ERDŐS PROBLEM 278: FORMULAS AND COMPLEXITY OF MAXIMUM COVERED DENSITY

MICHAEL SCHROEDER

Abstract. Given distinct moduli, how much of the integers can one cover by choosing one residue class for each? When classes with noncoprime moduli can be made disjoint, a signed independence-polynomial formula gives the maximum. An optimal residue vector is explicit if every prime p divides at most p moduli. Nevertheless, evaluating this formula is #P-hard for polynomially bounded squarefree three-prime moduli with a cubic noncoprime graph. Complete-covering decision is NP-hard for distinct moduli; threshold decision is deterministically NP-complete even on prime–modulus incidence trees. In contrast, exact optimization and optimizer recovery are fixed-parameter tractable in noncoprime-graph treewidth, without supplied factorizations. We also disprove an unrestricted extension of a coprime inclusion–exclusion bound. Saturated local trees and convex compression provide exact representations. A scoped Lean companion verifies structural results for supplied coprime coordinates; the complexity classifications and specified arithmetic extensions remain written proofs.

Date: 21 September 2026.

2020 *Mathematics Subject Classification.* Primary 11B25; Secondary 11A07, 68Q17, 68V20, 90B35.

Key words and phrases. Erdős problem 278, covering systems, extremal density, Chinese remainder theorem, Shearer polynomial, exact counting, fixed-parameter tractability, pinwheel scheduling, periodic maintenance, Lean, formal verification.

Independent researcher, Singapore. ORCID: 0009-0004-3249-0195.

Reserved DOI: 10.5281/zenodo.22874487.

1. The problem and the main results

Let $A = \{n_1, \dots, n_r\}$ be a finite set of distinct integers $n_i \geq 2$. For a residue vector $a = (a_1, \dots, a_r)$, write

$$\mathcal{U}_A(a) = \bigcup_{i=1}^{r} \{x \in \mathbb{Z} : x \equiv a_i \pmod{n_i}\}, \qquad \delta_A(a) = d(\mathcal{U}_A(a)),$$

and define

$$M(A) = \max_a \delta_A(a), \qquad U(A) = 1 - M(A). \tag{1.1}$$

An *optimizer* is a residue vector attaining $M(A)$. Periodicity ensures that the density exists. Erdős and Graham ask for $U(A)$ on page 28 of [12], now problem 278 [5]. The other half, maximizing missed density by aligning residues, is settled by Rogers' theorem [17, 39]. We study only $M(A)$, or equivalently $U(A)$.

For prescribed moduli, a failed covering search leaves a basic question: were the residues poorly chosen, or is coverage impossible? Computing $M(A)$ answers this and quantifies the best partial covering. Integer-programming searches already address both tasks [18, Section 6]. Coprime moduli force intersections; shared primes permit separation by residue digits.

Our organizing question is therefore: *how do arithmetic overlaps control the difficulty of choosing classes and evaluating the density they cover?* These are different tasks. Even when an optimizer is supplied, computing its exact density can be counting-hard. Conversely, suitable structure in the overlap graph makes both tasks tractable.

A running example. For $A = \{2, 4, 6\}$, the classes 1 (mod 2), 0 (mod 4) and 0 (mod 6) cover 5/6 of the integers, which is optimal (Example 7.2). We follow this same instance from its divisibility graphs to its residue-digit trees and finally through exact evaluation.

Three views of the same arithmetic. The *noncoprime graph* G_A has one vertex per modulus and joins pairs with gcd greater than one. The *prime–modulus incidence graph* connects each modulus to its prime divisors. The *prime co-occurrence graph* H_P joins two primes when they divide a common modulus. Figure 1 distinguishes the vertices and edges before these graphs enter the complexity statements. Call a residue vector *coprime-disjoint* (CD) when every pair of classes whose moduli share a prime is disjoint. Such vectors, when they exist, attain a particularly simple inclusion–exclusion formula.

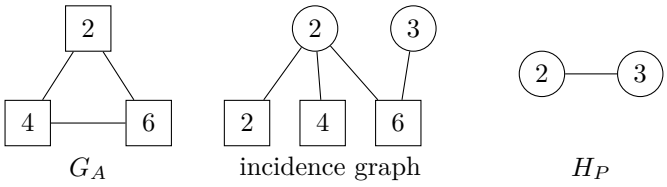

Figure 1. Three graphs for $\{2, 4, 6\}$. Rectangles are moduli; circles are primes. The noncoprime graph is a triangle, the incidence graph is a tree, and the prime co-occurrence graph is a single edge. Prime powers affect the local trees, but do not add edges to these graphs.

1.1. **Main results.**

Theorem A (The coprime-disjoint formula)**.** If A admits a coprime-disjoint residue vector, every such vector is optimal and

$$M(A) = 1 - \sum_{\substack{J \subseteq [r] \\ J \text{ independent in } G_A}} (-1)^{|J|} \prod_{j \in J} \frac{1}{n_j}.$$

In particular, this holds if every prime p divides at most p moduli. An optimizer is then obtained by assigning distinct residues modulo p to the moduli divisible by p, lifting to their prime powers and applying CRT.

Theorem B (Counting despite an explicit optimizer)**.** Exact evaluation is #P-hard, and supplied-vector hole counting is #P-complete, under polynomial-time Turing reductions. Hardness persists on polynomially bounded squarefree three-prime moduli with cubic G_A, every prime incident to exactly two moduli, and an explicit CD optimizer (Theorems 10.2 and 11.1).

Theorem 3.3 combines Adenwalla's CD density identity [1, Lemma 2.1] with a Shearer comparison proving optimality; Theorem 8.1 gives the construction. Formula availability and counting hardness overlap: the same formula applies to Theorem B's hard family. It removes the search over residue vectors, not the counting difficulty.

Theorem C (Distinct-modulus covering)**.** For general distinct moduli, deciding $U(A) = 0$ is NP-hard under deterministic polynomial-time many-one reductions, even with all moduli even and at most two of each positive 2-adic valuation.

Theorem D (Threshold hardness on incidence trees)**.** For pairwise-coprime integers $m_i > 1$ coprime to 3, deciding whether $U(\{3\} \cup \{3m_i\}) < \theta$ is NP-complete under deterministic polynomial-time many-one reductions. The rational threshold is encoded in binary, and the prime–modulus incidence graph is a tree. With distinct prime cofactors $m_i > 3$, the decision remains NP-hard under ZPP many-one reductions.

Theorem E (Noncoprime-graph tractability)**.** Without supplied factorizations, exact evaluation and optimizer recovery take $g(\Delta)2^{\tau+1}\,\mathrm{poly}(b)$ bit operations given a width-τ decomposition of G_A. Here b is the binary input length, $\Delta \le \tau + 1$ is the maximum prime incidence, and g is explicit in (12.2). Thus the problem is FPT in $\mathrm{tw}(G_A)$. Supplied-vector evaluation takes $2^{\tau+1}\,\mathrm{poly}(b)$ time.

Saturation provides finite local trees containing an optimizer (Theorem 4.2), and Proposition 6.2 gives safe convex compression. Across different arithmetic completions, however, every squarefree local partition can be required (Theorem 13.2). Corollary 13.6 constructs such witnesses in polynomial time when the private factors may be composite and pairwise coprime.

Which graph controls which task? The three graphs in Figure 1 record different kinds of interaction. Their algorithmic roles are correspondingly different:

Restriction	Exact task and conclusion
$\mathrm{tw}(G_A) = \tau$	Optimization and optimizer recovery are FPT in τ, without supplied factorizations (Theorem 12.3).
Prime–modulus incidence tree	Coprime-star threshold decision is deterministically NP-complete (Theorem 14.4).
H_P a star	Prime-star threshold decision is NP-hard under ZPP many-one reductions (Theorem 14.6).
Cubic G_A, $\Delta = 2$	Value evaluation is #P-hard under Turing reductions, even with a CD optimizer (Theorem 10.2).

The hard stars have complete G_A. Also, CD recognition is FPT in Δ alone (Corollary 12.4), whereas exact value evaluation is not unless FP = #P (Corollary 12.6). Small decision spaces do not remove the counting difficulty.

How the arguments fit together. Section 3 proves the CD formula; Section 8 gives its zero-overload construction. Saturation in Section 4 replaces general residue choices by finite local trees. From this common representation, Sections 5–7 combine local probability laws and compress the choice sets; Section 12 instead evaluates each choice directly on G_A, yielding the treewidth bound. Section 13 identifies a limit of completion-independent compression. Sections 9–11 and 14 show why choosing an optimizer and counting its covered points each require complexity analysis. The elementary star reduction can be read from Section 2, Lemma 13.5 and Section 14.

1.2. **Position within the literature.** Filaseta et al. study uncovered-density estimates [13]. Cambie gives the prime-cofactor balanced-partition formulas used in Section 14, a comparison-model obstruction, and covering hardness for binary-encoded *lists* [7]; his Remark 5 leaves distinct moduli untreated. Exact list-covering hardness is due to Kawamura, Kobayashi and Kusano [23]. Kleinberg and Mishra prove NP-completeness of dense scheduling [24, Corollary 5.1]; Kobayashi, Lin and Swernofsky strengthen dense hardness to unary input [25, Theorem 1.1]. At reciprocal density one, covering and exact packing coincide. These give complementary sources for our hole-preserving transfer to distinct moduli (Theorem 9.7). Section 14 gives separate threshold reductions with sparse incidence. Proposition 3.7 answers the unrestricted-extension part of [18, Problem 4] negatively.

Pairwise-disjoint congruence classes connect this problem to periodic maintenance and Exact Pinwheel Packing [9, 33, 40, 42, 28, 29]. Unary exact-packing hardness is attributed to Mok et al. [32, Theorem 1] in [25, Section 1]; our strong-hardness corollary uses the accessible dense theorem directly. Kobayashi and Lin give the residue/gcd formulation and a graph-colouring reduction [27, Section 1.2 and Theorem 6]; the journal version has Swernofsky as coauthor [26]. Section 9 refines this packing literature by a distinctness transfer and an incidence-two construction.

Exact Pinwheel Covering and Exact Pinwheel Packing are already FPT in the number of tasks [27, Theorems 4–5]. Onishi [34, Theorem 1.1(ii) and Proposition 5.1] gives a deterministic, factorization-free algorithm for exact evaluation of $M(A)$ and optimizer recovery in $2^{O(r^2)}\operatorname{poly}(b)$ bit operations, sharpened to $2^{O(m)}\operatorname{poly}(b)$ for $m = |E(G_A)|$. Theorem 12.3 instead exploits noncoprime-graph treewidth, giving polynomial-time computation on bounded-treewidth families even when m is unbounded. These guarantees are complementary; we do not claim an improved dependence on the number of moduli.

Adenwalla's coprime-disjoint terminology and inclusion–exclusion identity supply the attainment half of our CD calculation [1, Lemma 2.1]. Jia, Li and Liu characterize CD existence for the family of all divisors of an integer greater than one [21, Theorems 1–2]; our recognition problem allows arbitrary modulus sets. The extremal comparison uses Shearer's framework [38, 37, 19]; gcd dependency graphs and Shearer-type bounds already occur in covering systems [20, Section 2]. Bucket elimination [10, 14], coprime-basis refinement [3, 4] and tree-decomposition construction [6] supply standard algorithmic machinery.

1.3. **Conventions and input model.** Put $L = \operatorname{lcm}(n_1, \dots, n_r)$. For $A = \varnothing$, set $L = 1$, $M(A) = 0$ and $U(A) = 1$; otherwise $r \geq 1$. Modulus 1, if allowed, gives $M(A) = 1$ immediately. Empty products and least common multiples are one. An unlabelled tree forgets *digit labels*, not class indices. Write $u_A(a) = 1 - \delta_A(a)$ for the uncovered density of a fixed residue vector.

The input uses ordinary binary integers: $\log L \leq \sum_i \log n_i = O(b)$. Exponents do not give a succinct encoding. Some hardness statements supply complete prime factorizations; the algorithms instead use polynomial-time gcd refinement (Lemma 4.4).

We use three named computational problems. EXACT-E278 outputs $M(A)$ as a reduced fraction. E278-UB-EQUALITY asks whether $M(A) = \sum_i 1/n_i$, equivalently whether pairwise-disjoint classes exist. Star-Threshold asks whether $U(\{3\} \cup \{3m_i\}) < \theta$, for $n \geq 1$ pairwise-coprime integers $m_i > 1$ coprime to 3. A rational input means an integer numerator and positive denominator, both in binary, not necessarily a dyadic rational. Invalid star encodings are rejected by integer and gcd checks.

Notation	Meaning
$r,\ b,\ L$	Number of moduli, input bit length, common period
$u_A(a),\ U(A),\ M(A)$	Fixed uncovered density, its minimum, maximum coverage
$Z(A,a)$	Integer hole count $Lu_A(a)$
$G_A,\ H_P$	Noncoprime graph, prime co-occurrence graph
$I_p,\ \Delta,\ \tau$	Prime scope, maximum incidence, treewidth of G_A
$\beta_j,\ J_j$	Gcd-free base and its modulus scope
$q_S(A)$	Signed independence polynomial on S
$\varphi_\sigma,\ \mathcal{F}_I$	Avoidance profile and retained scope frontier

Throughout, τ denotes noncoprime-graph treewidth and $q_S(A)$ the signed independence polynomial. In Sections 13–14, t denotes an arithmetic scale; other auxiliary symbols are defined locally within each argument.

A *ZPP many-one reduction* here is a zero-error (Las Vegas) randomized map with polynomial expected running time: every returned instance has the correct yes/no answer. Deterministic many-one reductions have the usual worst-case polynomial bound. Exact-value consequences use one oracle query followed by a rational comparison.

1.4. **Formal verification.** The Lean 4 companion [30, 31] verifies structural theorems for supplied coprime coordinates and correctness of an exact optimizer against the original residue objective. The complexity classifications are written proofs. Appendix B gives the precise formalization map and reproduction instructions; Section 15.1 collects the mathematical scope and open questions.

2. The finite cyclic model

Put $L = \operatorname{lcm}(n_1, \ldots, n_r)$ and work on the uniform probability space $\mathbb{Z}/L\mathbb{Z}$.

Lemma 2.1 (Period reduction)**.** *For every residue vector,*

$$\delta_A(a) = \frac{1}{L}\left|\bigcup_i \{x \bmod L : x \equiv a_i \pmod{n_i}\}\right|.$$

Proof. Each congruence class is a union of residue classes modulo L. Its finite union is L-periodic, and the density of an L-periodic set is its fraction of occupied residues in one period. □

Lemma 2.2 (Generalized CRT compatibility)**.** *For nonempty $S \subseteq [r]$, the system $x \equiv a_i \pmod{n_i}$ for $i \in S$ is soluble if and only if*

$$a_i \equiv a_j \pmod{\gcd(n_i, n_j)} \qquad (i, j \in S). \tag{2.1}$$

A soluble system is one class modulo $\operatorname{lcm}(n_i : i \in S)$ and has density the reciprocal of that least common multiple.

Proof. Necessity follows by subtraction. For sufficiency, fix a prime and choose an index with largest exponent at that prime. The pairwise conditions prescribe all smaller prime-power residues consistently with that index. Do this independently at every prime and apply the ordinary CRT. The combined modulus is the stated least common multiple. □

Call S compatible when (2.1) holds, with the empty set compatible by convention.

