision bad events. Separately, n belongs to only $O(L(X))$ length-$L(X)$ windows, so it is contained in only $O(L(X)) = X^{o(1)}$ capacity bad events. Hence the overall dependency degree in the Moser–Tardos/LLL instance is subpolynomial.

Choosing deletion probability with slack, the Moser–Tardos algorithm produces a realization with no bad events, hence capacity and complete elimination of logarithmic collisions at scale X, with $|D_X| = o(X)$.

1.5 Deriving the main theorem from the modules

Lemma 1.6 (Dyadic $o(X)$ deletions preserve density)**.** *Let $A_0 \subset \mathbb{N}$ have density 1, and suppose that for all sufficiently large dyadic X we have $D_X \subset [X, 2X)$ with $|D_X| = o(X)$. Then the set A defined in (5) has asymptotic density 1.*

Proof. Summing over dyadic $X \leq N$ gives $\sum_{X \leq N} |D_X| = o(N)$, hence $\left|\bigcup_{X \leq N} D_X\right| = o(N)$. Removing an $o(N)$ subset from a density-1 set preserves density 1. □

Proof of Theorem 1.1 assuming Propositions 1.2–1.5. Let A_0 be as in Proposition 1.2. For each sufficiently large dyadic X, take D_X from Proposition 1.5 and define A by (5). By Lemma 1.6, A has density 1.

Let $d_1 < d_2 < \cdots$ be the increasing enumeration of A. Suppose for contradiction that a collision (2) occurs in A and let $m = d_M$ be its maximal term. Choose dyadic X with $m \in [X, 2X)$. By Proposition 1.2, for all sufficiently large m we have $d = \max(\ell, k + 1) \leq 10(\log m)^3 \leq L(X)$, so this is a logarithmic-regime collision at scale X. But by Proposition 1.5(ii), no such collision can occur in A_X, and since $A \cap [X, 2X) \subset A_X$ this is a contradiction. Therefore no collision exists, and (1) holds for (d_i). □

1.6 Where the module propositions are proved

- Proposition 1.2: Module I (sprinkling), Section 2.
- Proposition 1.3: Module II (capacity and encoding), Section 3.
- Proposition 1.4: Module III (Diophantine bounds), Section 4.
- Proposition 1.5: Module IV (dyadic stage deletion via Moser–Tardos), Section 5.

2 Module I: sprinkling and the polylog cutoff

This section constructs a density-1 set $A_0 \subset \mathbb{N}$ with a single crucial property: *all collisions with long block length are eliminated.* The key output is Corollary 2.1, which implies that after passing to A_0 the only possible collisions at dyadic scale X have

$$\max(\ell, k+1) \;\le\; L(X) \;:=\; 10\big(\log(2X)\big)^3.$$

All later modules work entirely inside A_0 and only address this remaining logarithmic regime.

2.1 Sprinkling model

For $n \ge 3$ set

$$\delta(n) \;:=\; \frac{1}{(\log n)^2}.$$

Let $S \subset \mathbb{N}$ be a random set obtained by deleting each integer $n \ge 3$ independently with probability $\delta(n)$ (and always keeping $1, 2$), and define the *sprinkled set*

$$A_0 \;:=\; \mathbb{N} \setminus S.$$

2.2 Main output

Corollary 2.1 (Polylog cutoff for collisions)**.** *With probability* 1*, there exists* m_0 *such that for all* $m \ge m_0$ *the following holds. Let* $d_1^{(0)} < d_2^{(0)} < \cdots$ *be the increasing enumeration of* A_0*. Then there is no solution to* (2) *with maximal term* $d_M^{(0)} = m$ *and*

$$\max(\ell, k+1) \;\ge\; 10(\log m)^3.$$

The proof has two independent parts: density bookkeeping and elimination of long collisions.

2.3 Density bookkeeping

Lemma 2.2 (Density of A_0)**.** *With probability* 1 *one has* $|S\cap[1,N]| = o(N)$ *as* $N\to\infty$*. Consequently* A_0 *has asymptotic density* 1*.*

Proof. Write $X_n := \mathbf{1}_{\{n\in S\}}$ for $n\geq 3$. Then X_n are independent Bernoulli variables with $\mathbb{E}X_n = \delta(n)$. For dyadic $N=2^j$ define

$$Z_N \;:=\; \sum_{N<n\leq 2N} X_n \;=\; |S\cap(N,2N]|.$$

Its expectation satisfies

$$\mu_N \;:=\; \mathbb{E}Z_N \;=\; \sum_{N<n\leq 2N}\delta(n) \;\asymp\; \frac{N}{(\log N)^2}.$$

A Chernoff bound gives $\mathbb{P}(Z_N\geq 2\mu_N)\leq \exp(-c\mu_N)$ for an absolute constant $c>0$. Since $\mu_N \gg N/(\log N)^2$, the sum $\sum_{j\geq 1}\mathbb{P}(Z_{2^j}\geq 2\mu_{2^j})$ converges, so by Borel–Cantelli we have $Z_{2^j}\leq 2\mu_{2^j}$ for all sufficiently large j almost surely. Therefore, for dyadic 2^J,

$$|S\cap[1,2^J]| \;=\; O(1)+\sum_{j<J} Z_{2^j} \;\ll\; \sum_{j<J}\frac{2^j}{j^2} \;\ll\; \frac{2^J}{J^2} \;=\; o(2^J),$$

almost surely. This implies $|S\cap[1,N]| = o(N)$ for all N, hence A_0 has density 1. □

2.4 Counting candidate collisions up to a maximal term

Lemma 2.3 (A uniform m^4 bound)**.** *Fix* $m\geq 1$ *and let* $E\subset\{1,2,\dots,m\}$ *be any set, with increasing enumeration* $e_1<e_2<\cdots<e_r$*. Then the number of* ordered pairs *of consecutive blocks*

$$(e_u\cdots e_v,\;\; e_{u'}\cdots e_{v'})$$

is at most m^4*. In particular, the number of candidate equalities*

$$\prod_{u\leq i\leq v} e_i \;=\; \prod_{u'\leq i\leq v'} e_i$$

between consecutive-block products is at most m^4*, uniformly in* E*.*

Proof. There are at most $r(r+1)/2\leq m^2$ choices for a consecutive block $[u,v]$, and likewise at most m^2 choices for $[u',v']$. Hence the number of ordered pairs of blocks is at most m^4. □

2.5 Eliminating long collisions: union bound and Borel–Cantelli

For each integer $m \geq 3$, let E_m be the event that there exists a collision (2) inside A_0 with maximal term m and with

$$\max(\ell, k+1) \;\geq\; 10(\log m)^3.$$

Lemma 2.4 (Long collisions are summably rare)**.** *There is an absolute constant $C > 0$ such that for all $m \geq 3$,*

$$\mathbb{P}(E_m) \;\leq\; C\, m^{-6}.$$

Proof. Fix $m \geq 3$ and condition on the random set $A_0 \cap [1, m]$, with increasing enumeration $e_1 < \cdots < e_r$. Any collision among consecutive blocks inside $A_0 \cap [1, m]$ yields an equality

$$\prod_{u \leq i \leq v} e_i \;=\; \prod_{u' \leq i \leq v'} e_i,$$

and after cancelling overlap one obtains a disjoint-block collision of the form (2). We do not need to count disjointness sharply; we bound the total number of ordered block-pairs.

By Lemma 2.3, the number of candidate ordered block-pairs is at most m^4. Fix one such candidate equality whose reduced form (2) has $\max(\ell, k+1) \geq 10(\log m)^3$. Then at least $10(\log m)^3$ distinct integers in $\{1, \ldots, m\}$ must lie in A_0 (namely the entire longer block). For each $n \leq m$ we have $\delta(n) \geq \delta(m)$, hence $\mathbb{P}(n \in A_0) = 1 - \delta(n) \leq 1 - \delta(m)$. By independence, the probability that all these required integers survive is at most

$$(1-\delta(m))^{10(\log m)^3} \;\leq\; \exp\Big(-\delta(m)\cdot 10(\log m)^3\Big) \;=\; \exp\Big(-\tfrac{10(\log m)^3}{(\log m)^2}\Big) \;=\; m^{-10}.$$

Taking a union bound over at most m^4 candidates yields

$$\mathbb{P}(E_m) \;\leq\; m^4 \cdot m^{-10} \;=\; m^{-6},$$

and we absorb small values of m into an absolute constant C. $\square$

Proof of Corollary 2.1. By Lemma 2.4, the series $\sum_{m \geq 3} \mathbb{P}(E_m)$ converges. By the Borel–Cantelli lemma, with probability 1 only finitely many of the events E_m occur. Equivalently, there exists m_0 such that for all $m \geq m_0$ there is no collision (2) inside A_0 with maximal term m and $\max(\ell, k+1) \geq 10(\log m)^3$. $\square$

Lemma 2.5 (Local regularity of A_0)**.** *With probability* 1*, for all sufficiently large dyadic* X *the following holds. Let*

$$L := L(X) = 10\big(\log(2X)\big)^3, \qquad U(X) := \lceil 100\log(2X)\rceil\,.$$

Then every integer interval $I \subset [X/2, 2X)$ *with* $|I| \le 2L$ *satisfies*

$$|I \setminus A_0| \;\le\; U(X).$$

Proof. Fix dyadic X and an interval $I \subset [X/2, 2X)$ of length $|I| \le 2L$. For $n \in I$ the indicators $\mathbf{1}_{\{n \notin A_0\}}$ are independent Bernoulli variables with

$$\mathbb{P}(n \notin A_0) = \delta(n) = \frac{1}{(\log n)^2} \;\le\; \frac{2}{(\log X)^2}$$

for large X and all $n \in [X/2, 2X)$. Hence

$$\mu := \mathbb{E}|I \setminus A_0| \;\le\; |I| \cdot \frac{2}{(\log X)^2} \;\le\; \frac{4L}{(\log X)^2} \;\ll\; \log X.$$

A Chernoff bound gives

$$\mathbb{P}\big(|I \setminus A_0| \ge U(X)\big) \;\le\; \exp(-c\log X)$$

for an absolute $c > 0$. The number of such intervals I is $O(X)$, so a union bound gives total probability

$$O(X) \cdot \exp(-c\log X) = O(X^{1-c}).$$

Choosing the constant 100 in $U(X)$ large enough makes $c > 3$, so the sum of these probabilities over dyadic X converges. By Borel–Cantelli, the conclusion holds for all large dyadic X almost surely. □

Proof of Proposition 1.2. For (i) combine Lemma 2.2 (density 1) with Corollary 2.1 (polylog cutoff).

(ii) is Lemma 2.5. □

Remark 2.6 (How the rest of the paper uses A_0). Fix one realization of A_0 for which Lemma 2.2 and Corollary 2.1 hold. All subsequent deletions are performed *within* this fixed A_0. In particular, for every sufficiently large dyadic X, any collision inside A_0 with maximal term in $[X, 2X)$ must satisfy

$$\max(\ell, k+1) \;\le\; L(X) \;=\; 10\big(\log(2X)\big)^3.$$

Thus later modules only need to eliminate collisions in this logarithmic regime.

3 Module II: capacity constraint and collision encoding

In this module we show that *local regularity* of the surviving set at a dyadic scale forces every remaining logarithmic-length collision to take a controlled algebraic form. Concretely: under a capacity constraint, any collision (2) with

$$\max(\ell, k+1) \;\le\; L(X)$$

can be encoded as an integer point on a separated-variable (split-product) curve

$$\prod_{i=0}^{\ell-1}(x-a_i) \;=\; \prod_{j=0}^{k}(y+b_j),$$

where the shift parameters a_i, b_j are uniformly bounded by a polylogarithmic quantity.

3.1 Dyadic setup and capacity

Fix a dyadic X and abbreviate

$$L := L(X) = 10\big(\log(2X)\big)^3.$$
$$U := U(X) = \lceil 100\log(2X)\rceil\,.$$

Let $B \subset [X, 2X) \cap \mathbb{Z}$ be the set of integers that survive at scale X (in the application, B will be the surviving set inside $[X, 2X)$ after sprinkling and dyadic deletion). Write the increasing enumeration as

$$b_1 < b_2 < \cdots < b_{|B|}.$$

By Lemma 2.5, any length-$\le 2L$ integer window at scale X contains at most U points missing from A_0. Thus for $B = ([X, 2X) \cap A_0) \setminus D_X$ and any integer interval $I \subset [X, 2X)$ with $|I| = L$,

$$|I \setminus B| \;\le\; |I \setminus A_0| + |D_X \cap I| \;\le\; U + T(X).$$

We impose a capacity constraint on the *missing* integers in B.

Definition 3.1 (Capacity at scale X)**.** Let $T = T(X)$ be an integer with $0 \le T \le L/2$. We say that B satisfies *capacity* (L, T) if for every integer interval

$$I = [u, u+L-1] \subset [X, 2X)$$

one has

$$|I \setminus B| \;\le\; T, \quad \text{equivalently} \quad |I \cap B| \;\ge\; L - T.$$

Remark 3.2 (How capacity is used later). The only purpose of capacity is to ensure that the predecessor/successor structure in B is controlled by a bounded value window: blocks of length at most L cannot reach far in value. This turns collisions into algebraic equations with uniformly bounded shifts. In Module IV we enforce capacity as part of the Moser–Tardos/LLL stage.