Proposition 2.3 (Compatibility expansion)**.**

$$\delta_A(a) = \sum_{\varnothing \neq S \subseteq [r]} \frac{(-1)^{|S|+1} \mathbf{1}[S \text{ compatible}]}{\operatorname{lcm}(n_i : i \in S)}, \tag{2.2}$$

$$u_A(a) = \sum_{S \subseteq [r]} \frac{(-1)^{|S|} \mathbf{1}[S \text{ compatible}]}{\operatorname{lcm}(n_i : i \in S)}. \tag{2.3}$$

Proof. Apply finite inclusion–exclusion and Lemma 2.2 to every intersection. □

Independent uniform choices of the residues, followed by averaging, and the union bound give

$$1 - \prod_i \left(1 - \frac{1}{n_i}\right) \le M(A) \le \min\left(1, \sum_i \frac{1}{n_i}\right). \tag{2.4}$$

The affine action $x \mapsto ux + t$, with $\gcd(u, L) = 1$, preserves density. In particular translation permits $a_1 = 0$ in exhaustive tests; no common-factor assumption is needed for that normalization.

3. The noncoprime graph and the Shearer–CD theorem

Let G_A have vertex set $[r]$ and edges ij exactly when $\gcd(n_i, n_j) > 1$. Define

$$q_S(A) = \sum_{\substack{J \subseteq S \\ J \text{ independent in } G_A}} (-1)^{|J|} \prod_{j \in J} \frac{1}{n_j}. \tag{3.1}$$

A residue vector is *coprime-disjoint* (CD) if

$$a_i \not\equiv a_j \pmod{\gcd(n_i, n_j)} \qquad (ij \in E(G_A)).$$

Following [1], the name records that an intersecting subfamily can contain only pairwise-coprime moduli.

Lemma 3.1 (Lopsided Shearer bound)**.** *Let E_i be events of probabilities p_i, with graph G such that*

$$\Pr\left(E_i \,\middle|\, \bigcap_{j \in J} \overline{E}_j\right) \le p_i$$

for every set J of nonneighbors of i, whenever conditioning has positive probability. Set

$$q_S = \sum_{\substack{J \subseteq S \\ J \text{ independent in } G}} (-1)^{|J|} \prod_{j \in J} p_j.$$

If all $q_S \ge 0$, then $\Pr(\bigcap_i \overline{E}_i) \ge q_{[r]}$.

Proof. We give the ratio proof, including the boundary step, in the framework of [19, 38, 37]. Write $P_S = \Pr(\bigcap_{j \in S} \overline{E}_j)$ and $P_\varnothing = q_\varnothing = 1$. The deletion recurrence is

$$q_S = q_{S \setminus \{i\}} - p_i q_{S \setminus (\{i\} \cup \Gamma(i))}.$$

First suppose every $q_S > 0$. Induct on $|S|$, maintaining $P_S > 0$ and

$$\frac{P_S}{P_{S \setminus \{i\}}} \ge \frac{q_S}{q_{S \setminus \{i\}}}. \tag{3.2}$$

Put $T = S \setminus (\{i\} \cup \Gamma(i))$ and $R = S \setminus \{i\}$. By induction, all probabilities and q-values on proper subsets of S are positive. The lopsided hypothesis gives

$$\Pr\left(E_i \cap \bigcap_{j \in R} \overline{E}_j\right) \le \Pr\left(E_i \cap \bigcap_{j \in T} \overline{E}_j\right) \le p_i P_T.$$

Adjoin the vertices of $R \setminus T$ one at a time and apply the already proved ratios on subsets of R. This yields $P_R/P_T \geq q_R/q_T$. Therefore

$$\frac{P_S}{P_R} \geq 1 - p_i \frac{P_T}{P_R} \geq 1 - p_i \frac{q_T}{q_R} = \frac{q_S}{q_R} > 0.$$

This establishes the induction. Multiplying ratios along an ordering gives $P_S \geq q_S$.

For the boundary, if $0 < t < 1$, direct expansion gives

$$q_S(tp) = \sum_{R \subseteq S} t^{|R|}(1-t)^{|S \setminus R|} q_R(p) > 0.$$

Indeed the coefficient of each independent monomial on both sides agrees, and the term $R = \varnothing$ is strictly positive. Independently retain each event with probability t: let $E_i^{(t)} = E_i \cap \{F_i^{(t)} = 1\}$. To verify the same lopsided graph, condition on the thinning variables on a nonneighbor set J. Avoiding the thinned J-events requires avoiding precisely those original events whose thinning variables equal one. The original conditional inequality applies to that subset; averaging preserves it, and the independent thinning of i contributes a factor t. Thus the strict argument applies to tp. Couple $F_i^{(t)} = \mathbf{1}[V_i \leq t]$ with independent uniform V_i. As $t \uparrow 1$, continuity of probabilities and of the polynomial gives the result. □

Theorem 3.2 (Residue-free Shearer bound)**.** *If all $q_S(A) \geq 0$, then every residue vector satisfies*

$$u_A(a) \geq q_{[r]}(A), \qquad M(A) \leq 1 - q_{[r]}(A).$$

Proof. If J consists of nonneighbors of i, then n_i is coprime to $\operatorname{lcm}(n_j : j \in J)$. CRT makes the event $x \equiv a_i \pmod{n_i}$ independent of every Boolean combination of the J-events. The hypotheses of Lemma 3.1 hold with $p_i = 1/n_i$. □

Theorem 3.3 (CD exactness and strict equality)**.** *If A admits a CD vector, then every $q_S(A) \geq 0$ and*

$$M(A) = 1 - q_{[r]}(A).$$

If all $q_S(A) > 0$, a residue vector attains equality in the Shearer bound if and only if it is CD.

Proof. For a CD vector, every subfamily containing an edge has empty intersection. Every independent subfamily intersects by CRT and has density $\prod_{j \in J} 1/n_j$. Inclusion–exclusion gives uncovered density $q_{[r]}(A)$. Restricting the same vector to any S realizes $q_S(A)$ as an actual probability. Theorem 3.2 then proves optimality. This attainment calculation is the identity in [1, Lemma 2.1]; the extremal comparison is the additional step.

For necessity in the strict regime, suppose $P_{[r]} = q_{[r]}$. All ratio lower bounds (3.2) are positive. Their product is an equality, so each ratio in any chosen ordering is an equality. For an arbitrary edge ij, choose an ordering beginning with i, j. It follows that $P_{\{i,j\}} = q_{\{i,j\}}$. But

$$P_{\{i,j\}} = 1 - p_i - p_j + \Pr(E_i \cap E_j), \qquad q_{\{i,j\}} = 1 - p_i - p_j.$$

Hence $\Pr(E_i \cap E_j) = 0$. A nonempty congruence intersection has positive density, so this edge is disjoint. Repeating for every edge proves CD. Sufficiency was already established. □

Corollary 3.4. *If G_A is a forest, a CD vector exists and the CD formula is exact.*

Proof. Root each component. Choose a root residue arbitrarily. For each child, choose a residue different from its parent's modulo their nontrivial gcd. Such a residue exists, and these are all graph-edge constraints. □

Corollary 3.5. *For pairwise-coprime moduli, $M(A) = 1 - \prod_i (1 - 1/n_i)$, independently of the residues.*

Proof. CRT makes all class events mutually independent. □

Example 3.6 (CD beyond zero overload)**.** For $A = \{2, 4, 8\}$, the classes 1 (mod 2), 2 (mod 4) and 0 (mod 8) are pairwise disjoint. Thus $M(A) = 7/8$, although the prime 2 is incident to three moduli. Zero overload is sufficient, not necessary, for the CD formula.

Proposition 3.7 (The full polynomial alone is not a bound)**.** *There is a set of 75 distinct moduli and a residue vector a for which $0 < u_A(a) < q_{[r]}(A)$, with no pairwise-coprime triple of moduli. Consequently the proposed unrestricted covered-density bound $C_1 - C_2 + C_3$ fails, where C_k is the sum of reciprocal products over pairwise-coprime k-subsets.*

Proof. Take the disjoint blocks

$$A_0 = \{2^j : 1 \le j \le 16\} \cup \{254\}, \qquad B_0 = \{3(2j+1) : 0 \le j \le 57\}.$$

Every within-block pair is noncoprime, whereas every cross-block pair is coprime: every element of B_0 is odd, and its cofactor $2j + 1 \le 115 < 127$ is not divisible by 127. Put $x = \sum_{n \in A_0} 1/n$ and $y = \sum_{n \in B_0} 1/n$. Then $C_1 = x + y$, $C_2 = xy$ and $C_k = 0$ for $k \ge 3$. Moreover $x - 1 = 1/254 - 1/65536 > 1/256$, and the integer certificate

$$\sum_{j=0}^{57} \left\lfloor \frac{2^{20}}{3(2j+1)} \right\rfloor = 1052742 > 1052672 = 2^{20}(1 + 1/256)$$

gives $y - 1 > 1/256$. Hence $q_{[r]} = (1 - x)(1 - y) > 1/65536$.

Assign residues $2^{j-1} - 1$ modulo 2^j, residue 253 modulo 254, and zero on B_0. The dyadic classes leave only -1 (mod 65536); the class modulo 254 removes $1/127$ of that fibre. The B_0 classes have union 0 (mod 3). Their periods are coprime to the first block's period, so CRT gives

$$u_A(a) = \frac{1}{65536}\frac{126}{127}\frac{2}{3} = \frac{21}{2080768} < \frac{1}{65536}.$$

The induced polynomials $1 - x$ and $1 - y$ are negative, so the hypothesis of Theorem 3.2 fails. All higher odd coprime truncations coincide here. □

Harrington et al. prove the three-term upper bound when the prime support is contained in $\{2, 3, 5\}$ and explicitly extend their argument to any three primes [18, Theorem 1.9 and the paragraph after Problem 4]. The proposition answers the unrestricted-extension question in their Problem 4; it does not contradict their restricted theorem. The two-block construction needs both reciprocal sums to exceed one. If exactly one does, $(1 - x)(1 - y) \le 0$ gives a trivial uncovered-density bound. If both are at most one, the blockwise union bounds and CRT give $u_A(a) \ge (1 - x)(1 - y)$. The 75 moduli provide explicit positive margins, not a claim of smallest possible size.

4. p-adic partition trees

The aim is to retain the choices that affect overlap and forget the names of the digits. A residue modulo p^e specifies a path of e digits, read from least to most significant. Classes with a common initial path share a node; their next digits determine its children.

The running example as a tree. Index $2, 4, 6$ by $1, 2, 3$. At the prime 2 their required depths are $(1, 2, 1)$. For the residue vector $(1, 0, 0)$, index 1 takes first digit 1, while indices 2 and 3 take first digit 0. Indices 1 and 3 then stop: their moduli contain only one factor of 2. Index 2 needs a second digit, chosen as 0. Figure 2 records these choices. The root uses both available digits, and the one continuing index uses one child. This is the *saturation* condition: use as many distinct next digits as the live indices and the base permit.

Formally, fix $p \mid L$, and write $v_i(p) = v_p(n_i)$, $I_{p,h} = \{i : v_i(p) \ge h\}$ and $I_p = I_{p,1}$. The maximum prime incidence is $\Delta = \max_{p|L} |I_p|$, with empty maximum zero. At digit h, a live node $B \subseteq I_{p,h}$ consists of indices sharing their first $h - 1$ digits. Its children partition B by the next digit. Below a child D, only $D^+ = D \cap I_{p,h+1}$ continues; indices of exponent exactly h stop. A node is *saturated* when it has exactly $\min(p, |B|)$ nonempty children. The nested-divisibility geometry has scheduling precedents [29, 11], but the following fibre argument

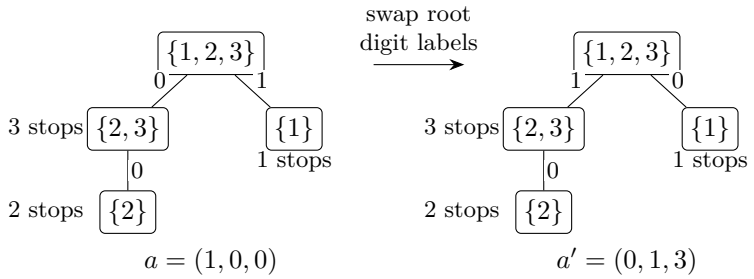

FIGURE 2. One partition tree, two labellings, for $\{2, 4, 6\}$ at the prime 2. Nodes record class indices; edges record digits, least significant first. Only occupied children are drawn. Keeping the residue of class 3 equal to zero modulo 3 gives the displayed CRT vectors. Both have the same membership-mask law and covered density.

applies to arbitrary prime powers and arbitrary incidence cycles. For the running example, the three saturated root partitions are 12|3, 13|2 and 23|1. Here 23|1 means that indices 2 and 3 share a first digit different from index 1's. Example 7.2 evaluates all three choices.

Lemma 4.1 (Fibre splitting)**.** *For measurable C, X, Y in a finite probability space,*

$$\mu(C \cup Y) + \mu(C \cup X) - \mu(C \cup X \cup Y) - \mu(C) = \mu((X \setminus C) \cap (Y \setminus C)) \geq 0.$$

Proof. Subtract the common baseline C and apply the two-set union identity. □

Theorem 4.2 (Global p-adic saturation)**.** *For every finite A, some optimal residue vector has a saturated local tree at every prime and every live node.*

Proof. Start with a global optimizer, which exists because the residue-vector space is finite. Fix one prime and process nodes from shallow digits to deep digits. At a node of depth h, freeze every earlier digit, all other prime coordinates, all deeper digits, and all data outside the node. Conditional on its parent prefix, the ambient space is $\mathbb{Z}/p\mathbb{Z} \times \Omega$ with uniform first coordinate.

A class that stopped before digit h is independent of the current digit inside this prefix. A class still alive with an incompatible earlier prefix contributes nothing there. A class not divisible by p is also independent of this digit. Thus all outside contributions form the same baseline $C \subseteq \Omega$ in every child fibre. A live class contributes $\{c_i\} \times X_i$, with its deeper and other-prime restrictions contained in X_i.

If fewer than $\min(p, |B|)$ child digits are occupied, some occupied digit contains an index i and a nonempty collection of other indices, and some digit is unused. Let Y be the union of that collection's remaining-coordinate sets. Moving only the current digit of i to the unused digit changes the two affected fibre coverages from

$$\mu(C \cup X_i \cup Y) + \mu(C) \quad \text{to} \quad \mu(C \cup Y) + \mu(C \cup X_i).$$

Lemma 4.1 makes the change nonnegative; the omitted ambient factor $1/p$ is common and positive. Repeat until the node is saturated. Later changes at deeper digits do not alter earlier child partitions. Changes at another prime do not alter any digit at this prime. There are finitely many primes and digits, so the process terminates. Coverage never decreases from its global maximum, and hence remains maximal. □

Let $\mathcal{S}_p(A)$ be the saturated local trees with class indices retained and digit labels forgotten. Recursively, at a live node,

$$\mathcal{S}_p(B, h) = \coprod_{\substack{\pi \text{ a partition of } B \\ |\pi| = \min(p, |B|)}} \prod_{D \in \pi} \mathcal{S}_p(D \cap I_{p,h+1}, h + 1).$$

The empty or singleton continuation has one state; for a singleton its original exponent remains stored, and its unrecorded digits may be filled with zeros.

Corollary 4.3 (Saturated prime-power kernel)**.** *Choose one $\sigma_p \in \mathcal{S}_p(A)$ for each prime. Label each set of children injectively by digits and combine the resulting residues by CRT. If $a(\sigma)$ is any such realization, then*

$$M(A) = \max_{\sigma \in \prod_p \mathcal{S}_p(A)} \delta_A(a(\sigma)).$$

Proof. First consider the relabelling in Figure 2. Swapping the root digits sends $z_0 + 2z_1$ to $(1 - z_0) + 2z_1$ modulo 4. This is a permutation of the four equally likely residues. It sends every selected prefix cylinder to the corresponding relabelled cylinder, including the shorter cylinders of classes 1 and 3. With the prime-3 coordinate fixed, CRT gives the vectors $(1, 0, 0)$ and $(0, 1, 3)$; their covered points are permuted, not changed in number.

The general argument is the same at every node. Any two injective labellings of its occupied children extend to a permutation of all p children. Choose these permutations successively from the root to depth $\max_i v_i(p)$, including singleton tails. They define a bijection of the full p-ary tree that takes each selected prefix cylinder to its relabelled cylinder. All leaves have equal probability, so this bijection preserves the membership-mask law. Independent such bijections at the different primes combine by CRT and preserve global density.

Every state product has a valid CRT realization. Conversely Theorem 4.2 supplies an optimizer with one such product, and the relabelling argument allows its canonical representative. This proves the equality. □

Lemma 4.4 (Gcd-free coordinates)**.** *From binary moduli one can compute, in polynomial time, pairwise-coprime integers $\beta_1, \dots, \beta_s > 1$ and nonnegative exponents e_{ij} such that*

$$n_i = \prod_{j=1}^{s} \beta_j^{e_{ij}}.$$

Put $E_j = \max_i e_{ij}$ and $J_j = \{i : e_{ij} > 0\}$. CRT identifies the period space with $\prod_j \mathbb{Z}/\beta_j^{E_j}\mathbb{Z}$. Every J_j is a clique of G_A, and $\max_j |J_j| = \Delta$.

Proof. Coprime-basis refinement is classical [3, 4]; the following elementary version suffices. Start with the set of distinct input integers. Whenever distinct current entries x, y have $h = \gcd(x, y) > 1$, replace them by $h, x/h, y/h$, omitting ones and merging equal entries. Every original modulus remains a product of current entries. The product of the distinct current entries decreases by a factor of at least two, so there are at most $\sum_i \log_2 n_i$ refinements. The number and bit lengths of all entries are bounded by this same sum. Pairwise gcd searches therefore take polynomial time. At termination the entries are pairwise coprime; repeated exact division recovers the exponents.

For $p \mid \beta_j$ one has $v_p(n_i) = v_p(\beta_j)e_{ij}$. Thus p divides precisely the moduli indexed by J_j. This proves both the clique assertion and equality of the two maximum incidences. The bases themselves need not be prime or squarefree.

□

Proposition 4.5 (Transport to composite bases)**.** *For the coordinates of Lemma 4.4, the following assertions hold.*

(i) Some optimizer has a saturated tree in every base β_j: a live set B occupies $\min(\beta_j, |B|)$ children.

(ii) Forgetting digit labels preserves the membership-mask law. Canonical labels followed by CRT realize every state product and recover an optimizer from an optimal one.

(iii) The local cylinders are laminar; a coordinate of incidence d has at most $d + 1$ positive membership atoms. Independent coordinate laws support the same tensor, fusion and convex-pruning arguments as prime coordinates.

(iv) There are at most $(d+1)^{\binom{d}{2}}$ *saturated unlabelled states in that coordinate, and one if* $\beta_j \geq d$*. At most* $\pi(\Delta-1)$ *coordinates have a choice of state. These states and their laws can be generated without factoring the bases or iterating through their numerical digit alphabets.*

Proof. For any integer $\beta > 1$, the map

$$(z_0, \ldots, z_{E-1}) \longmapsto \sum_{h=0}^{E-1} z_h \beta^h$$

is a bijection from $\{0, \ldots, \beta-1\}^E$ to $\mathbb{Z}/\beta^E\mathbb{Z}$. Thus a uniform coordinate has independent uniform digits. A congruence modulo β^e fixes precisely its first e digits, regardless of the prime factors of β.

At a live node, classes that have stopped give a common baseline in every child; incompatible prefixes contribute nothing. Moving one class from a shared child to an unused child changes coverage by the nonnegative fibre difference in Lemma 4.1, now scaled by $1/\beta$. Processing shallow digits first proves (i), exactly as in Theorem 4.2. Permuting the β children independently at each prefix gives a measure-preserving permutation of all β^E leaves. Any two injective labels of the occupied children extend to such permutations, proving (ii).

Two prefix cylinders are nested or disjoint. For each distinct selected cylinder, subtract its maximal proper selected subcylinders; together with the complement of the top cylinders these pieces partition the coordinate. Each piece has constant membership mask, and its mass is obtained by subtracting the child masses β^{-e} from its parent mass. There are at most $d+1$ pieces. CRT makes the different base coordinates independent. The tensor and fusion identities use only that independence, while pruning uses only nonnegative outside coefficients. This proves (iii).

For (iv), while two live indices stay together through a saturated split, their common child has fewer live indices than its parent. They therefore separate by digit $d-1$, or one expires first. Their first-separation digits, with one symbol for no separation before expiry, determine the indexed tree from the fixed exponents. The pair-code argument of Lemma 12.1 gives the stated bound. A base at least d separates all indices immediately. Bases with choices are less than Δ; their least prime divisors are distinct because the bases are pairwise coprime, so there are at most $\pi(\Delta-1)$.

Generate partitions of the live index sets, recording only occupied children and storing singleton tails by their exponents. Assign the occupied digits $0, 1, \ldots$; neither generation nor atom computation enumerates β children or β^E leaves. All stored exponents, residues and rational masses have polynomial bit length in the input. This supplies the composite-base input to the bit-complexity analysis in Section 12. □

5. The nonnegative membership-mask tensor

We now evaluate a fixed collection of local trees. At each prime, the only information needed about a sampled residue is *which classes it hits.* These random subsets are membership masks. Independence of prime coordinates then turns global coverage into a finite calculation with their local laws. For a fixed local state, take X_p uniform modulo p^{e_p}, where $e_p = \max_i v_i(p)$, and define

$$B_p = \{i \in I_{p,1} : X_p \equiv a_i \pmod{p^{v_i(p)}}\}, \qquad \mu_{p,\sigma_p}(B) = \Pr(B_p = B).$$

The preceding tree automorphism argument makes this law independent of the chosen digit labelling. CRT makes the laws at different primes independent. A point belongs to class i precisely when $i \in B_p$ at every prime dividing n_i.

Theorem 5.1 (Exact nonnegative tensor formula)**.**

$$U(A) = \min_{\sigma \in \prod_p \mathcal{S}_p(A)} \sum_{(B_p)} \left(\prod_p \mu_{p,\sigma_p}(B_p) \right) \prod_{i=1}^{r} \left(1 - \prod_{p \mid n_i} \mathbf{1}[i \in B_p] \right). \tag{5.1}$$

All entries are nonnegative exact rationals. Prime coordinates use supplied factorizations; the corresponding gcd-free-coordinate formula is computable directly from binary moduli by Lemma 4.4 and Proposition 4.5.

Proof. For fixed σ, the product of the local laws is the distribution of the independent prime masks. The last product is exactly the indicator of missing every class. Its expectation is $u_A(a(\sigma))$. Apply Corollary 4.3. □

In general $\sum_x \min_s f(s,x) \neq \min_s \sum_x f(s,x)$: a residue decision cannot depend on the sampled CRT point. Sum chance variables before minimizing decisions, or enumerate complete decisions and evaluate each expectation as in Section 12.

Lemma 5.2 (Sparse fixed-state atoms)**.** *For a prime incident to d classes, any fixed local residue vector has at most $d+1$ positive membership atoms. Its law can be computed using the inclusion tree of its distinct cylinders, without enumerating p^{e_p} points. Every atom count is an integer over denominator p^{e_p}.*

Proof. Prime-power cylinders are nested or disjoint. Merge equal cylinders, retaining their index labels, and add the full space as root. The inclusion tree has at most d nonroot nodes. Subtract each node's maximal proper selected subcylinders. These remainders partition the space and have fixed masks: the remainder of D belongs exactly to the cylinders containing D. Omitting zero remainders leaves at most $d+1$ atoms. For a node of depth v with child depths v_j, its probability is

$$p^{-v} - \sum_j p^{-v_j};$$

Use $v = 0$ at the root. Multiplying by p^{e_p} gives integer counts. Exponent comparisons and congruence tests determine containment in polynomial time. □

The bound is for one fixed state; the union over all states can have 2^d masks.

6. Scope fusion and avoidance profiles

The next step reduces repeated work in this representation. Primes involving the same set of classes can be combined into one probability law. We then compare candidate laws by how they behave under every possible choice of the remaining coordinates. This route compresses the decision space; the separate treewidth route in Section 12 evaluates decisions directly. Put $I_p = \{i : p \mid n_i\}$, $\mathcal{I}(A) = \{I_p : p \mid L\}$ and $\Gamma_I = \{p : I_p = I\}$. Outside this group, a state $\sigma = (\sigma_p)_{p \in \Gamma_I}$ is visible only through

$$C_\sigma = \bigcap_{p \in \Gamma_I} B_p.$$

For $S \subseteq I$ set $z_{p,\sigma_p}(S) = \Pr(S \subseteq B_p)$. Independence, subset inversion and inclusion–exclusion give respectively

$$z_\sigma(S) = \prod_{p \in \Gamma_I} z_{p,\sigma_p}(S), \tag{6.1}$$

$$\nu_\sigma(C) = \sum_{S \supseteq C} (-1)^{|S \setminus C|} z_\sigma(S), \tag{6.2}$$

$$\varphi_\sigma(T) = \Pr(C_\sigma \cap T = \varnothing) = \sum_{S \subseteq T} (-1)^{|S|} z_\sigma(S). \tag{6.3}$$

The coordinate $\varphi_\sigma(T)$ is the probability that the fused mask avoids every class indexed by T. These transforms are invertible, so equal profiles mean equal laws. Here an *outside*

completion ξ means a choice of local states at the primes outside Γ_I, with A fixed. Those outside coordinates determine which subsets T contribute to global avoidance. The convex test below applies mixed-strategy domination [8] uniformly over these choices.

Theorem 6.1 (Completion-safe convex domination)**.** *Fix every state outside Γ_I, denoting that choice by ξ. There are coefficients $\lambda_\xi(T) \geq 0$ such that*

$$u_A(\sigma, \xi) = \sum_{T \subseteq I} \lambda_\xi(T) \varphi_\sigma(T). \tag{6.4}$$

Merge equal profiles. A target φ_ρ can then be deleted if

$$\varphi_\rho \geq \sum_{j=1}^{k} \alpha_j \varphi_j \quad \textit{coordinatewise}, \qquad \alpha_j \geq 0, \quad \sum_j \alpha_j = 1, \tag{6.5}$$

using only different, currently available *profiles. The optimum and an arithmetic representative are preserved.*

Proof. Condition on the outside masks. A point covered by a class not using this group contributes zero. Otherwise let $T \subseteq I$ be the indices satisfying every outside requirement; avoidance is $C_\sigma \cap T = \varnothing$. Grouping configurations by T gives (6.4), with $\lambda_\xi(T) \geq 0$ and $\sum_T \lambda_\xi(T) \leq 1$.

For (6.5), multiply coordinatewise by λ_ξ and sum:

$$u_A(\rho, \xi) \geq \sum_j \alpha_j u_A(\varphi_j, \xi) \geq \min_j u_A(\varphi_j, \xi).$$

The alternative may depend on ξ; excluding the target prevents vacuous deletion. □

For the finite set Q_I of distinct profiles, define the canonical frontier

$$\mathcal{P}_I = \operatorname{conv}(Q_I) + \mathbb{R}^{2^{|I|}}_{\geq 0}, \qquad \mathcal{F}_I = \operatorname{vert}(\mathcal{P}_I). \tag{6.6}$$

Each vertex retains one original state representative.

Proposition 6.2 (Canonical frontier and certificates)**.** *Every vertex is an original profile. A profile $\varphi \in Q_I$ is a nonvertex exactly when it is coordinatewise above a convex combination of $Q_I \setminus \{\varphi\}$. Deleting all nonvertices preserves every nonnegative linear minimum, independently of deletion order. Each retained profile is uniquely exposed by a strictly positive rational functional.*

Proof. A certificate $\varphi \geq z \in \operatorname{conv}(Q_I \setminus \{\varphi\})$ expresses φ either as a nontrivial convex combination or as the midpoint of z and $z + 2(\varphi - z)$. Conversely, expand a nontrivial midpoint representation of φ into original profiles and nonnegative slack. The averaged coefficient of φ is less than one: equality would force zero slack and both midpoint endpoints to equal φ. Rearrangement gives the certificate.

A vertex must belong to Q_I, since nonzero slack or a nontrivial profile mixture admits a midpoint decomposition. Each certified deletion leaves $\mathcal{P}_I$ unchanged. Iterating the characterization on the remaining set therefore leaves precisely its vertices, and gives retained-only certificates for every deletion.

A retained φ lies outside the closed polyhedron $\operatorname{conv}(Q_I \setminus \{\varphi\}) + \mathbb{R}^{2^{|I|}}_{\geq 0}$. Strict separation gives $\lambda \geq 0$ with $\lambda \cdot (\varphi' - \varphi) > 0$ for every alternative; the orthant forces nonnegativity. Finitely many rational strict inequalities allow a rational λ. Adding a sufficiently small positive rational all-ones vector preserves every gap and makes all coefficients positive, so upward slack also increases the value. For a singleton profile set, use the all-ones functional directly. □

A retained profile need not optimize the present completion: retention preserves every nonnegative interface objective. Section 13 realizes the exposing objectives arithmetically for its partition family, not for arbitrary interfaces.

Proposition 6.3 (A singleton-frontier criterion)**.** *In a fused scope I, let $\alpha_i = \Pr(i \in C_\sigma)$, which is independent of σ. If the candidate laws include one for which these events are pairwise disjoint, then $\mathcal{F}_I$ has one profile:*

$$\varphi_*(T) = 1 - \sum_{i \in T} \alpha_i \qquad (T \subseteq I).$$

This applies to distinct moduli $n_i = p^{u_i}\ell^{v_i}$, with p, ℓ distinct primes, $u_i, v_i \geq 1$, and $\max(p, \ell) < r \leq p\ell$. Both prime scopes are overloaded.

Proof. The union bound gives $\varphi_\sigma \geq \varphi_*$, with equality for the disjoint law; hence the upper hull is $\varphi_* + \mathbb{R}^{2^{|I|}}_{\geq 0}$.

For $p \leq \ell$, start with the ℓ first-digit pairs $(j \bmod p, j)$, $0 \leq j < \ell$, and add unused pairs until there are r. They cover every digit in both coordinates and separate every class pair. Saturating deeper digits preserves these saturated first splits and disjointness, so the law is a candidate. □

Corollary 6.4 (Exponential profile compression)**.** *For every prime $\ell \geq 3$, the family $n_i = 2\ell^i$, $1 \leq i \leq r = 2\ell$, has at least $2^{r/2}$ distinct fused profiles but a singleton frontier.*

Proof. The two prime scopes are both $[r]$, and Proposition 6.3 gives the singleton frontier. Fix a saturated ℓ-tree whose first split has the block $\{1, \ldots, \ell+1\}$ and $\ell-1$ singletons. Index 1 expires after the first digit; the remaining ℓ indices of that block separate at the second digit. At prime 2, put index 1 in digit 0 and index $\ell + 2$ in digit 1. Choose the digits of indices $2, \ldots, \ell + 1$ independently. All 2^ℓ states are saturated. For each such j, its class is contained in class 1 if their binary digits agree, and disjoint from it otherwise. Thus

$$\varphi(\{1, j\}) = 1 - \frac{1}{n_1} - \frac{1}{n_j} + \frac{\mathbf{1}[\text{the binary digits agree}]}{n_j}$$

recovers each choice. The profiles are therefore distinct. □

This is a compression bound, not an end-to-end speedup: generating all candidate profiles may still be exponential. The CD identity already gives this family's value.