3.2 Predecessors and successors are local

Assume B satisfies capacity (L, T) in the sense of Definition 3.1. For $z \in B$ and an integer $r \geq 1$, define $\mathrm{Pred}_r(z)$ (resp. $\mathrm{Succ}_r(z)$) to be the r-th predecessor (resp. successor) of z in B whenever it exists.

Lemma 3.3 (Local window for predecessors and successors). *Assume $B \subset [X, 2X) \cap A_0$ satisfies capacity (L, T) with $T \leq L/2$ in the sense of Definition 3.1. Then for every $z \in B$ and every $1 \leq r \leq L$ for which $\mathrm{Pred}_r(z)$ exists,*

$$z - \mathrm{Pred}_r(z) \;\leq\; L + 2T + U.$$

Similarly, for every $z \in B$ and every $1 \leq r \leq L$ for which $\mathrm{Succ}_r(z)$ exists,

$$\mathrm{Succ}_r(z) - z \;\leq\; L + 2T + U.$$

Proof. Fix $z \in B$ and $1 \leq r \leq L$.

Predecessors. Consider the interval

$$J = [z - (L + 2T + U),\, z - 1].$$

For all sufficiently large X one has $2T + U < L$, hence $|J| < 2L$. If $z < X + (L + 2T + U)$ we intersect J with $[X, 2X)$, which only strengthens the bound.

Cover J by two consecutive length-L intervals

$$I_1 = [z - (L + 2T + U),\, z - (L + 1)], \qquad I_2 = [z - L,\, z - 1].$$

By capacity, each I_i has at most T missing integers from B, so

$$|J \setminus B| \leq 2T.$$

By Lemma 2.5, at most U integers of J fail to lie in A_0. Therefore

$$|J \cap B| \geq |J| - 2T - U = (L + 2T + U) - 2T - U = L.$$

Thus J contains at least L elements of B, hence at least r predecessors of z. This gives $z - \mathrm{Pred}_r(z) \leq L + 2T + U$.

Successors. The argument is identical with the interval $[z + 1,\, z + (L + 2T + U)]$. □

Remark 3.4. Lemma 3.3 is deliberately coarse and uses only $T \leq L/2$. In applications we will take $T \ll L$ (e.g. a small power of $\log X$), so $L+2T \asymp L$.

3.3 Encoding collisions by split-product curves

We now show that any logarithmic-length collision inside B gives rise to a curve equation with uniformly bounded shifts. Recall the disjoint-block collision form (2):

$$d_{M-\ell+1} \cdots d_M \;=\; d_N \cdots d_{N+k} \qquad \text{with } M < N.$$

Inside this module we will simply write the collision in terms of the enumeration of B.

Definition 3.5 (Block collision in B)**.** Let $b_1 < b_2 < \cdots$ enumerate B. A *collision of type* (ℓ, k) in B is an equality

$$b_{M-\ell+1} \cdots b_M \;=\; b_N \cdots b_{N+k} \tag{9}$$

for some indices $M < N$. We call it *logarithmic-length* if $\max(\ell, k+1) \leq L$.

Proposition 3.6 (Module II: encoding under capacity)**.** *Assume B satisfies capacity (L, T) with $T \leq L/2$. Let* (9) *be a logarithmic-length collision in B, and set*

$$x := b_M, \qquad y := b_N.$$

Define shift vectors by

$$a_0 := 0, \quad a_i := x - b_{M-i} \;\; (1 \leq i \leq \ell-1), \qquad b_0 := 0, \quad b_j := b_{N+j} - y \;\; (1 \leq j \leq k).$$

Then:

(i) The shifts are strictly increasing:

$$0 = a_0 < a_1 < \cdots < a_{\ell-1}, \qquad 0 = b_0 < b_1 < \cdots < b_k.$$

(ii) The shifts are uniformly bounded:

$$a_{\ell-1} \;\leq\; L + 2T + U, \qquad b_k \;\leq\; L + 2T + U.$$

(iii) The collision (9) *is equivalent to the split-product equation*

$$\prod_{i=0}^{\ell-1} (x - a_i) \;=\; \prod_{j=0}^{k} (y + b_j). \tag{10}$$

Proof. (i) Since $b_{M-i} < b_M = x$ for $i \geq 1$, each $a_i = x - b_{M-i}$ is positive and increases strictly with i. Similarly $b_{N+j} > b_N = y$ for $j \geq 1$, so $b_j = b_{N+j} - y$ is positive and increases strictly with j.

(ii) The quantity $a_{\ell-1} = x - b_{M-(\ell-1)}$ is the distance from x to its $(\ell-1)$-st predecessor in B. Since $\ell - 1 \leq \ell \leq L$, Lemma 3.3 gives $a_{\ell-1} \leq L + 2T + U$. Likewise $b_k = b_{N+k} - y$ is the distance from y to its k-th successor in B, and since $k \leq L$ we obtain $b_k \leq L + 2T + U$.

(iii) By definition,

$$b_{M-i} = x - a_i \quad (0 \leq i \leq \ell - 1), \qquad b_{N+j} = y + b_j \quad (0 \leq j \leq k).$$

Substituting these identities into (9) yields (10), and conversely (10) implies (9) by the same substitution. ☐

Remark 3.7 (Admissible parameter space and coefficient size). Proposition 3.6 shows that every logarithmic-length collision in B is encoded by data

$$(\ell, k, \mathbf{a}, \mathbf{b}, x, y)$$

with $\ell, k + 1 \leq L$ and with $0 = a_0 < a_1 < \cdots < a_{\ell-1} \leq L + 2T + U$ and $0 = b_0 < b_1 < \cdots < b_k \leq L + 2T + U$. The associated polynomial

$$F_{\mathbf{a},\mathbf{b}}(x, y) := \prod_{i=0}^{\ell-1}(x - a_i) - \prod_{j=0}^{k}(y + b_j)$$

has degree $d = \max(\ell, k + 1) \leq L$. Moreover, expanding the products shows that its coefficients are (signed) elementary symmetric polynomials in the shifts, hence are bounded in absolute value by $(L + 2T + U)^{O(d)}$. In particular, when L and T are polylogarithmic in X, the coefficient height is $\exp(\text{polylog} \log X)$, which is the regime needed for the uniform determinant-method bounds in Module III.

Remark 3.8 ("Admissible" means "arising from the encoding"). In this paper we will never use a converse statement of the form "every solution to a split-product equation with small shifts comes from a collision". Instead, throughout Modules III–IV we interpret "admissible" as: *the shift data and the split-product equation arise from the predecessor/successor encoding in the current surviving set under the capacity hypothesis.* This one-way interpretation is sufficient for the Lovász local lemma stage, because we only need:

$$\text{collision in the surviving set} \implies \text{occurrence of a corresponding cylinder/witness event.}$$

Remark 3.9 (What Module II does not do). Module II does not count solutions to (10). It only shows that under capacity, the collision problem reduces to integer points on a family of split-product curves with uniformly bounded shifts. All quantitative counting (Bombieri–Pila, CCDN, and the Bilu–Tichy/Hajdu–Tijdeman degeneracy filter) enters in Module III.