Example 6.5 (Fusion with two overloaded primes)**.** For $A = \{6, 12, 18, 24\}$, primes 2 and 3 have seven and six saturated states. Their 42 products give six distinct fused profiles and one retained profile. The disjoint vector $(0, 4, 2, 9)$ gives $M(A) = 1/6+1/12+1/18+1/24 = 25/72$ by Theorem 3.3; the additional content is compression, not the value formula.

7. Exact scope-fused optimization

Combining the local laws and retained representatives gives the first exact algorithm. Its order of operations is important: evaluate the expected uncovered density of each decision before comparing decisions. The worked example below completes the calculation begun with the trees in Section 4.

Theorem 7.1 (Exact scope-fused optimization formula)**.** *For every finite A, with $\mathcal{F}_I$ defined by* (6.6)*,*

$$U(A) = \min_{(\sigma_I) \in \prod_{I \in \mathcal{I}(A)} \mathcal{F}_I} \sum_{(C_I)} \left(\prod_{I \in \mathcal{I}(A)} \nu_{I,\sigma_I}(C_I) \right) \cdot \prod_{i=1}^{r} \left(1 - \prod_{\substack{I \in \mathcal{I}(A) \\ i \in I}} \mathbf{1}[i \in C_I] \right). \tag{7.1}$$

This finite rational formula recovers an optimizer from retained representatives and CRT, and holds equally on gcd-free coordinates. Chance masks are summed before decisions are minimized.

Proof. Fuse Theorem 5.1: each class tests $i \in C_I$, and disjoint groups of prime coordinates remain independent. Apply Proposition 6.2 scope by scope. Equation (6.4) holds for *every* outside choice, including already-pruned scopes, so each replacement preserves the global minimum. Retained profiles have original state representatives, not artificial mixtures. Backtracking these representatives and applying Corollary 4.3 gives an attaining residue vector. Proposition 4.5 transports the argument to gcd-free coordinates. □

Section 12 gives a different evaluation route with a parameterized running-time bound.

Example 7.2 (The full pipeline for $A = \{2, 4, 6\}$)**.** Continue the index order and trees of Figure 2. At the prime 2, the depths are $(1, 2, 1)$ and the saturated root partitions are $12|3$, $13|2$ and $23|1$. Writing e_B for unit mass at mask B, their laws are

$$\mu_{2,12|3} = \tfrac{1}{4}e_{\{1,2\}} + \tfrac{1}{4}e_{\{1\}} + \tfrac{1}{2}e_{\{3\}},$$
$$\mu_{2,13|2} = \tfrac{1}{2}e_{\{1,3\}} + \tfrac{1}{4}e_{\{2\}} + \tfrac{1}{4}e_{\varnothing},$$
$$\mu_{2,23|1} = \tfrac{1}{4}e_{\{2,3\}} + \tfrac{1}{4}e_{\{3\}} + \tfrac{1}{2}e_{\{1\}}.$$

The prime 3 has scope $\{3\}$ and forced law $\frac{1}{3}e_{\{3\}} + \frac{2}{3}e_{\varnothing}$. Singleton avoidance coordinates at the prime 2 are $(1/2, 3/4, 1/2)$, and the remaining coordinates are

	$\varphi(\{1,2\})$	$\varphi(\{1,3\})$	$\varphi(\{2,3\})$	$\varphi(\{1,2,3\})$
$12\|3$	$1/2$	0	$1/4$	0
$13\|2$	$1/4$	$1/2$	$1/4$	$1/4$
$23\|1$	$1/4$	0	$1/2$	0

All three are retained: coordinates $(13, 23)$ force any mixture below the first profile to equal the first; $(13, 12)$ do the same for the third, and $(12, 23)$ for the second. Here ij denotes $\{i, j\}$, and the table's common minimum in each indicated coordinate forces zero weight on its larger alternatives.

The uncovered densities are respectively $1/3$, $1/4$ and $1/6$, so

$$U(\{2, 4, 6\}) = \frac{1}{6}, \qquad M(\{2, 4, 6\}) = \frac{5}{6}.$$

The chosen state $23|1$ is realized by $1 \pmod 2$, $0 \pmod 4$, $0 \pmod 6$. Only the last two classes intersect, at density $1/12$, and $1/2 + 1/4 + 1/6 - 1/12 = 5/6$ checks the result directly.

8. Zero hierarchical overload

When each prime supplies enough distinct first digits, a CD optimizer can be constructed without comparing local states. The following overload statistic makes this sufficient condition explicit. Define

$$H(A) = \sum_{p} \sum_{h \geq 1} \max(0, |I_{p,h}| - p).$$

Because the sets $I_{p,h}$ decrease with h,

$$H(A) = 0 \quad \Longleftrightarrow \quad |I_p| \leq p \text{ for every prime } p.$$

In this case every saturated tree separates its indices at the first digit. Subsequent singleton tails have one unlabelled state; this is uniqueness of the *state*, not uniqueness of the labelled residue vector.

Theorem 8.1 (Complete zero-overload formula)**.** *If $H(A) = 0$, then $M(A) = 1 - q_{[r]}(A)$. An optimizer is obtained by assigning distinct first p-adic digits to all indices in I_p, lifting to their prescribed powers and applying CRT.*

Proof. Each noncoprime pair disagrees at a shared prime's first digit, hence is disjoint. Each independent subfamily consists of pairwise-coprime moduli and intersects by CRT. The resulting vector is CD, so Theorem 3.3 applies. □

Constructing this optimizer is elementary with suitable coordinates; its exact value can still be hard to count.

Example 8.2. For $A = \{6, 10, 15\}$, G_A is a triangle and every prime has incidence two. Thus $U(A) = 1 - 1/6 - 1/10 - 1/15 = 2/3$ and $M(A) = 1/3$.

Example 8.3 (The graph formula needs a hypothesis)**.** For $\{2, 4, 6\}$, the complete-graph expression $1 - q_{[r]}$ gives $11/12$, not the actual $5/6$ (Example 7.2). Some hypothesis is necessary; zero overload suffices but is not necessary, as $\{2, 4, 8\}$ shows.

9. Packing and covering with distinct moduli

9.1. **Packing and CD recognition.** E278-UB-EQUALITY is Exact Pinwheel Packing with distinct periods: task i runs at $a_i + n_i\mathbb{Z}$, and collision is determined by the gcd. Unlike star threshold decision, packing also supplies the CD-recognition hardness used in Section 12.

Lemma 9.1 (Private-prime padding)**.** *For positive periods $n_1, \dots, n_r$, possibly repeated, choose distinct primes t_i dividing none of them. Then $N_i = n_i t_i$ are distinct,*

$$\gcd(N_i, N_j) = \gcd(n_i, n_j) \quad (i \neq j),$$

and packing feasibility is preserved. Squarefreeness and old prime incidences are preserved; new primes have incidence one. Primes $t_i > r$ can be found in polynomial time, ensure $\sum_i 1/N_i < 1$, and hence give strict Shearer inequalities when the periods share a prime.

Proof. Freshness proves the gcd and structural claims. Lift residues to N_i, or reduce them to n_i; the pairwise gcd inequalities are unchanged. Also $1/N_i < 1/r$, and a common prime makes G_A complete, with $q_S = 1 - \sum_{i \in S} 1/N_i > 0$.

A b-bit input has at most b distinct prime divisors. Among the first $b + r$ primes beyond r, at least r divide no period. These primes have polynomial numerical size by the standard bound for the kth prime, so enumeration and divisibility tests take polynomial time. □

Padding the squarefree common-3 construction of Kobayashi and Lin [27, Theorem 6] gives distinct, squarefree, strict-Shearer packing hardness. Proposition 9.3 additionally bounds all other prime incidences by two; that restriction does not follow from their construction.

Corollary 9.2 (Strong hardness of distinct packing)**.** *The unary dense-hardness theorem of the preprint* [25, Theorem 1.1] *implies that E278-UB-EQUALITY is NP-complete on unary-encoded distinct moduli, hence strongly NP-complete. Exact computation of $M(A)$ is NP-hard on these inputs.*

Proof. On an explicit unary source list, test $\sum_i 1/n_i = 1$ exactly, sending a non-dense input to the fixed no-instance $\{2, 3\}$ of E278-UB-EQUALITY. With $P = \prod_i n_i$, the test is $\sum_i P/n_i = P$ and has polynomial bit cost. At density one, feasibility is equivalent to pairwise-disjoint residue classes [25, Lemma 2.2]; see also [27, Observation 1, p. 47:3]. For a source of unary length $S = r + \sum_i n_i$, its binary length is $O(S)$. The primes in Lemma 9.1 can therefore be chosen of size $O(S \log(S + 2))$, so

$$\sum_i N_i \leq (\max_i t_i) \sum_i n_i = O(S^2 \log(S + 2)).$$

Padding is thus a polynomial-time unary reduction, preserving feasibility. A residue certificate is checked by all pairwise gcd inequalities, proving NP membership. An exact-value oracle decides equality with the rational $\sum_i 1/N_i$, giving the final assertion. □

Padding preserves the indexed gcd pattern, G_A and all old prime incidences, while adding incidence-one primes. It preserves squarefreeness *if the source has it*; the corollary does not inherit the further promises of the following proposition.

Proposition 9.3 (An incidence-two refinement of packing hardness)**.** *E278-UB-EQUALITY is NP-complete even for distinct squarefree moduli with a complete noncoprime graph, strict Shearer inequalities, exactly one overloaded prime and incidence at most two at every other prime, with factorizations supplied.*

Proof. A polynomial-bit residue certificate is checked by the pairwise tests

$$a_i \not\equiv a_j \pmod{\gcd(n_i, n_j)} \quad (i \neq j), \tag{9.1}$$

by Lemma 2.2; a nonempty intersection has positive density. Thus union-bound equality is in NP.

For a 3-colourability instance $G = (V, E)$, add isolated vertices to ensure $N \geq 4$. Choose mutually distinct private primes s_v and nonedge primes r_{uv}, all greater than $10N^2$ and of polynomial numerical size. Set

$$n_v = 3s_v \prod_{\substack{u \neq v \\ \{u,v\} \notin E}} r_{uv}. \tag{9.2}$$

Then

$$\gcd(n_u, n_v) = \begin{cases} 3, & uv \in E, \\ 3r_{uv}, & uv \notin E. \end{cases}$$

Disjoint classes yield a proper colouring $a_v \bmod 3$. Conversely, prescribe a proper colour modulo 3, residues 0 and 1 at each nonedge's endpoints, and arbitrary private residues. CRT separates edges at 3 and nonedges at r_{uv}.

Private primes make the squarefree moduli distinct. Prime 3 has incidence $N > 3$, separators have incidence two, and private primes incidence one. Thus only 3 is overloaded. Finally $\sum_v 1/n_v < N/(30N^2) < 1$, so every complete-graph q_S is positive. Prime enumeration and all output bit lengths are polynomial; factorizations are supplied. □

Corollary 9.4. *For* (9.2)*, deciding whether some saturated partition at prime 3 admits a pairwise-disjoint completion is NP-complete, although all other local states are forced. Exact computation of $M(A)$ is NP-hard under all the restrictions of Proposition 9.3.*

Proof. All nonedge endpoints separate at forced local states, so feasibility at 3 means partitioning V into three independent blocks. A colouring using fewer blocks can be split to three since $N \geq 4$. This is 3-colourability with a polynomially checked partition certificate. An exact value decides union-bound equality, proving the final assertion. □

9.2. **Removing repetitions without changing the minimum hole count.** For a nonempty indexed list $B = (d_i)$, $d_i \geq 2$, allow repeated moduli and write $L_B = \operatorname{lcm}_i d_i$, $U_{\mathrm{list}}(B) = \min_a u_B(a)$, and $Z(B, a) = L_B u_B(a)$. Repetitions cannot simply be deleted: two classes modulo 2 may cover, while one cannot. The following construction preserves their roles.

Lemma 9.5 (Encoding binary digits in odd coordinates)**.** *Write $d_i = 2^{e_i} m_i$ with m_i odd, $E = \max_i e_i$, and $L_0 = \operatorname{lcm}_i m_i$. Choose distinct odd primes $\ell_1, \ldots, \ell_E$ not dividing L_0 and put $Q_e = \prod_{j \leq e} \ell_j$, $Q_0 = 1$. Form an odd list O containing $\ell_j - 2$ copies of ℓ_j for each j, together with one entry $m_i Q_{e_i}$ for each i. Then $L_O = Q_E L_0$, and*

$$\min_b Z(O, b) = \min_a Z(B, a), \qquad U_{\mathrm{list}}(O) = \frac{2^E}{Q_E} U_{\mathrm{list}}(B).$$

A canonical lift preserves the hole count for every source vector. Every target vector admits a source vector with no more holes. Both maps and the moduli can be constructed in deterministic polynomial bit time.

Proof. Use the $\ell_j - 2$ guard copies to cover residues $2, \ldots, \ell_j - 1$. Given a_i, the data class fixes $a_i \pmod{m_i}$ and, at ℓ_j for $j \leq e_i$, its $(j-1)$st low binary digit. CRT realizes these prescriptions. The complement of the guards is $\{0,1\}^E \times \mathbb{Z}/L_0\mathbb{Z}$, in bijection with $\mathbb{Z}/(2^E L_0)\mathbb{Z}$.

A source congruence fixes precisely the corresponding initial digits and odd residue. Thus the bijection preserves uncovered points.

For arbitrary target residues, each guard family leaves at least two symbols. Choose two and label them 0 and 1. On their product, each data class is either absent or fixes an initial binary string and a residue modulo m_i. Decode the latter by CRT to a residue modulo d_i; restore absent source classes with arbitrary residues. This restoration can only cover additional points. Therefore the extracted source has at most as many holes as the target has in that selected product, and hence no more than its total hole count. The canonical lift and this inequality give equality of the minima.

For input length b, $E \le b$ and L_0 has at most b distinct prime divisors. Among the first $b + E + 1$ primes, at least E are odd and avoid L_0. These primes have polynomial numerical size, so enumeration, primality and divisibility tests, and explicitly writing every guard copy take polynomial time. All products and CRT data have polynomial bit length. No factorization or enumeration of the period is required. When $E = 0$, this is the identity map. □

Lemma 9.6 (Distinctification of an odd list)**.** *Let $O = (o_1, \dots, o_M)$ be a nonempty odd list with $o_i > 1$. Choose nonnegative heights h_i so that (o_i, h_i) are distinct, and put $T = \max_i h_i$. Then the set*

$$D = \{2, 4, \dots, 2^T\} \cup \{2^{h_i} o_i : 1 \le i \le M\}$$

has period $L_D = 2^T L_O$ and minimum hole count equal to that of O. A canonical lift preserves every source hole count; arbitrary target vectors can be decoded without increasing it. The maps take polynomial time in the explicit input and output lengths.

Proof. Distinct odd parts and heights distinguish the data moduli, none of which is a pure power of two. Assign the mask modulo 2^j the residue $2^{j-1} - 1$. The masks leave exactly $-1 \pmod{2^T}$. Lift each source class into that fibre by prescribing $-1 \pmod{2^{h_i}}$ and its original residue modulo o_i. CRT gives one target hole per source hole.

Conversely, arbitrary masks have total density $1 - 2^{-T} < 1$, so some binary fibre is unmasked. Within it, retain each active data class modulo o_i and restore inactive source classes arbitrarily. The extracted source holes inject into target holes in this fibre. It remains to find such a fibre efficiently.

Read binary words from the least significant digit. The mask modulo 2^j forbids one prefix of length j. Discard any forbidden prefix extending a shorter one, leaving a prefix antichain $\mathcal{W}$ of at most T words. Its forbidden cylinders are disjoint. For a prefix v of length at most T, let $N_{\text{free}}(v)$ count the unmasked length-T words extending it. If a word in $\mathcal{W}$ is a prefix of v, this count is zero; otherwise

$$N_{\text{free}}(v) = 2^{T-|v|} - \sum_{\substack{w \in \mathcal{W} \\ v \text{ is a prefix of } w}} 2^{T-|w|}.$$

The root has positive count by the density bound above. At each level, $N_{\text{free}}(v) = N_{\text{free}}(v0) + N_{\text{free}}(v1)$, so choose a child with positive count. After T choices the resulting word is an unmasked fibre. Each count uses at most T stored prefixes and $T + 1$-bit integers; prefix comparisons and all T choices therefore take polynomial time. For $T = 0$ the empty word already describes the sole fibre.

For example, masks $0 \pmod 2$, $1 \pmod 4$, $3 \pmod 8$ forbid the low-digit-first words $0, 10, 110$. The search follows $\varnothing \to 1 \to 11 \to 111$ and returns the fibre $7 \pmod 8$. □

Theorem 9.7 (Hole-preserving list-to-set reduction)**.** *Every nonempty binary-encoded list B of moduli at least two can be transformed in deterministic polynomial time into a set $D(B)$ of distinct moduli with*

$$\min_c Z(D(B), c) = \min_a Z(B, a), \qquad U(D(B)) = \frac{L_B}{L_{D(B)}} U_{\text{list}}(B).$$

Canonical lifting preserves each integer hole count, and arbitrary target residue vectors admit polynomial-time extraction with no greater source hole count. The target may be required to have only even moduli and at most two moduli of each positive 2*-adic valuation.*

Proof. Compose Lemmas 9.5 and 9.6. Their pointwise maps compose as well. In the second stage choose $h_i = i$ and $T = M$. This gives the last promises: at valuation j there is one mask and one data modulus. The intermediate list has polynomial length and polynomial-bit entries, so writing all shifted entries and masks costs polynomially many bits, including the $O(M^2)$ bits of the mask list. Alternatively, smaller heights may be reused for different odd parts. In either case the stated period identities give the density scaling. The moduli depend only on B, not on a source optimizer. □

Corollary 9.8 (Distinct-modulus covering hardness)**.** *Deciding $U(A) = 0$ for distinct binary moduli is NP-hard under deterministic polynomial-time many-one reductions, even with the final promises of Theorem 9.7. The general problem belongs to Σ_2^P.*

Proof. Exact Pinwheel Covering on lists is NP-hard [23], as recorded in [7, Theorem 3]. There is also a direct route from Dense Pinwheel Scheduling [24, Corollary 5.1], whose unary strengthening is [25, Theorem 1.1]. For either dense source, first test $\sum_i 1/d_i = 1$ exactly as above; otherwise output the fixed no-instance $\{2\}$, which satisfies the target promises. At density one, residue classes cover if and only if they are pairwise disjoint: equality in the union bound forces every pairwise intersection to be empty, since a nonempty congruence intersection has positive density. Dense scheduling admits exactly such residue schedules [27, Observation 1]. Hence these dense sources already imply exact list-covering hardness. A modulus 1 is a trivial yes-case for the exact-covering source; for the dense source, handle it only after the density test, when the list must be (1). Replace either yes-case by $(2, 2)$ before distinctification. An empty list maps to $\{2\}$. The positive scaling in Theorem 9.7 preserves exactly whether the optimum is zero. For the upper bound, guess the polynomial-bit residue vector and universally quantify the $O(\log L)$-bit positions in one period; positions outside $[0, L)$ are accepted automatically. Each coverage test is polynomial-time. This is a Σ_2^P predicate, not an NP certificate. □

For a supplied list of residue classes, existence of a hole is the NP-complete Simultaneous Incongruences problem [15, AN2, p. 249]; verifying coverage is its coNP-complete complement. This explains the difficulty with the naive residue certificate, not a proof that the existential distinct-modulus problem is outside NP.

The transfer proves binary-input hardness, not strong hardness. In the printed construction $T = M$, and even reusing heights requires $T \geq m - 1$ for an odd part repeated m times. The masks include 2^T, which has polynomial bit length but can have exponential numerical magnitude. The repeated cell periods in [25] do not avoid this blowup. Unary hardness of distinct covering is not established here.

Example 9.9 (An even source)**.** For $B = (2, 3, 6)$, take $\ell_1 = 5$. The odd list is $(5, 5, 5, 5, 3, 15)$. Using successive heights within each repeated odd part gives $D = \{2, 4, 8, 5, 10, 20, 40, 3, 15\}$, with periods 6 and 120 and minimum hole count one: the first two source classes cover 2/3, and the last adds at most 1/6. The source residues $(0, 0, 1)$ attain this bound and lift, in that order, to $(0, 1, 3, 2, 3, 19, 15, 0, 1)$, leaving only 71 (mod 120).

In these instances the masks alone leave density at most 2^{-T}, and the scaling L_B/L_D may be exponentially small; see Section 15.1 for the distinction between exact and approximate optimization.

10. Exact evaluation is #P-hard with a unique local state

We reduce counting independent sets in 3-regular graphs to exact density evaluation even though a CD optimizer is explicit. Edge primes encode the graph, and modular interpolation recovers the source count; the preliminary lemma supplies enough prime labels in suitable residue classes. Write $\pi(x)$ for the prime-counting function.

Lemma 10.1 (Populated prime classes at a polynomial cutoff)**.** *Let $N, K \geq 1$, $s = N+K+2$, $B = 1024s^2$, $Y = 2^{27}s^4$ and $X = 2^{40}s^8$. There are at least $N+1$ primes ℓ with $B < \ell \leq Y$. For each such ℓ, more than $4(N+1)$ nonzero classes modulo ℓ contain at least K primes from $(\ell, X]$. These fields and classes can be found in deterministic polynomial time in $N+K$.*

Proof. For an integer $z \geq 2$, a largest binomial coefficient in row z is at least $2^z/(z+1)$. Its exponent at a prime p is at most $\lfloor \log_p z \rfloor$: each term of the factorial-valuation difference is zero or one. Hence that coefficient is at most $z^{\pi(z)}$, giving the elementary bound

$$\pi(z) \log z \geq z \log 2 - \log(z+1).$$

Using $1/2 \leq \log 2 \leq 1$ and $\log s \leq s-1$ yields

$$\pi(Y) \geq 2^{20}s^3, \qquad \pi(X) \geq 2^{33}s^7.$$

For example, $\log X \leq 40s$ and $\log(X+1) \leq 41s$, so $X/2 - 41s \geq 40s \cdot 2^{33}s^7$ proves the second inequality; the first follows from $\log Y \leq 27s$ and $\log(Y+1) \leq 28s$. Since $\pi(B) \leq B$, the first bound leaves at least $N+1$ eligible fields.

Fix such an ℓ. A residue class contains at most $\lfloor X/\ell \rfloor + 1 \leq 2^{30}s^6 + 1$ integers in the pool. If at most $4(N+1)$ classes contain K primes, then counting the remaining classes with fewer than K primes gives

$$\pi(X) \leq (\ell+1) + 4s(2^{30}s^6+1) + \ell K \leq 2^{32}s^7 + 2^{27}s^4 + 2^{27}s^5 + 4s + 1 < 2^{33}s^7,$$

a contradiction. The zero class contains no prime greater than ℓ. Finally, X has polynomial numerical size in $N+K$. Enumerating integers, testing primality even by trial division, and sorting primes into residue classes therefore takes polynomial time. The large explicit constants are convenient bounds, not proposed practical cutoffs. □

Theorem 10.2 (#P-hardness with supplied prime factorizations)**.** *EXACT-E278 is #P-hard under polynomial-time Turing reductions, even for binary-encoded moduli polynomially bounded in their number r, with complete prime factorizations supplied and the following promises:*

(i) every occurring prime divides exactly two moduli and exceeds the number of labels;

(ii) each modulus is squarefree with exactly three distinct prime factors;

(iii) $H(A) = 0$ and the saturated state product is a singleton;

(iv) the noncoprime graph is 3-regular;

(v) the instance is in the strict Shearer region and an explicit CD optimizer is available.

Proof. Counting independent sets in *3-regular* graphs is #P-complete [16, Theorem 3.1]. Start with such a nonempty simple graph $G = (V, E)$ on N vertices, and write

$$I_G(z) = \sum_{k=0}^{N} i_k z^k, \qquad K = |E| = 3N/2.$$

The empty source graph is handled directly. Apply Lemma 10.1 with these N, K.

Sampling points and prime labels. For an eligible field $\mathbb{F}_\ell$, call a nonzero residue d populated if at least K pool primes have that residue. The map $d \mapsto -d^{-3}$ has fibres of size at

most three: an equation $-d^{-3} = z \neq 0$ is a nonzero cubic equation. More than $4(N+1)$ populated classes therefore provide at least $N+1$ distinct sample points. Choose them, and for each d select K distinct primes $R_j \in (\ell, X]$ congruent to d modulo ℓ. These primes exceed K because $\ell > B > K$.

Arithmetic graph encoding. Give each edge a distinct prime label R_e and put

$$n_v = \prod_{e \ni v} R_e. \tag{10.1}$$

Each modulus has three distinct factors. Two different vertices in a simple graph share at most their joining edge; their three-edge incidence sets cannot coincide. Unique prime factorization therefore makes all moduli distinct, without private labels. Adjacent moduli have gcd R_e, whereas nonadjacent moduli are coprime, so $G_A = G$. Orient each edge and prescribe residue 0 at its tail and 1 at its head modulo R_e. CRT constructs a CD vector.

Every label has incidence two and exceeds K, so zero overload holds and each local saturated state is unique. Moreover,

$$\sum_v \frac{1}{n_v} \leq \frac{N}{(K+1)^3} < 1.$$

On any induced subset, the CD identity and the union bound give $q_S(A) \geq 1 - \sum_{v \in S} 1/n_v > 0$. All promises follow. Adenwalla's identity, together with CD optimality, gives

$$U(A) = \sum_{S \text{ independent in } G} (-1)^{|S|} \prod_{v \in S} \frac{1}{n_v}.$$

Exact recovery and costs. Since $n_v \equiv d^3 \pmod{\ell}$ and $\ell \nmid L$,

$$U(A) \equiv I_G(-d^{-3}) \pmod{\ell}. \tag{10.2}$$

Query the exact oracle at the $N+1$ distinct samples and interpolate to recover each $i_k \bmod \ell$. Use $N+1$ distinct eligible fields. Their product exceeds 2^N, while $0 \leq i_k \leq 2^N$, so integer CRT recovers every coefficient and hence $I_G(1)$.

There are $(N+1)^2$ queries. All labels are at most X, all moduli at most X^3, and the prime search has polynomial numerical cutoff. Field arithmetic, rational reduction, interpolation and integer CRT involve polynomially many polynomial-bit integers. Thus the entire construction and recovery, including the supplied factorizations, form a deterministic polynomial-time Turing reduction. □

The moduli are polynomially bounded in their number: $r = N$, $K = 3N/2$ and $n_v \leq X^3 = 2^{120}(5N/2+2)^{24}$. Thus the same counting hardness holds with unary moduli.

Remark 10.3 (Oracle encoding)**.** Attainable densities are C/L, $0 \leq C \leq L$, so reduced answers have $O(b)$ bits. For $U = 1 - M = u/v$, reduction modulo ℓ is uv^{-1}, defined since $v \mid L$ and $\ell \nmid L$. The reduction is Turing, not parsimonious or many-one.

Remark 10.4 (Arithmetic content checked in Lean)**.** Adding one distinct private prime at every vertex gives a four-prime variant, with $K = 5N/2$ and samples $-d^{-4}$; the same argument proves its hardness. The companion checks the populated-class, query-promise, recovery and input-size arithmetic for that four-prime variant. Appendix B.1 distinguishes these components from the written complexity classification and three-prime specialization.

Remark 10.5 (What the exact hardness separates)**.** The hard instances have an explicit optimizer; Adenwalla's CD identity makes their value an independence-polynomial evaluation. For any supplied vector, $O(\varepsilon^{-2}\log(1/\rho))$ uniform samples modulo L estimate coverage within additive error ε with failure probability at most ρ. Each test has polynomial bit cost; binary rejection generates a uniform residue in fewer than two expected trials. This also estimates $M(A)$ when a CD optimizer is supplied. Approximate optimization for general inputs remains an open direction (Section 15.1).

Remark 10.6 (Coprime bases and prime promises)**.** Composite coprime labels suffice for the algorithms, but not for the squarefree three-prime promise. The pool construction supplies genuine primes.

11. A natural #P-complete integer output

Counting uncovered residues gives an integer output with a direct counting-machine interpretation. Multiplication by L transfers the independent-set encoding of Section 10 while retaining its arithmetic promises. For any supplied residue vector define

$$Z(A,a) = |\{x \bmod L : x \not\equiv a_i \pmod{n_i} \text{ for all } i\}|.$$

Theorem 11.1 (Zero-overload hole count)**.** *The total function $Z(A,a)$ is #P-complete under polynomial-time Turing reductions. Hardness persists for moduli polynomially bounded in their number, with squarefree three-prime factorizations, zero overload and a CD optimizer supplied, as in Theorem 10.2. On these inputs, $Z(A,a) = L\,U(A)$.*

Proof. Compute L by gcd and multiplication. Guess exactly $\lceil \log_2 L \rceil$ bits, interpret them as x, reject $x \geq L$, and otherwise accept precisely when every modular avoidance test succeeds. Each uncovered residue has exactly one accepting path, and all tests are polynomial in the binary input length. This proves membership for the total function, without requiring recognition of the promise.

For hardness use Theorem 10.2, supplying its explicit CD vector. Since every distinct prime label occurs, $L = \prod_{j=1}^{K} R_j$, and $\ell \nmid L$. A hole-count answer gives

$$Z(A,a)L^{-1} \equiv I_G(-d^{-3}) \pmod{\ell}.$$

The same interpolation and integer CRT recover the source independent-set count. Every query satisfies the promise. CD exactness proves $Z = L\,U(A)$ there. □

Remark 11.2 (Promise convention)**.** The total count is in #P; hardness queries only promised inputs. No membership claim is made for minimizing over vectors. The integer answer has $O(\log L)$ bits.

Remark 11.3 (Unrestricted counting versus the arithmetic promises)**.** For unrestricted lists, CRT already encodes independent sets: assign each vertex a distinct polynomially bounded prime $p_v > 2$; classes modulo p_v forbid residues $2, \ldots, p_v - 1$, and one class modulo $p_u p_v$ forbids $(1,1)$ for each edge. Every prime occurs, even at isolated vertices, so holes correspond bijectively to independent sets. This parsimonious encoding repeats moduli, has high incidence and supplies no CD optimizer; Theorem 11.1 preserves the stronger arithmetic promises.