4 Module III: Diophantine bounds for the split-product curve family

In Module II we showed that, under the capacity constraint at dyadic scale X, every logarithmic-regime collision with

$$d := \max(\ell, k+1) \leq L(X)$$

is equivalent to an integer solution $(x, y) \in [X, 2X)^2$ of a split-product equation

$$\prod_{i=0}^{\ell-1}(x - a_i) \;=\; \prod_{j=0}^{k}(y + b_j), \tag{11}$$

where $a_0 = b_0 = 0$, $\ell, k+1 \leq L(X)$, and the shifts satisfy $a_i, b_j \ll L(X) + T(X)$. This module supplies the Diophantine input used later in Module IV: uniform bounds for the number of integer solutions in dyadic boxes, with special care for uniformity in the degree.

4.1 Admissible parameters

Fix a dyadic X and abbreviate

$$L := L(X) = 10\big(\log(2X)\big)^3, \qquad d_0 := d_0(X) = \left\lceil (\log X)^{1/3} \log\log X \right\rceil.$$

$$U = U(X) := \lceil 100 \log(2X) \rceil$$

Let $T := T(X)$ be the capacity parameter from Module II and set

$$H := H(X) := C_*(L + T + U),$$

where $C_* \geq 1$ is the absolute constant from Module II controlling the shift sizes. We call a choice

$$\ell \geq 1, \qquad k \geq 0, \qquad d = \max(\ell, k+1) \leq L,$$

together with strictly increasing shifts

$$0 = a_0 < a_1 < \cdots < a_{\ell-1} \le H, \qquad 0 = b_0 < b_1 < \cdots < b_k \le H,$$

admissible at scale X. For admissible data define

$$F_{\mathbf{a},\mathbf{b}}(x,y) := \prod_{i=0}^{\ell-1}(x-a_i) - \prod_{j=0}^{k}(y+b_j) \ \in\ \mathbb{Z}[x,y].$$

Then (11) is exactly $F_{\mathbf{a},\mathbf{b}}(x,y) = 0$.

4.2 Linear components are harmless

The determinant-method bounds we use later are intended for components of degree ≥ 2. We isolate linear components separately; in our application they never contribute relevant dyadic solutions with $0 < x < y$.

Lemma 4.1 (Linear components are harmless)**.** *Let $X \ge 10$ and let $(\ell, k, \mathbf{a}, \mathbf{b})$ be admissible at scale X. Suppose $F_{\mathbf{a},\mathbf{b}}(x,y) = 0$ has an absolutely irreducible component C which is a line. Then every integer point $(x,y) \in C(\mathbb{Z})$ satisfies either $y \le x$ or $y \le 0$. In particular, there is no point of $C(\mathbb{Z})$ with $0 < x < y$.*

Proof. Write $F_{\mathbf{a},\mathbf{b}}(x,y) = f(x) - g(y)$ with

$$f(t) := \prod_{i=0}^{\ell-1}(t-a_i), \qquad g(t) := \prod_{j=0}^{k}(t+b_j).$$

If a line C is an irreducible component, then there exist $(u,v,w) \in \mathbb{Q}^3 \setminus \{(0,0,0)\}$ such that C is given by $ux+vy+w=0$ and $ux+vy+w$ divides $f(x)-g(y)$ in $\mathbb{Q}[x,y]$. If $v=0$, then C is vertical and contributes at most finitely many points, all with fixed x. If $u=0$, then C is horizontal and similarly contributes finitely many points with fixed y, hence no points in $[X,2X)^2$ for large X.

Assume now $uv \ne 0$ and solve for $y = \alpha x + \beta$ with $\alpha, \beta \in \mathbb{Q}$. Then $f(x) = g(\alpha x + \beta)$ as polynomials in x. Comparing leading terms gives $1 = \alpha^{k+1}$, hence $\alpha = 1$ if $k+1$ is odd, and $\alpha = \pm 1$ if $k+1$ is even. Thus C has the form $y = x + t$ or $y = -x + t$ with $t \in \mathbb{Q}$.

If C is $y = -x + t$, then any integer point has $y \le t$, hence for X large there is no point with $y \in [X, 2X)$; in particular no point with $0 < x < y$. If C is $y = x + t$, then any integer point satisfies $y \le x$ when $t \le 0$. If instead

$t > 0$, then $g(y) = g(x+t)$ has all its roots at $y = -b_j$, so as a polynomial in x the right-hand side has roots at $x = -t - b_j < 0$. But $f(x)$ has all its roots at $x = a_i \geq 0$. Hence $f(x) = g(x+t)$ forces both sides to have the same multiset of roots, impossible unless all roots lie in $\{0\}$, which contradicts admissibility for X large. Therefore $t \leq 0$, hence $y \leq x$ on integer points of C. □

4.3 Module III output

Proposition 4.2 (Module III)**.** *There exists a function $\Delta(X) = X^{o(1)}$ such that for all sufficiently large dyadic X the following hold.*

(i) ***Short logarithmic regime (endpoint sparsity).*** *Let $S_{\text{short}}(X)$ be the set of integers $x \in [X, 2X) \cap \mathbb{Z}$ for which there exist admissible data with degree $d < d_0$ and some $y \in [X, 2X) \cap \mathbb{Z}$ such that $F_{\mathbf{a},\mathbf{b}}(x, y) = 0$. Then $|S_{\text{short}}(X)| = o(X)$.*

(ii) ***Long logarithmic regime (uniform dyadic point bounds).*** *Fix admissible data with collision degree $d_0 \leq d \leq L$ and let C be any absolutely irreducible component of the curve $F_{\mathbf{a},\mathbf{b}}(x, y) = 0$. If C is a line then $C(\mathbb{Z}) \cap [X, 2X)^2 = \varnothing$ for large X by Lemma 4.1. Otherwise,*

$$|C(\mathbb{Z}) \cap [X, 2X)^2| \; \leq \; \Delta(X).$$

Moreover, if C is exceptional *in the sense that it contains infinitely many integer points, then C is forced (in our split-root setting) into Pell/Dickson-type families via Bilu–Tichy and Hajdu–Tijdeman, and in particular*

$$|C(\mathbb{Z}) \cap [X, 2X)^2| \; \ll \; \log X,$$

which is absorbed into $\Delta(X)$.

All bounds are uniform over admissible parameters at scale X.