12. Noncoprime-graph treewidth and exact algorithms

Return the empty instance immediately. Otherwise put $d_p = |I_p|$ and $\Delta = \max_p d_p$. For a gcd-free basis use $d_j = |J_j|$; Lemma 4.4 gives the same maximum Δ. Every prime or basis scope is a clique of G_A, so

$$\Delta \leq \omega(G_A) \leq \mathrm{tw}(G_A) + 1. \tag{12.1}$$

The two tasks are to choose a bounded set of local decisions and to evaluate each resulting residue vector.

Lemma 12.1 (State bounds)**.** *For a prime or coprime-base coordinate of incidence d, the number of saturated local trees is at most $(d+1)^{\binom{d}{2}}$. If its base is at least d, that number is one. Consequently the entire state product has size at most*

$$s(\Delta) = (\Delta+1)^{\binom{\Delta}{2}}, \qquad g(\Delta) = s(\Delta)^{\pi(\Delta-1)}. \tag{12.2}$$

For a squarefree prime coordinate the exact count is the Stirling number $S(d, \min(p,d))$; in particular it is at most p^d.

Proof. Let the fixed exponents be e_i. For each pair, record its first separating digit ℓ_{ij}, or ∞ if its prefixes agree through $\min(e_i, e_j)$ digits. While a pair continues together through a saturated split, its child has fewer live indices than its parent. Starting with d indices, it must therefore separate by digit $d-1$ or one index expires first. The codes lie in a palette of at most $d+1$ symbols, independently of the exponents.

Here is the inverse encoding. At depth h, restrict to indices with $e_i \geq h$. They share a live node exactly when $\ell_{ij} \geq h$, and share a child exactly when $\ell_{ij} > h$, with the usual convention for ∞. These relations recover every indexed partition; the exponents specify which indices stop after that digit. Singleton tails have fixed lengths and no unlabelled choices. Thus the encoding is injective, giving the stated pair-code bound. For example, the state $12|3$ for depths $(1, 2, 1)$ has $\ell_{12} = \infty$ and $\ell_{13} = \ell_{23} = 1$: index 1 expires before separating from index 2.

If the base is at least d, the first split separates all indices. Only bases smaller than Δ can carry choices. Pairwise-coprime bases have distinct least prime divisors, so there are at most $\pi(\Delta - 1)$ such bases, giving the product bound without factoring them. For squarefree prime coordinates there is only one digit: the states are exactly the partitions into $\min(p, d)$ nonempty blocks. □

Lemma 12.2 (Signed evaluation on the noncoprime graph)**.** *Given binary moduli, a residue vector a, and a width-τ tree decomposition of G_A, the hole count $Z(A, a)$ is computable in $2^{\tau+1}\,\mathrm{poly}(b)$ bit operations, without prime factorization. Intermediate integers have $O(r + \log L) = O(b)$ bits.*

Proof. Factorization into small scopes. Compute a gcd-free basis. For $x \in \{0,1\}^r$, let $S(x) = \{i : x_i = 1\}$ and set

$$c_{ij}(x_i, x_j) = \mathbf{1}[x_i x_j = 0 \text{ or } a_i \equiv a_j \pmod{\gcd(n_i, n_j)}].$$

The compatibility expansion in Proposition 2.3, multiplied by L, is

$$Z(A, a) = \sum_{x \in \{0,1\}^r} \prod_i (-1)^{x_i} \prod_{ij \in E(G_A)} c_{ij}(x_i, x_j) \prod_{j=1}^{s} \beta_j^{E_j - \max_{i \in J_j}(e_{ij} x_i)}. \tag{12.3}$$

The last exponent uses maximum zero when no selected index belongs to J_j. Generalized CRT proves the compatibility indicators, and the last product is $L/\operatorname{lcm}(n_i : i \in S(x))$.

Every factor scope is a clique of G_A: a vertex, an edge, or J_j. A clique lies in a bag of every tree decomposition. An elimination order obtained from the decomposition therefore involves at most $\tau+1$ binary variables in any bucket, giving tables with at most $2^{\tau+1}$ entries.

Elimination and its invariant. To eliminate x_i, gather all current factors containing it, multiply their tables and sum over $x_i \in \{0, 1\}$. Replace these factors by the resulting table, called a message, on the remaining variables. This is distributivity: factors not depending on x_i stay outside that sum.

For each message, keep track of its set D_m of eliminated variables and its set $\mathcal{F}_m$ of original factors. For boundary values y, the invariant is

$$m(y) = \sum_{x \in \{0,1\}^{D_m}} \prod_{f \in \mathcal{F}_m} f(x, y).$$

Initially an original factor has $D_m = \varnothing$ and owns only itself. Distinct current messages own disjoint original factors and disjoint eliminated variables: a factor is consumed once, and all occurrences of a variable enter its elimination bucket together. Multiplication unions these sets; summing out x_i adds it to D_m. This proves the invariant inductively, including for temporary products, and the final scalar is (12.3).

Integer sizes. Each expanded summand has absolute value at most $\prod_j \beta_j^{E_j} = L$: it uses each base-weight factor at most once, while all other factors have absolute value at most one. Hence

$$|m(y)| \leq 2^{|D_m|} L \leq 2^r L.$$

The tables can be signed; the absolute-value bound applies before any cancellation. Their $O(r+\log L)$-bit entries and polynomially many $2^{\tau+1}$-entry tables prove the time and storage bounds. Computing the basis and initial factor entries takes polynomial additional work per entry. □

An elimination in the running example. For $A=\{2,4,6\}$ and $a=(1,0,0)$, the edge factors are $c_{12}=1-x_1x_2$, $c_{13}=1-x_1x_3$ and $c_{23}=1$. Eliminate x_1 together with the base-2 factor and its sign:

$$m(x_2,x_3)=\sum_{x_1=0}^{1}(-1)^{x_1}(1-x_1x_2)(1-x_1x_3)\,2^{2-\max(x_1,2x_2,x_3)}.$$

The four entries are $m(0,0)=m(0,1)=2$ and $m(1,0)=m(1,1)=1$. The remaining base-3 factor and signs give

$$Z(A,a)=\sum_{x_2,x_3=0}^{1}(-1)^{x_2+x_3}3^{1-x_3}m(x_2,x_3)=6-2-3+1=2.$$

Thus this residue vector leaves $2/12=1/6$ uncovered, matching Example 7.2. The same table operations apply when the numerical period is too large to list.

Theorem 12.3 (Factorization-free treewidth algorithm)**.** *On ordinary binary inputs, given a width-τ decomposition of G_A, exact evaluation of $M(A)$ and recovery of an optimizing residue vector take the following number of bit operations:*

$$g(\Delta)2^{\tau+1}\,\mathrm{poly}(b),\qquad \Delta\le\tau+1, \tag{12.4}$$

Processing state vectors sequentially uses $2^{\tau+1}\,\mathrm{poly}(b)$ working bits. Without a supplied decomposition, the additional cost is $f(\tau)\,\mathrm{poly}(b)$ for a computable function f. Thus the problem is fixed-parameter tractable in $\mathrm{tw}(G_A)$ alone.

Proof. Compute a gcd-free basis by Lemma 4.4 and enumerate its at most $g(\Delta)$ saturated state vectors (Lemma 12.1). Every partial partition extends to a saturated tree, so depth-first generation has polynomial delay in the stored tree size. Generate only occupied children and store singleton tails by their exponents; no iteration over a large base's digit alphabet is needed.

Label occupied children canonically and apply CRT to obtain a representative a. Evaluate its hole count by Lemma 12.2, retaining the best representative and discarding each state vector after evaluation. Proposition 4.5(i)–(ii) guarantees that an enumerated state and its canonical representative attain the optimum. Stored trees and residues have polynomial size in b, and exact CRT has polynomial bit cost. The signed-table bound completes the time and space estimates.

Finally, the fixed-parameter algorithm for treewidth [6] constructs a suitable decomposition in $f(\tau)\,\mathrm{poly}(b)$ time. The length of a supplied decomposition is included in b; alternatively one uses a reduced decomposition with a polynomial number of bags. □

Lean counterparts of the supplied-basis steps in Proposition 4.5(i)–(ii) are mapped in Appendix B.1.

Corollary 12.4 (CD recognition)**.** *Existence of a CD residue vector, and recovery of one when it exists, take $g(\Delta)\,\mathrm{poly}(b)$ time without prime factorization. No Shearer-positivity promise is required.*

Proof. A saturation move sends one live index from a crowded child to an unused child of the same prefix. It cannot make a previously incompatible pair compatible. A pair that separated above this prefix stays separated; an index ending at or above it sees unchanged shared digits; and every other live index at the prefix avoids the newly chosen child. This argument holds for each coprime-base coordinate. Thus saturation preserves every CD edge inequality, and canonical relabelling preserves all prefix-equality relations.

Consequently a CD residue vector exists if and only if one of the at most $g(\Delta)$ canonical saturated vectors is CD. Generate each vector as above and check $(a_i - a_j) \not\equiv 0 \pmod{\gcd(n_i, n_j)}$ on every edge. Generation, CRT and testing have polynomial bit cost. General recognition is NP-complete by Proposition 9.3, whose noncoprime graph is complete; bounded incidence makes it tractable. □

Corollary 12.5 (CD value evaluation)**.** *If a CD residue vector exists, the scalar value $M(A)$ is computable from a width-τ decomposition of G_A in $2^{\tau+1}\,\mathrm{poly}(b)$ time. No CD witness or factorization is needed for this value computation.*

Proof. By Theorem 3.3, $U(A) = q_{[r]}(A)$. This is an independence-polynomial evaluation: each selected vertex has weight $-1/n_i$ and each edge forbids selecting both endpoints. Binary-variable elimination has the stated table bound. To use integers throughout, multiply by $D = \prod_i n_i$, giving unary weights n_i for an unselected vertex and -1 for a selected one. Intermediate absolute values are at most $2^r D$, so their bit lengths are polynomial. □

Corollary 12.6 (Local parameters alone do not suffice)**.** *Unless* FP = #P*, no exact evaluation algorithm, even on supplied factorizations, has running time $f(\Delta)\,\mathrm{poly}(b)$. The same obstruction applies to a parameter tuple consisting only of Δ, $H(A)$, the maximum local state count and the maximum fixed-state support size.*

Proof. The hardness family in Theorem 10.2 has $\Delta = 2$, $H = 0$, local state count one, and fixed-state support size at most three. Such an algorithm would evaluate all oracle instances in polynomial time and collapse the stated counting classes. □

Corollary 12.7 (Two factors versus three at incidence two)**.** *If every modulus has at most two distinct prime factors and every prime divides at most two moduli, exact optimization and optimizer recovery take polynomial time without supplied factorizations. Allowing three squarefree factors instead already gives the #P-hardness of Theorem 10.2.*

Proof. Each vertex of G_A has degree at most two. Its components are paths, cycles or isolated vertices, so its treewidth is at most two. Apply Theorem 12.3; arbitrary prime powers cause no change to this argument. □

Remark 12.8 (Comparison with prime-coordinate elimination)**.** Let H_P join primes occurring in a common modulus. With supplied factorizations, enumerating local states and then eliminating the at most $\Delta + 1$ positive atoms of each fixed prime law gives $g(\Delta)(\Delta+1)^{w+1}\,\mathrm{poly}(b)$ time for a width-w decomposition of H_P. This is an alternative implementation, not the parameter in Theorem 12.3. In fact,

$$\mathrm{tw}(G_A) \le \Delta\big(\mathrm{tw}(H_P) + 1\big) - 1.$$

Replace each bag of an H_P decomposition by the union of the incidence scopes of its primes. This covers every edge of G_A. For each modulus, the subtrees of bags containing its prime factors meet in a common bag, since those factors form a clique; their union is connected. Bag sizes give the bound. Conversely, pairwise-coprime moduli with many prime factors give an edgeless G_A but arbitrarily large cliques in H_P.

13. Arithmetic realization and limits of local state deletion

The next results limit deletion rules based only on local incidence data. Here an *arithmetic completion* of a squarefree prime scope means choosing pairwise-coprime cofactors m_i, coprime to p, to form moduli pm_i. Unlike Section 6's outside completion, this varies A while preserving the common-prime incidence data. Already $p = 2$ gives $S(d, 2) = 2^{d-1} - 1$ saturated bipartitions, each realizable as a completion's unique optimizer. In general a squarefree prime scope of size $d \ge p$ has states $\pi \in \Pi_{d,p}$, the partitions into p nonempty blocks. A uniform block is its mask, so

$$\varphi_\pi(\{i,j\}) = \begin{cases} 1 - 1/p, & i, j \text{ lie in one block,} \\ 1 - 2/p, & i, j \text{ lie in different blocks.} \end{cases}$$

Theorem 13.1 (Exposed partition profiles)**.** *Every φ_π, $\pi \in \Pi_{d,p}$, is a distinct retained vertex. The nonnegative functional*

$$L_\pi(\varphi) = \sum_{\substack{i<j \\ i,j \text{ in different } \pi\text{-blocks}}} \varphi(\{i,j\})$$

is uniquely minimized by φ_π among the partition profiles and their convex hull*. A sufficiently small strictly positive perturbation exposes φ_π uniquely also in the upper-closed polyhedron of Section 6.*

Proof. The target attains the smallest value $1-2/p$ in every summand. If another partition attains the same total, it must separate every target cross-block pair. Hence each of its blocks is contained in a target block. Both partitions have p nonempty blocks, so they coincide. Pair coordinates also recover the same-block relation, giving distinctness.

For any different profile the functional gap is at least $1/p$. All 2^d avoidance coordinates lie in $[0,1]$. Therefore

$$L_\pi(\varphi) + \varepsilon \sum_{T\subseteq[d]} \varphi(T), \qquad \varepsilon = \frac{1}{p2^{d+1}},$$

retains a positive gap of at least $1/(2p)$ between the target and every alternative. Its coefficients are strictly positive, so any nonzero upward slack also increases the value. It uniquely exposes the target in the entire upper-closed hull. This distinction matters because the unperturbed L_π has zero coefficients and need not expose a unique point on an unbounded upward ray. □

For $d = 2p$ the number of profiles is at least the number of perfect-pair partitions,

$$|\Pi_{2p,p}| \geq \frac{(2p)!}{2^p p!} = (2p-1)!! = d^{\Omega(d)}.$$

The last $\lfloor p/2 \rfloor$ factors are at least p. This growth requires varying p; for fixed p, $|\Pi_{d,p}| \leq p^d$.

Theorem 13.2 (Arithmetic necessity of every local partition)**.** *Fix a prime p, $d > p$, and $\pi \in \Pi_{d,p}$. There exist distinct primes $m_1, \ldots, m_d$, all different from p, such that for*

$$A_\pi = \{pm_1, \ldots, pm_d\}$$

the unique optimal partition of the classes among the p residue fibres is π, up to a permutation of fibre labels. Competing residue vectors with empty fibres are included. The balancing weights can be chosen explicitly, and a positive rational separation bound is available.

Proof. Explicit weights. Choose an anchor a_B in each target block B. Enumerate its nonanchor indices, across all blocks, as $i_1, \ldots, i_h$, where $h = d - p$. Set

$$w_{i_j} = 3^{-j}, \qquad w_{a_B} = 1 - \sum_{i_j \in B} 3^{-j}. \tag{13.1}$$

These are positive rationals, every target block has weight one, and all weights have denominator dividing $Q = 3^h$. They are *not* required to be distinct: singleton blocks have anchor weight one.