4.4 Short logarithmic regime: Bombieri–Pila and a parameter count

We use the quantitative bound of Bombieri–Pila.

Theorem 4.3 (Bombieri–Pila [2, Theorem 5])**.** *Let $G(x, y) \in \mathbb{Z}[x, y]$ be absolutely irreducible of degree $d \geq 2$. Then for all integers $N \geq \exp(d^6)$ one has*

$$\#\{(x, y) \in \mathbb{Z}^2 : \; 1 \leq x \leq N, \; 1 \leq y \leq N, \; G(x, y) = 0\} \; \leq \; N^{1/d} \exp\Big(12\sqrt{d \log N \log\log N}\Big).$$

Lemma 4.4 (Counting admissible choices)**.** *For each $D \geq 1$, the number of admissible parameter choices with collision degree $d \leq D$ is*

$$\leq \sum_{d \leq D} (H+1)^{O(d)}.$$

In particular, since $H = \mathrm{polylog}(X)$ and $D = d_0(X) = o(\log X)$, this number is $X^{o(1)}$.

Proof. For fixed d there are $O(d^2)$ choices of (ℓ, k) with $\max(\ell, k+1) = d$. Given (ℓ, k), the number of strictly increasing ℓ-tuples (a_i) in $[0, H]$ with $a_0 = 0$ is $\leq (H+1)^{\ell-1}$, and similarly the number of (b_j) is $\leq (H+1)^k$. Since $(\ell - 1) + k \leq 2d$, the number of choices is $\leq (H+1)^{O(d)}$. □

Proposition 4.5 (Short-logarithmic endpoints are sparse)**.** *For $X \to \infty$ one has $|S_{\text{short}}(X)| = o(X)$.*

Proof. Fix $\varepsilon > 0$. For each admissible choice with collision degree $d < d_0$, decompose $F_{\mathbf{a},\mathbf{b}}(x, y) = 0$ into absolutely irreducible components. Linear components contribute no points with $0 < x < y$ by Lemma 4.1, hence are irrelevant for collisions and can be ignored in the count. For each remaining component of degree ≥ 2, apply Theorem 4.3 with $N = 2X$. Since $d < d_0(X)$ and X is large, the hypothesis $N \geq \exp(d^6)$ holds. Thus each component contributes at most

$$(2X)^{1/d} \exp\Big(12\sqrt{d \log(2X) \log\log(2X)}\Big) \;\leq\; X^{1/d + o(1)} \;\leq\; X^{1/2 + o(1)}.$$

Summing over all admissible choices with $d < d_0$ using Lemma 4.4 gives a total of $X^{1/2+o(1)} \cdot X^{o(1)} = o(X)$ integer points in $[X, 2X)^2$ arising from the short-log regime. Projecting to the x-coordinate can only decrease cardinality, hence $|S_{\text{short}}(X)| = o(X)$. □

4.5 Long logarithmic regime: CCDN uniform determinant-method bounds

We now use the uniform determinant-method bound of Castryck–Cluckers–Dittmann–Nguyen. For planar curves they record the convenient specialization

$$N_{\text{aff}}(C, B) \;\ll\; \delta^3 B^{1/\delta}(\log B + \delta),$$

where $\delta = \deg(C)$, see [3, (4.3.2)]; this is deduced there from their main curve bound [3, Theorem 3].

Theorem 4.6 (Castryck–Cluckers–Dittmann–Nguyen [3, Theorem 3])**.** *There exists an absolute constant $c' \geq 1$ such that the following holds. Let $C \subset \mathbb{A}^2_{\mathbb{Q}}$ be an integral affine curve of degree $\delta \geq 2$. Then for all $B \geq 2$,*

$$\#\big(C(\mathbb{Z}) \cap [-B, B]^2\big) \;\leq\; c'\, \delta^3\, B^{1/\delta}\big(\log B + \delta\big). \tag{12}$$

Proposition 4.7 (Nonlinear components have $X^{o(1)}$ dyadic points)**.** *Let X be sufficiently large dyadic and let C be an absolutely irreducible component of $F_{\mathbf{a},\mathbf{b}}(x, y) = 0$ which is* not *a line. If $\deg(C) \geq d_0(X)$ then*

$$|C(\mathbb{Z}) \cap [X, 2X)^2| \;\leq\; X^{o(1)}.$$

Proof. Let $\delta := \deg(C) \geq d_0(X)$. Apply Theorem 4.6 with $B = 2X$. Since $[X, 2X)^2 \subset [-2X, 2X]^2$, we obtain

$$|C(\mathbb{Z}) \cap [X, 2X)^2| \;\leq\; c'\, \delta^3\, (2X)^{1/\delta}\big(\log(2X) + \delta\big).$$

Because $\delta \geq d_0(X) \to \infty$, we have $X^{1/\delta} = X^{o(1)}$. Also $\delta \leq d \leq L(X) = \mathrm{polylog}(X)$, hence $\delta^3(\log X + \delta) = X^{o(1)}$. Multiplying gives $|C(\mathbb{Z}) \cap [X, 2X)^2| \leq X^{o(1)}$. □

4.6 Exceptional components: Bilu–Tichy + Hajdu–Tijdeman and Pell/Dickson sparsity

Write again $F_{\mathbf{a},\mathbf{b}}(x, y) = f(x) - g(y)$ with f, g split into *distinct* linear factors over $\mathbb{Q}$. We use the standard classification of infinite-solution separated-variable equations.

Theorem 4.8 (Bilu–Tichy [5, Theorem 1.1])**.** *Let $f, g \in \mathbb{Q}[t]$ be nonconstant. If the Diophantine equation $f(x) = g(y)$ has infinitely many integer solutions, then, up to composing f and g with linear polynomials and a common polynomial, the pair (f, g) belongs to the finite list of standard pairs in the sense of Bilu–Tichy.*

In our split-root setting one can sharpen the conclusion using Hajdu–Tijdeman.