In fact, a subset has total weight one if and only if it is exactly a target block. For a subset T, expand its weight as

$$\sum_{i\in T} w_i = k + \sum_{j=1}^{h} \epsilon_j 3^{-j}, \qquad \epsilon_j \in \{-1, 0, 1\},$$

where k is the number of anchors in T and $\epsilon_j = \mathbf{1}[i_j \in T] - \mathbf{1}[a_{B(i_j)} \in T]$. The fractional correction has absolute value less than $1/2$. Weight one forces $k = 1$ and a zero correction. In a nonzero signed base-three sum, the first nonzero term has magnitude greater than the sum of all later magnitudes, so cancellation to zero is impossible. Thus every $\epsilon_j = 0$: a

nonanchor is included exactly when its anchor is included. There is one anchor, so T is one whole target block.

Extend any competing partition ρ to exactly p fibres by adding empty fibres and put $W_B = \sum_{i\in B} w_i$. The total is p, whence

$$\sum_{B\in\rho} W_B^2 - p = \sum_{B\in\rho} (W_B - 1)^2. \tag{13.2}$$

It vanishes only for $\rho = \pi$ up to fibre labels. If it is positive, each deviation is an integer multiple of $1/Q$ and the deviations sum to zero; at least one is positive and one negative. Therefore the gap is at least $2/Q^2$, including partitions with empty fibres.

Distinct prime cofactors in certified windows. Put

$$\eta = \frac{1}{8p^2Q^2}, \qquad t \geq 4p^3Q^2.$$

For sufficiently large t, every interval

$$\left[\frac{t}{w_i}, \frac{(1+\eta)t}{w_i}\right] \tag{13.3}$$

contains at least d primes. This follows from the prime number theorem, since p, d, w_i, η are fixed and each interval has fixed positive relative length. Select one prime m_i from each interval greedily while avoiding previously selected primes. All $m_i > p$ for large t. Repeated weights pose no problem: at most $d-1$ primes have been used when the next one is chosen.

Exact comparison, with a uniform remainder bound. Put $x_i = 1/m_i$ and $z_i = tx_i$. Then

$$\frac{w_i}{1+\eta} \leq z_i \leq w_i, \qquad |z_iz_j - w_iw_j| \leq 2\eta w_iw_j.$$

By CRT the uncovered density for a fibre partition ρ is

$$\frac{F(\rho)}{p}, \qquad F(\rho) = \sum_{B\in\rho} \prod_{i\in B} (1 - x_i),$$

with empty-fibre product one. Expand the finite products as

$$F(\rho) = p - \sum_i x_i + E_2(\rho; x) + R_\rho, \qquad E_2(\rho; x) = \sum_{B\in\rho} \sum_{\{i,j\}\subseteq B} x_ix_j.$$

For the ideal weights,

$$E_2(\rho; w) - E_2(\pi; w) = \frac{1}{2}\left(\sum_{B\in\rho} W_B^2 - p\right) \geq \frac{1}{Q^2}$$

whenever $\rho \neq \pi$. Replacing w by z changes this difference by at most

$$2\eta \sum_{i<j} w_iw_j \leq \eta p^2 = \frac{1}{8Q^2}.$$

Hence $E_2(\rho; x) - E_2(\pi; x) \geq 7/(8Q^2t^2)$.

Let $X = \sum_i x_i \leq p/t$. The sum of all terms of degree at least three in absolute value, over any fibre partition, is at most

$$\sum_{k\geq 3} \frac{X^k}{k!} \leq \frac{X^3e^X}{6}.$$

Thus $|R_\rho - R_\pi| \leq X^3e^X/3 < p^3/t^3 \leq 1/(4Q^2t^2)$, using $t \geq p$ and $e < 3$. The constant and linear terms are exactly independent of the partition, so

$$F(\rho) - F(\pi) > \frac{5}{8Q^2t^2} > 0. \tag{13.4}$$

In particular the uncovered-density gap exceeds $1/(2pQ^2t^2)$. This proves unique optimality without relying on an unquantified Taylor remainder. Every residue vector has one such

fibre partition, and the private-prime residues do not affect the CRT products. Conversely CRT realizes every partition. The theorem follows. □

The Lean proof uses the audited PrimeNumberTheoremAnd dependency [35]; finite witnesses are checked by primality tests and rational inequalities.

Corollary 13.3 (Polynomial witness bit length)**.** *The cofactors in Theorem 13.2 may each be chosen with $O(d^2)$ bits, uniformly over $p < d$ and the target partition. Their total encoding length is $O(d^3)$. This is an existence bound, not a polynomial-time prime-search algorithm.*

Proof. The explicit Trudgian estimate (14.5) implies

$$|\vartheta(x) - x| \le C_0 x (\log x)^{1/4} e^{-c\sqrt{\log x}}$$

for absolute positive constants and all sufficiently large x. Here $\log Q = O(d)$, $\log(1/\eta) = O(d)$ and $Q^{-1} \le w_i \le 1$. Choose $t = \lceil \exp(Cd^2) \rceil$ with a sufficiently large absolute C. For $x = t/w_i$, uniformly $t \le x \le Qt$. The relative errors at x and $(1+\eta)x$ are then at most $\eta/8$, because their exponential decay is $e^{-c\sqrt{C}\,d}$ up to a polynomial factor in d. Increasing C handles the finitely many small dimensions and also ensures $t \ge 4p^3Q^2$.

It follows that

$$\vartheta((1+\eta)x) - \vartheta(x) \ge \eta x/2.$$

Each prime contributes at most $\log((1+\eta)Qt)$, so every window contains at least $\eta t/(2\log((1+\eta)Qt)) \ge d$ primes, after the same absolute enlargement of C. Greedy selection gives distinct cofactors at most $(1+\eta)Qt$, whose logarithms are $O(d^2)$. All the hypotheses of the arithmetic-necessity proof hold. □

Corollary 13.4 (Completion-independent state deletion)**.** *Fix a prime p, an integer $d > p$, and the common prime's squarefree incidence pattern. Suppose a preprocessing rule retains a set $R \subseteq \Pi_{d,p}$ depending only on this local data, and guarantees an optimal retained state for every completion by distinct private-prime cofactors. Then $R = \Pi_{d,p}$.*

Proof. If a partition π were omitted, choose its completion from Theorem 13.2. Every optimizer has partition π up to fibre labels, so no retained partition attains the optimum. This contradicts the rule's guarantee. □

This applies to Section 6's frontier, not completion-aware optimization, implicit encodings or arbitrary algorithms. Different partitions require different numerical completions.

Allowing composite cofactors makes realization deterministically constructive. The following lemma also supplies the star reduction's cofactors.

Lemma 13.5 (Deterministic coprime windows)**.** *Let $d \ge 1$. Given an integer $c \ge 2$, rationals $0 < w_i \le 1$ for $1 \le i \le d$, a rational $0 < \eta \le 1$, and an integer $s \ge 2d/\eta$, one can construct pairwise-coprime integers $m_i > c$, all coprime to c, and $t \ge cs$ such that*

$$t/w_i \le m_i \le (1+\eta)t/w_i.$$

The construction has polynomial running time and output length in the binary inputs.

Proof. Choose a_i as the first unused integer at or above $\lceil s/w_i \rceil$. Then $s/w_i \le a_i < s/w_i + d$. Set

$$R = c\prod_{i<j} |a_i - a_j|, \qquad m_i = Ra_i + 1, \qquad t = sR.$$

This is the standard difference-product device for forcing coprimality: every m_i is coprime to R, so a common divisor of m_i, m_j divides $a_i - a_j$, which divides R, and must be one. Also $R \ge c$, so $t \ge cs$ and $m_i > c$. The lower window bound is immediate, and

$$m_i < t/w_i + dR + 1 \le t/w_i + 2dR \le (1+\eta)t/w_i,$$

using $\eta s \ge 2d$ and $w_i \le 1$. With $w_{\min} = \min_i w_i$, each $a_i < s/w_{\min} + d$. There are at most d trials per choice and $\binom{d}{2}$ differences to multiply, all with polynomially many bits. Integer arithmetic gives the claimed cost, without factoring or prime search. □

Corollary 13.6 (Deterministic coprime realization)**.** *Given a prime $p < d$ and $\pi \in \Pi_{d,p}$, one can construct in deterministic polynomial time pairwise-coprime integers $m_i > p$, all coprime to p, such that π is the unique optimal fibre partition for $A = \{pm_1, \dots, pm_d\}$, up to fibre labels. Each cofactor has $O(d^3)$ bits and their total length is $O(d^4)$. An explicit uncovered-density gap separates every other partition, including empty-fibre competitors.*

Proof. Use the weights (13.1), put $Q_0 = 3^{d-p}$, and apply Lemma 13.5 with

$$c = p, \qquad \eta = \frac{1}{8p^2Q_0^2}, \qquad s = 16dp^2Q_0^2.$$

Then $\eta s = 2d$ and $t \geq ps \geq 4p^3Q_0^2$. CRT gives the same product formula $F(\rho)/p$ as in Theorem 13.2, because the cofactor coordinates are pairwise coprime; their primality is not used in its exact comparison. The window inequalities therefore give (13.4) with $Q = Q_0$, and an uncovered-density gap greater than $1/(2pQ_0^2t^2)$.

The weights lie in $[Q_0^{-1}, 1]$, so $\log a_i = O(d)$. Thus $\log R = O(d^3)$ and each m_i has $O(d^3)$ bits. The weights and all outputs are computable in polynomial time. This gives a constructive coprime completion, not prime or squarefree cofactors; the prime witness bound in Corollary 13.3 remains stronger in those respects. □

14. Threshold complexity on the constant-gcd star

Here the difficulty lies in splitting classes between residue fibres. For the centre modulus 3, we encode PARTITION into this choice through reciprocals of pairwise-coprime cofactors, then strengthen the construction to prime cofactors. Let p be prime and let $m_1, \dots, m_n > 1$ be pairwise coprime and coprime to p. Put

$$A = \{p\} \cup \{pm_i : 1 \leq i \leq n\}.$$

The cofactors need not be prime. The basis $\{p, m_1, \dots, m_n\}$ has a star co-occurrence graph, while G_A is complete on $n+1$ vertices. Prime cofactors also make the prime co-occurrence graph H_P a star; general coprime cofactors need not. Normalize the class modulo p to fibre 0. A class modulo pm_i placed there adds no coverage. Moving it to another fibre cannot decrease coverage, so an optimizer places all these classes in the other $p-1$ fibres. They define a partition $[n] = B_1 \sqcup \cdots \sqcup B_{p-1}$, with empty parts allowed.

Proposition 14.1 (Exact star formula)**.**

$$U(A) = \frac{1}{p} \min_{[n]=B_1\sqcup\cdots\sqcup B_{p-1}} \sum_{b=1}^{p-1} \prod_{i\in B_b} \left(1 - \frac{1}{m_i}\right).$$

Proof. Fibre 0 is covered. In another fibre, CRT makes the pairwise-coprime cofactor coordinates independent and uniform. One residue out of m_i is covered by its class, giving the displayed uncovered product. Each fibre has density $1/p$, and every partition can be realized. Empty parts contribute one. □

For $p = 3$ this is balanced partition on the weights $-\log(1 - 1/m_i)$. Cambie gives the prime-cofactor formulas and a comparison-model obstruction [7, Proposition 2 and Section 4].

14.1. **An elementary deterministic reduction.** Recall the Star-Threshold language and rational encoding from Section 1.3. Its cofactor conditions are checked by gcd computations.

Lemma 14.2 (Coprime cofactors in narrow windows)**.** *Let $n \geq 1$ and let positive integers $W_1, \dots, W_n$ sum to $2T$, with $W_i \leq T$. Set $w_i = W_i/T$ and $\eta = 1/(64T^2)$. In deterministic polynomial time one can construct $t \geq 64T^2$ and pairwise-coprime integers $m_i > 3$, all coprime to 3, satisfying*

$$\frac{t}{w_i} \leq m_i \leq \frac{(1+\eta)t}{w_i}. \tag{14.1}$$

Each has $O(n^2(1+\log n+\log T))$ bits.

Proof. Apply Lemma 13.5 with $c=3$ and $s=128nT^2$, so $\eta s=2n$ and $t\ge 3s\ge 64T^2$. For the encoding bound, $a_i<sT+n$ gives $O(1+\log n+\log T)$ bits per auxiliary integer. The product R, and hence t, m_i, have $O(n^2(1+\log n+\log T))$ bits. □

Lemma 14.3 (A uniform rational separation)**.** *With the weights above, let $t\ge 64T^2$ and let pairwise-coprime $m_i>1$, coprime to 3, satisfy* (14.1)*. Put $x_i=1/m_i$, $X=\sum_i x_i$, and*

$$\theta=\frac{4t^2-2t^2X-\sum_i w_i^2+2+T^{-2}}{6t^2}. \tag{14.2}$$

There is a subset of the W_i summing to T if and only if $U(\{3\}\cup\{3m_i\})<\theta$. More precisely, for $A=\{3\}\cup\{3m_i\}$,

$$\begin{aligned} &\textit{YES:}\quad U(A)\le\theta-\frac{1}{8t^2T^2},\\ &\textit{NO:}\quad U(A)\ge\theta+\frac{1}{8t^2T^2}. \end{aligned} \tag{14.3}$$

Proof. Write $z_i=tx_i$ and $X\le 2/t$. For a bipartition B, B^c, put

$$F_B=\prod_{i\in B}(1-x_i)+\prod_{i\notin B}(1-x_i),\qquad S_B=\sum_{i\in B}w_i.$$

Proposition 14.1 gives $3U(A)=\min_B F_B$, including empty parts. Expanding the products gives

$$F_B=2-X+E_2(B;x)+R_B,\qquad |R_B|\le X^3e^X/6<4/t^3,$$

where E_2 sums same-part pair products. The windows imply

$$|E_2(B;z)-E_2(B;w)|\le 2\eta\sum_{i<j}w_iw_j\le 4\eta,\quad 2E_2(B;w)=S_B^2+(2-S_B)^2-\sum_i w_i^2.$$

Uniformly over every bipartition,

$$\left|2t^2(F_B-2+X)+\sum_i w_i^2-(S_B^2+(2-S_B)^2)\right|\le 8\eta+8/t\le\frac{1}{4T^2}.$$

Taking minima preserves this error bound. Thus

$$D=6t^2U(A)-4t^2+2t^2X+\sum_i w_i^2 \tag{14.4}$$

is at most $2+1/(4T^2)$ if a balanced partition exists. Otherwise every S_B-1 is a nonzero integer multiple of $1/T$, so $D\ge 2+7/(4T^2)$. The separating test $D<2+1/T^2$ is exactly $U(A)<\theta$. Its distance from either bound is $3/(4T^2)$; division by $6t^2$ gives (14.3). □

Theorem 14.4 (Deterministic star threshold NP-completeness)**.** Star-Threshold *is NP-complete under deterministic polynomial-time many-one reductions. Consequently EXACT-E278 is NP-hard on this family by a single exact-value query. The gcd-free basis can be supplied as part of the reduction.*

Proof. For membership in NP, a certificate is a bipartition B. Compute $F_B/3$ exactly and compare it with θ. The product denominator divides $\prod_i m_i$, so the verifier uses polynomial-bit integer arithmetic. Proposition 14.1 proves both completeness and soundness of this certificate. The input conditions are also decidable in polynomial time.