Theorem 4.9 (Hajdu–Tijdeman [6, Theorem 9.1])**.** *Assume f and g have only simple rational roots. If $f(x) = g(y)$ has infinitely many rational solutions with bounded denominator (in particular, infinitely many integer solutions), then (f, g) is forced into the PTE/Dickson-type configurations described in [6, Theorem 9.1]. In particular, any infinite integer-solution family arising from such a component is governed by a Pell/Dickson-type parametrization, and the sizes of solutions grow exponentially along the family.*

Lemma 4.10 (Pell/Dickson families are dyadically sparse). *Let* $\{(x_n, y_n)\}_{n\geq 0} \subset \mathbb{Z}^2$ *satisfy an exponential growth condition: there exist* $c > 1$ *and* n_0 *such that for all* $n \geq n_0$*,*

$$|x_{n+1}| + |y_{n+1}| \;\geq\; c\big(|x_n| + |y_n|\big).$$

Then

$$\#\{n : \; (x_n, y_n) \in [X, 2X)^2\} \;\ll\; \log X.$$

Proof. If $(x_n, y_n) \in [X, 2X)^2$ then $|x_n| + |y_n| \in [2X, 4X)$. By exponential growth, the quantity $|x_n| + |y_n|$ can lie in a fixed dyadic interval for at most $O(\log X)$ indices. □

Proposition 4.11 (Exceptional components contribute at most $O(\log X)$ dyadic points). *Let* X *be sufficiently large dyadic and let* C *be an absolutely irreducible component of* $F_{\mathbf{a},\mathbf{b}}(x, y) = 0$ *which contains infinitely many integer points. Then*

$$|C(\mathbb{Z}) \cap [X, 2X)^2| \;\ll\; \log X.$$

Proof. If C contains infinitely many integer points, then $f(x) = g(y)$ has infinitely many integer solutions. Apply Theorem 4.8 and Theorem 4.9 to deduce that the integer solutions on C lie in finitely many Pell/Dickson-type parametrized families with exponential growth. Lemma 4.10 then gives $O(\log X)$ solutions in $[X, 2X)^2$ for each family. □

4.7 Proof of Proposition 4.2

Proof of Proposition 4.2. Part (i) is Proposition 4.5.

For part (ii), fix admissible data with collision degree $d_0 \leq d \leq L$ and let C be an absolutely irreducible component of $F_{\mathbf{a},\mathbf{b}}(x, y) = 0$.

If C is a line, then $C(\mathbb{Z})$ contains no point with $0 < x < y$ by Lemma 4.1, hence no point in $[X, 2X)^2$ for large X.

Assume C is not a line. If C contains infinitely many integer points, Proposition 4.11 yields

$$|C(\mathbb{Z}) \cap [X, 2X)^2| \;\ll\; \log X.$$

Otherwise $C(\mathbb{Z})$ is finite, hence $|C(\mathbb{Z}) \cap [X, 2X)^2| = O(1)$.

Finally, if $\deg(C) \geq d_0(X)$ then Proposition 4.7 yields $|C(\mathbb{Z}) \cap [X, 2X)^2| \leq X^{o(1)}$ uniformly. Taking

$$\Delta(X) := \max\big\{X^{o(1)}, \; \log X\big\} = X^{o(1)}$$

concludes part (ii). □

Remark 4.12 (Fiber bound (occasionally convenient)). For any fixed admissible data and fixed $x \in \mathbb{Z}$, the equation $F_{\mathbf{a},\mathbf{b}}(x,y) = 0$ is a polynomial equation in y of degree $\le k+1 \le d$, hence has at most $d \le L(X)$ integer solutions y.

5 Module IV: dyadic stage deletion via Moser–Tardos

Fix a sufficiently large dyadic X. Throughout this module we abbreviate

$$L := L(X) = 10\big(\log(2X)\big)^3, \qquad T := T(X) = \Big\lceil (\log X)^{2/3} \Big\rceil, \qquad p := \frac{T}{10L}.$$

5.1 Universe and random deletions

Let

$$V_X^{\mathrm{buf}} := [X/2, 2X) \cap A_0, \qquad V_X := [X, 2X) \cap A_0.$$

We use V_X^{buf} only to define local windows (so that predecessor lists near X are visible without boundary artefacts), but *we only delete inside* V_X.

Expose independent Bernoulli variables $(\xi_n)_{n \in V_X}$ with

$$\mathbb{P}(\xi_n = 1) = p.$$

Define the deletion set and the surviving sets by

$$D_X := \{n \in V_X : \ \xi_n = 1\}, \qquad B_X := V_X \setminus D_X, \qquad B_X^{\mathrm{buf}} := V_X^{\mathrm{buf}} \setminus D_X.$$

Thus B_X^{buf} is the dyadic "buffered" survivor set (points below X are never deleted at scale X, but may be used as predecessors).

5.2 Capacity bad events

For each integer interval $I \subset [X, 2X)$ of length L, define overflow and underflow events

$$\mathsf{B}_I^+ := \big\{|D_X \cap I| > T\big\}, \qquad \mathsf{B}_I^- := \big\{|D_X \cap I| < T/100\big\}, \qquad \mathsf{B}_I := \mathsf{B}_I^+ \cup \mathsf{B}_I^-.$$

The overflow condition is exactly the capacity requirement for B_X at scale X, since $|I \setminus B_X| = |D_X \cap I|$. The underflow condition is not needed for capacity, but is a convenient technical device: it forces local deletion patterns to contain $\Omega(T)$ ones, making cylinder-event probabilities extremely small.

5.3 Local windows and predecessor functionals

For $u \in [X, 2X)$ define the left window

$$W_u := [u - (L + 2T + U),\ u) \cap V_X^{\mathrm{buf}}.$$

Write $\sigma_u \in \{0,1\}^{W_u \cap V_X}$ for the restriction $(\xi_n)_{n \in W_u \cap V_X}$. (So σ_u records *only* the random bits in the window; points of W_u below X carry no variables at scale X.)

Lemma 5.1 (No overflow implies enough kept points in windows)**.** *Assume $\bigcap_I \neg \mathsf{B}_I^+$ (no overflow on any length-L interval in $[X, 2X)$). Then for every $u \in [X, 2X)$ the window W_u contains at least L kept points of B_X^{buf}, i.e.*

$$|W_u \cap B_X^{\mathrm{buf}}| \ \geq\ L.$$

Proof. Let $\widetilde{W}_u := [u - (L + 2T + U),\ u)$ be the underlying integer interval. Since $|\widetilde{W}_u| \leq 2L$ for all sufficiently large X, Lemma 2.5 gives

$$|\widetilde{W}_u \setminus A_0| \ \leq\ U,$$

hence

$$|W_u| = |\widetilde{W}_u \cap A_0| \geq |\widetilde{W}_u| - U = L + 2T.$$

Moreover, $\widetilde{W}_u \cap [X, 2X)$ has integer length at most $L + 2T + U \leq 2L$, hence it is contained in the union of two length-L intervals in $[X, 2X)$. Under $\bigcap_I \neg \mathsf{B}_I^+$, each such length-$L$ interval contains at most T elements of D_X, so

$$|D_X \cap \widetilde{W}_u| \ \leq\ 2T.$$

Therefore

$$|W_u \cap B_X^{\mathrm{buf}}| = |W_u| - |D_X \cap W_u| \geq (L + 2T) - 2T = L.$$

□

Lemma 5.2 (Local predecessor functional)**.** *Assume $\bigcap_I \neg \mathsf{B}_I^+$. Then for each $u \in [X, 2X)$ and each $1 \leq t \leq L$, the t-th predecessor of u among the kept points in B_X^{buf} is a deterministic function of the pattern σ_u on $W_u \cap V_X$. Consequently, for each $1 \leq r \leq L$ the list of the first r predecessors of u in B_X^{buf} is determined by σ_u.*

Proof. By Lemma 5.1, W_u contains at least L kept points of B_X^{buf}. Since the only randomness in B_X^{buf} comes from deletions inside V_X, the set $W_u \cap B_X^{\mathrm{buf}}$ is determined exactly by which points of $W_u \cap V_X$ are deleted. That is, it is determined by σ_u. The predecessor list is obtained by scanning leftward through W_u and selecting successive kept points, so it is a deterministic function of σ_u. □

5.4 Admissible local patterns and their count

Definition 5.3 ((Two-sided) admissible patterns)**.** Fix $u \in [X, 2X)$. We call $\sigma_u \in \{0,1\}^{W_u \cap V_X}$ (L,T)-*admissible* if for every integer interval

$$J \subset [u - (L + 2T + U),\, u) \quad \text{with} \quad |J| = L$$

one has

$$\frac{T}{200} \;\le\; \sum_{n \in J \cap V_X} \sigma_u(n) \;\le\; \frac{T}{2}.$$

We write Σ_u for the set of (L,T)-admissible patterns at u.

Lemma 5.4 (Pattern count)**.** *For each $u \in [X, 2X)$ one has*

$$|\Sigma_u| \;\le\; \sum_{j \le CT} \binom{|W_u \cap V_X|}{j} \;=\; \exp\bigl(O(T \log(L/T))\bigr) \;=\; X^{o(1)}$$

for an absolute constant $C \ge 1$.

Proof. Cover W_u by $O(1)$ length-L integer intervals J. The upper admissibility bound $\sum_{n \in J \cap V_X} \sigma_u(n) \le T/2$ implies that σ_u has Hamming weight $O(T)$ on $W_u \cap V_X$. Counting patterns of weight $\le CT$ on a set of size $O(L)$ gives the binomial sum. Since $T \log(L/T) = o(\log X)$ for our choice of parameters, this is $X^{o(1)}$. □

5.5 Collision cylinder events

Fix $x, y \in [X, 2X) \cap \mathbb{Z}$ with $x < y$, fix lengths $1 \le \ell \le L$ and $1 \le k+1 \le L$, and fix patterns $\sigma_x \in \Sigma_x$ and $\sigma_y \in \Sigma_y$.

By Lemma 5.2, the pattern σ_x determines the predecessor list of x in B_X^{buf} and hence determines shifts $a_0 = 0 < a_1 < \cdots < a_{\ell-1}$ for the left block of length ℓ. Similarly, σ_y determines shifts $b_0 = 0 < b_1 < \cdots < b_k$ for the right block of length $k+1$. Therefore the split-product polynomial

$$F_{(\sigma_x,\ell),(\sigma_y,k)}(X,Y) := \prod_{i=0}^{\ell-1}(X - a_i) - \prod_{j=0}^{k}(Y + b_j)$$

is fixed once $(x, y, \ell, k, \sigma_x, \sigma_y)$ are fixed.

Define the collision index set $\mathcal{T}_X$ to consist of all tuples

$$\tau = (x, y, \ell, k, \sigma_x, \sigma_y)$$

with $x < y$, $x, y \in [X, 2X) \cap \mathbb{Z}$, $1 \le \ell \le L$, $1 \le k+1 \le L$, $\sigma_x \in \Sigma_x$, $\sigma_y \in \Sigma_y$, and such that

$$F_{(\sigma_x,\ell),(\sigma_y,k)}(x, y) = 0. \tag{13}$$

For each $\tau \in \mathcal{T}_X$ define the *collision cylinder event* A_τ by

$$\mathsf{A}_\tau := \big\{(\xi_n)_{n\in W_x\cap V_X} = \sigma_x \ \text{ and } \ (\xi_n)_{n\in W_y\cap V_X} = \sigma_y\big\}.$$

Under A_τ, the predecessor lists at x and y are exactly those encoded by σ_x and σ_y, so (13) becomes an actual collision identity between consecutive blocks in B_X^{buf}.

5.6 Dependency degree

Two bad events are adjacent if they depend on a common variable ξ_n.

Proposition 5.5 (Subpolynomial dependency degree)**.** *For all sufficiently large dyadic X, the maximum dependency degree among the events*

$$\{\mathsf{A}_\tau : \tau \in \mathcal{T}_X\} \ \cup \ \{\mathsf{B}_I : |I| = L\}$$

is $X^{o(1)}$.

Proof. Fix $n \in V_X$.

(i) Capacity events. The variable ξ_n appears in B_I only when $n \in I$. There are $O(L)$ length-L intervals $I \subset [X, 2X)$ containing n.

(ii) Cylinder events. If ξ_n appears in A_τ, then $n \in (W_x \cap V_X) \cup (W_y \cap V_X)$. For fixed n, there are only $O(L)$ choices of $u \in [X, 2X)$ with $n \in W_u$ since W_u has length $O(L)$. Hence there are $O(L^2)$ possible endpoint pairs (x, y) with $n \in W_x \cup W_y$. For each such (x, y) there are at most L^2 choices of (ℓ, k) and at most

$$|\Sigma_x||\Sigma_y| \ = \ X^{o(1)}$$

choices of (σ_x, σ_y) by Lemma 5.4. Therefore the number of cylinder events adjacent to ξ_n is $X^{o(1)}$. □

5.7 Coverage: collisions force cylinder events

Lemma 5.6 (Coverage)**.** *Assume $\bigcap_I \neg\mathsf{B}_I^+$ (capacity holds). If there exists a logarithmic-length collision in B_X^{buf} with maximal term in $[X, 2X)$, then there exists $\tau \in \mathcal{T}_X$ such that the cylinder event A_τ occurs.*

Proof. Let $y \in [X, 2X)$ be the maximal element of the collision and let $x \in [X, 2X)$ be the endpoint of the other block, so $x < y$ and the block lengths satisfy $\max(\ell, k+1) \le L$.

Under $\bigcap_I \neg \mathsf{B}_I^+$, Lemma 5.2 implies that the predecessor lists needed to form these blocks are determined by the realized window restrictions

$$\sigma_x := (\xi_n)_{n \in W_x \cap V_X}, \qquad \sigma_y := (\xi_n)_{n \in W_y \cap V_X}.$$

Define $\tau = (x, y, \ell, k, \sigma_x, \sigma_y)$. By construction, the collision identity is exactly (13), so $\tau \in \mathcal{T}_X$, and the event A_τ holds by definition. □

5.8 Moser–Tardos at scale X

Theorem 5.7 (Moser–Tardos output at scale X)**.** *For all sufficiently large dyadic X, there exists a choice of $(\xi_n)_{n \in V_X}$ and hence a deletion set $D_X \subset V_X$ such that:*

(i) No cylinder event occurs: A_τ fails for every $\tau \in \mathcal{T}_X$.

(ii) No overflow occurs: B_I^+ fails for every length-L interval $I \subset [X, 2X)$.

(iii) $|D_X| = o(X)$.

Consequently, B_X satisfies capacity (L, T) and B_X^{buf} contains no logarithmic-length collision with maximal term in $[X, 2X)$.

Proof. We use the product measure on $(\xi_n)_{n \in V_X}$ with $\mathbb{P}(\xi_n = 1) = p = T/(10L)$.

(A) Probabilities of bad events. Fix a length-L interval $I \subset [X, 2X)$. Then $\mathbb{E}|D_X \cap I| \asymp pL = T/10$. A Chernoff bound gives

$$\mathbb{P}(\mathsf{B}_I^+) \le \exp(-cT), \qquad \mathbb{P}(\mathsf{B}_I^-) \le \exp(-cT)$$

for an absolute constant $c > 0$.

Now fix $\tau \in \mathcal{T}_X$. Since $\sigma_x \in \Sigma_x$ and $\sigma_y \in \Sigma_y$, admissibility forces $\Omega(T)$ ones inside $W_x \cap V_X$ and also inside $W_y \cap V_X$. Therefore the cylinder event A_τ fixes at least $c_0 T$ variables to equal 1, for some absolute $c_0 > 0$. Consequently

$$\mathbb{P}(\mathsf{A}_\tau) \;\le\; p^{c_0 T} \;=\; \exp\bigl(-\Omega(T \log(L/T))\bigr).$$

(B) Apply the asymmetric Lovász local lemma and Moser–Tardos. Assign LLL-weights

$$x(\mathsf{A}_\tau) := \exp\bigl(-\Omega(T \log(L/T))\bigr), \qquad x(\mathsf{B}_I) := \exp(-\Omega(T)).$$

By Proposition 5.5, each bad event has at most $X^{o(1)}$ neighbors. Since $T\log(L/T) \to \infty$ and $T \to \infty$, we have

$$\sum_{\mathsf{F}\sim\mathsf{E}} x(\mathsf{F}) = o(1)$$

uniformly over bad events E, so the asymmetric LLL conditions hold for large X. By the Moser–Tardos theorem [4], there exists an assignment with no bad events. This gives (i) and (ii).

(C) Size bound. On $\bigcap_I \neg\mathsf{B}_I^+$, every length-$L$ interval contains at most T deleted points. Therefore

$$|D_X| \le \Big(\lceil X/L \rceil + 1\Big) T = o(X),$$

which gives (iii). ☐

5.9 Conclusion of Module IV

Proposition 5.8 (Module IV: dyadic deletion)**.** *For all sufficiently large dyadic X, there exists a deletion set $D_X \subset [X, 2X) \cap A_0$ such that:*

(i) $B_X := ([X, 2X) \cap A_0) \setminus D_X$ *satisfies capacity* $(L(X), T(X))$*;*

(ii) $B_X^{\text{buf}} := ([X/2, 2X) \cap A_0) \setminus D_X$ *contains no logarithmic-length collision* (2) *with maximal term in* $[X, 2X)$*;*

(iii) $|D_X| = o(X)$*.*

Proof. Apply Theorem 5.7. Item (i) is exactly the failure of all overflow events B_I^+. If a logarithmic-length collision with maximal term in $[X, 2X)$ existed in B_X^{buf}, then by Lemma 5.6 some cylinder event A_τ would occur, contradicting Theorem 5.7(i). Finally, Theorem 5.7(iii) gives $|D_X| = o(X)$. ☐

6 Assembling dyadic scales and completing the proof of the main theorem

We now build the final set

$$A := A_0 \setminus \bigcup_{X \text{ dyadic}} D_X.$$

Lemma 6.1 (Dyadic $o(X)$ deletions preserve density)**.** *If* A_0 *has density* 1 *and* $|D_X| = o(X)$ *for each dyadic* X*, then* A *has density* 1*.*

Proof. For N large, summing $|D_X| = o(X)$ over dyadic $X \le N$ gives

$$\Big|\bigcup_{X\le N} D_X\Big| = o(N),$$

so removing $\bigcup_{X\le N} D_X$ from a density-1 set preserves density 1. □

Proof of Theorem 1.1 assuming Modules I–IV. Let $d_1 < d_2 < \cdots$ enumerate A. By Lemma 6.1, the set A has density 1.

Suppose for contradiction that a collision (2) occurs in A. Let m be the maximal term of this collision and choose dyadic X such that $m \in [X, 2X)$.

By Proposition 1.2 (Module I), for all sufficiently large m this collision lies in the logarithmic regime at scale X, i.e. $d = \max(\ell, k+1) \le L(X)$.

Apply Proposition 5.8 (Module IV) at this dyadic scale X. It provides a set $D_X \subset [X, 2X) \cap A_0$ such that

$$B_X^{\mathrm{buf}} := ([X/2, 2X) \cap A_0) \setminus D_X$$

contains no logarithmic-length collision with maximal term in $[X, 2X)$.

Since later dyadic deletions occur in disjoint ranges above $2X$, we have

$$A \cap [X/2, 2X) \;\subset\; B_X^{\mathrm{buf}}.$$

Thus the assumed collision cannot occur in A, a contradiction.

Hence no collision exists in A, and (1) holds. □

References

[1] T. Bloom, Erdős Problem #421 `https://www.erdosproblems.com/421`

[2] E. Bombieri and J. Pila, The number of integral points on arcs and ovals, *Duke Math. J.* **59** (1989), no. 2, 337–357. (link)

[3] W. Castryck, R. Cluckers, P. Dittmann, and K. H. Nguyen, The dimension growth conjecture, polynomial in the degree and without logarithmic factors, *Algebra Number Theory* **14** (2020), no. 8, 2261–2294. (arXiv:1904.13109)

[4] R. A. Moser and G. Tardos, A constructive proof of the general Lovász local lemma, *J. ACM* **57** (2010), no. 2, Art. 11. (arXiv:0903.0544)

[5] Y. F. Bilu and R. F. Tichy, The Diophantine equation $f(x) = g(y)$, *Acta Arith.* **95** (2000), no. 3, 261–288. (pdf)

[6] L. Hajdu and R. Tijdeman, The Diophantine equation $f(x) = g(y)$ for polynomials with simple rational roots, preprint (arXiv:2204.12345). (arXiv:2204.12345)

[7] A. Hildebrand, Integers without large prime factors, *J. Théor. Nombres Bordeaux* **5** (1993), 411–484. (eudml)