Reduce from binary PARTITION [22]. The empty instance is yes; an odd total or an item exceeding half the total is no. Map these cases respectively to the fixed valid instances $(A,\theta)=(\{3,15\},2/3)$ and $(\{3,15\},1/3)$, since $U(\{3,15\})=3/5$. Otherwise Lemma 14.2 constructs the cofactors, and Lemma 14.3 supplies the separating rational (14.2). Its numerator and denominator have polynomial bit length: all products contain only the polynomial-bit constructed integers and the original weights. This is a deterministic many-one map, not a computation using an unknown optimum. The basis $\{3,m_1,\dots,m_n\}$ is already pairwise

coprime; the moduli are distinct. Computing the exact value and comparing it with θ gives the stated function-hardness consequence. □

14.2. **The prime-cofactor strengthening.**

Lemma 14.5 (Efficient sampling in the required prime windows)**.** *Let $n \geq 1$ and let positive integers $W_1, \ldots, W_n$ sum to $2T$, with $W_i \leq T$, and set*

$$w_i = W_i/T, \quad \eta = \frac{1}{64T^2}, \quad k = \lceil \log_2(2T) \rceil, \quad t = 2^{4096k^2}.$$

Distinct primes $m_i > 3$ satisfying the windows (14.1) *can be found by a zero-error randomized algorithm in expected time polynomial in the binary input length. Each m_i has $O(k^2)$ bits.*

Proof. Trudgian's explicit estimate [41, Theorem 1] gives, for $x \geq 149$,

$$|\vartheta(x) - x| \leq x\varepsilon(x), \qquad \varepsilon(x) = \sqrt{\frac{8}{17\pi}}\, s^{1/2} e^{-s}, \quad s = \sqrt{\frac{\log x}{6.455}}. \tag{14.5}$$

For $x \geq t$, $s > 16k$, hence $\varepsilon(x) \leq e^{-s/2} < e^{-8k} \leq 2^{-(2k+7)} \leq \eta/8$. Here $s^{1/2} \leq e^{s/2}$, $1/2 < \log 2 < 3/4$, and $T \leq 2^{k-1}$ justify the bounds. Therefore, for $x = t/w_i$ and $y = (1+\eta)x$,

$$\vartheta(y) - \vartheta(x) \geq \eta x/2, \qquad \#\{p \text{ prime} : x < p \leq y\} \geq \frac{\eta x}{2 \log y}.$$

Since $t \leq x \leq Tt$ and $\log y \leq 4097k^2$, this count is at least $4T \geq 2n$: indeed $2^{4096k^2 - 3k - 5} \geq 8194k^2$ for $k \geq 1$. The interval contains at most $\eta x + 1 \leq 2\eta x$ integers. Even after excluding the at most $n - 1$ previously selected primes, a uniform integer sample succeeds with probability at least $1/(8 \log((1+\eta)Tt)) = \Omega(k^{-2})$.

Sample using binary rejection, test primality deterministically [2], and reject previously used primes. Every accepted value is valid; expected sampling and primality-testing costs are polynomial in n and k. The rational endpoints and all candidates have $O(k^2)$ bits. No enumeration up to the numerical height t is used. □

The constant 4096 is deliberately unoptimized; it makes all window estimates uniform, including the smallest inputs.

Theorem 14.6 (Prime-star threshold hardness)**.** *Restrict* Star-Threshold *to distinct primes $m_i > 3$. This problem is in NP and is NP-hard under ZPP many-one reductions. The reduction supplies complete factorizations of the distinct squarefree moduli. EXACT-E278 on this family is therefore NP-hard under ZPP Turing reductions making one exact-value query.*

Proof. Use the same fixed yes/no instances as in Theorem 14.4, and otherwise substitute Lemma 14.5 for the coprime construction. Lemma 14.3 applies unchanged and gives the threshold (14.2). Each cofactor has $O(k^2)$ bits, so the moduli and the rational threshold have polynomial encoding length. Every completed sample gives a correct instance, and expected runtime is polynomial. Primality is decidable in polynomial time [2]; hence the bipartition verifier also establishes membership in NP for this restricted language. The exact period is $3\prod_i m_i$, so one exact-value query has a polynomial-bit answer and suffices for the threshold test. □

Remark 14.7 (Parameter and reduction boundaries)**.** The supplied basis has star co-occurrence and tree incidence graphs. Actual prime–modulus incidence is also a tree: only 3 is shared. Theorem 14.4 excludes deterministic FPT exact algorithms in any of these treewidths unless P = NP, on ordinary binary input without supplied prime factorizations.

Composite cofactors can create cliques in H_P. For prime cofactors, H_P is a star, but a deterministic FPT algorithm would yield only the weaker consequence NP $\subseteq$ ZPP. A Cramér-type $O(\log^2 x)$ prime-gap bound would derandomize the window search: scan with

primality tests and allow n additional gaps for previous choices. See [36] for the prime-search distinction.

The explicit margin (14.3) shrinks with $t^{-2}T^{-2}$. Section 15.1 separates the consequences for exact computation, approximation and numerical input size.

15. Structure, complexity and open questions

The common starting point is local: saturation reduces residue choices to finite trees. From there, completion-safe frontier compression (Theorem 7.1) reduces the decision space, while signed elimination (Theorem 12.3) exploits the overlap graph to evaluate each decision. The two routes address different costs. Counting hardness with a supplied optimizer shows that a simple choice can have a difficult value; star threshold hardness shows that even a tree of arithmetic coordinates can encode a difficult choice. Hole-preserving distinctification also transfers the hardness of deciding complete coverage.

15.1. **Scope and limitations.** The hardness results concern exact computation. Supplied-vector density is additively approximable by sampling (Remark 10.5), and the reductions give no fixed additive gap for optimizing $M(A)$. The star instances always have $U(A) > 0$; exact-zero hardness comes from the separate list-to-set construction. Distinct-modulus covering is NP-hard and lies in Σ_2^P, but membership in NP remains unsettled.

Numerical size also matters. Using the unary dense-hardness theorem of the preprint [25], Corollary 9.2 establishes strong NP-hardness of packing. The three-prime counting construction also has polynomially bounded moduli. By contrast, the distinctification and fixed-centre star reductions produce polynomial-length binary encodings with potentially exponential numerical moduli; they do not establish strong hardness for those two families.

The realization theorem concerns explicit local states retained independently of their completion. Different completions require different states. Its deletion-model consequence does not bound arbitrary algorithms, implicit representations or approximate compression. Likewise, Wilf's distinction between finite exact procedures and perspicuous general answers [43] is a useful frame, not an impossibility theorem for symbolic expressions. Appendix B.1 records the formal-verification boundary, separating checked structural results from written algorithmic and complexity proofs.

Open directions. The remaining questions include the precise complexity of distinct-modulus covering (NP-hard and in Σ_2^P here, with NP membership unsettled), its complexity on unary-encoded distinct moduli, additive approximation of $M(A)$, unconditional deterministic prime-star hardness, fixed-centre star-threshold hardness with polynomial-magnitude moduli, and stronger completion-aware compression criteria.

Appendix A. Exact algorithms and certificates

A.1. **Scope-fused master algorithm.** The implementation follows Theorem 7.1: compute gcd-free states and their laminar laws, fuse scopes, reduce frontiers and contract exactly. Each profile stores a state representative for CRT reconstruction. Singleton tails store exponents rather than expanded digits; no step enumerates the numerical period.

Deletion certificates use rational mixtures of retained *other* profiles; retention certificates give strictly positive rational exposing functionals. Python computes these by two-phase rational simplex with Bland's rule and checks them separately without simplex. The proved Lean solver instead uses Fourier–Motzkin feasibility. Implementation entry points and detailed run instructions are in the package README.

A.2. **Signed noncoprime-graph algorithm.** The companion implements Theorem 12.3 and the prime-graph comparison route of Remark 12.8. Signed elimination uses binary subset indicators, not masks or decisions, and needs no frontier. Its order finder is a *min-fill heuristic*, not Bodlaender's algorithm; supplied orders are also accepted and their induced

widths measured. The theorem's guarantees use a supplied decomposition or the cited fixed-parameter finder.

A.3. **Instances beyond literal period enumeration.** Take the certified private primes

$$\begin{aligned}(p_i) = (&10000010975622725633,\ 10000065951204114433,\\ &10000092339483181057,\ 10000129722878525441,\\ &10000138518971547649).\end{aligned}$$

Each p_i-1 is divisible by $2^{40} > \sqrt{p_i}$. With witnesses $(3,5,5,3,7)$, respectively, the companion checks $a^{p_i-1} \equiv 1 \pmod{p_i}$ and $\gcd(a^{(p_i-1)/2}-1, p_i) = 1$. These certify primality: every prime divisor of p_i must be 1 mod 2^{40} by the multiplicative-order condition, hence exceed $\sqrt{p_i}$. The solver receives only the moduli, not their factorizations.

The mixed instance

$$A = \{2p_1, 4p_2, 6p_3, 9p_4\}$$

has a 78-digit period and a 79-digit residue-vector count. No CD residue vector exists: the first three moduli have pairwise gcd 2. Neither star formula applies, since all four moduli have gcd 1. The graph G_A is a triangle with one pendant edge, of treewidth 2. The signed algorithm examines three saturated state products, with buckets of at most eight entries. The fused algorithm, with and without pruning, gives the same exact value; factorization is disabled during these runs.

For a hand-checkable answer, put $\alpha_i = 1/n_i$. The prime 3 separates indices 3 and 4; at prime 2 exactly one pair among the first three is merged. Inclusion–exclusion therefore gives

$$\begin{aligned}M(A) = &\sum_i \alpha_i - \alpha_4(\alpha_1+\alpha_2)\\ &- \min\{2\alpha_1\alpha_2(1-\alpha_4),\ 2\alpha_1\alpha_3,\ 2\alpha_2\alpha_3\}.\end{aligned}$$

The last loss is smallest for these primes. The factor $1-\alpha_4$ in the first loss accounts for the compatible triple $\{1,2,4\}$; all other triples vanish. A separate implementation enumerates all 72 vectors modulo the compatibility moduli $(2,2,6,3)$ and uses generalized-CRT inclusion–exclusion, independently confirming the value and recovered residues. This benchmark avoids literal period enumeration without relying on CD or the star formula; it is not a lower bound against other algorithms.

As controls, the five-modulus star $\{3p_i : 1 \le i \le 5\}$ has a 96-digit period, 25 saturated states and a separately checked 41-partition formula. Its winning partition is $\{1\} \mid \{2,5\} \mid \{3,4\}$.

For increasing numbers of moduli, let $n_i = \ell_i\ell_{i+1}$ with consecutive primes beginning at 2. The graph G_A is a path, every local state is forced, and every signed bucket has at most four entries. Independent checking uses the recurrence $Q_0 = 1$, $Q_1 = 1 - 1/n_1$, and $Q_i = Q_{i-1} - Q_{i-2}/n_i$. For 32, 128 and 512 moduli, the periods have 53, 301 and 1,567 digits respectively. Every output agrees exactly with the recurrence and every recovered vector is CD. The companion records full results and illustrative timings.

A.4. **Arithmetic-necessity certificates.** For Theorem 13.2, check the weights, Q, η, the lower bound on t, distinct prime cofactors $m_i \neq p$, and the rational windows. These certify the gap $1/(2pQ^2t^2)$ without competitor enumeration. For Corollary 13.6, replace primality by gcd checks and verify the explicit difference-product construction. Small fixtures additionally check every partition into at most p blocks.

Appendix B. Formal verification and computational reproducibility

B.1. **Scope of the Lean companion.** The entry point `lean/PublicationScope.lean` audits the selected closure's transitive axioms. `PUBLICATION_SCOPE.md` maps manuscript claims to complete endpoint types and hypotheses.

Checked structural theory. The companion checks density and CRT; Shearer–CD including boundary cases; saturation, membership laws, frontier optimization and reconstruction for supplied pairwise-coprime integer bases; zero overload; the prime-cofactor star formula; and arithmetic necessity. The supplied bases include exponents and modulus product identities. The LCM identity is derived, and the formal optimizer attains the original residue minimum. This covers Proposition 4.5(i)–(iii) and its local state-count bound. Arithmetic necessity includes distinct prime cofactors, every competitor and a positive rational gap, with the analytic premise discharged.
Checked components. Checked components include packing padding, graph-colouring correspondence, four-prime counting queries and recovery, hole certificates, prime-coordinate state counts, the generic prime-pool lemma and supplied-decomposition table bounds.
Written proofs. The complexity classifications, completeness of the source problems, reduction and verifier machines, and runtime bounds are written proofs. So are basis construction; global composite-base state-product and generation bounds; singleton-frontier compression; signed G_A algorithms and CD recognition; witness-bit and deletion-model bounds; deterministic coprime windows and realization; the coprime-star extension and both threshold reductions; three-prime counting, including its unary consequence; two-factor tractability; the density counterexample; hole-preserving distinctification and its consequences; and strong packing. Python is not Lean-verified; matching formal endpoint types to mathematical claims remains a review obligation.

B.2. **Dependencies, trust and reproduction.** The audit uses Lean `v4.32.2` (official release), with mathlib pinned to `905b95818eb32af7874a58b427f50c1711a5e96c`. Other pins are in `lake-manifest.json`. The prime-window input uses the complete `WeakPNT` proof of `prime_between` from the pinned PrimeNumberTheoremAnd source [35], with LeanArchitect annotations. Unused unfinished upstream material is excluded. `ANALYTIC_DEPENDENCY.md` retains provenance and licenses for these external proofs.

The audit permits only `propext`, `Classical.choice` and `Quot.sound`, including at the analytic endpoint. It rejects `sorryAx` and native-evaluation axioms. The selected closure also passes fresh proof replay with the stock kernel of that release; this is neither an independent kernel nor verification of the kernel itself.

From the companion root, run `make lean` to build and audit the selected imports, and `make lean-replay` for fresh-kernel replay. Every command must exit zero. The declaration map records the expanded commands and source hashes; files outside its closure support no claim here.

B.3. **Separate finite computational audits.** The fused and signed-graph routes agree with raw-residue optimization on all 1,023 nonempty subsets of $\{2, \dots, 12\}$ of size at most five, with CRT checks of recovered vectors. Audits cover frontier certificates, both counting constructions, CD recognition, star thresholds, constructive realization, noncanonical distinctification extraction and the density counterexample; Appendix A.3 gives large-period examples, including an independently enumerated mixed-prime benchmark. The README and JSON records supply counts, seeds and exact inputs/outputs. All decisions are rational. Prime-star tests use certified finite windows, not the worst-case search height. These are finite correctness checks; shared generators limit their independence.

Acknowledgments

This work benefited from research assistance by AI systems developed by OpenAI and Anthropic, including support for proof exploration, proof development, and exact computational checks among others.

References

[1] S. Adenwalla, *A question of Erdős and Graham on covering systems*, Integers **26** (2026), A52. arXiv:2501.15170.

[2] M. Agrawal, N. Kayal and N. Saxena, *PRIMES is in P*, Annals of Mathematics **160** (2004), no. 2, 781–793. doi:10.4007/annals.2004.160.781.

[3] E. Bach, J. Driscoll and J. Shallit, *Factor refinement*, Journal of Algorithms **15** (1993), no. 2, 199–222. doi:10.1006/jagm.1993.1038.

[4] D. J. Bernstein, *Factoring into coprimes in essentially linear time*, Journal of Algorithms **54** (2005), no. 1, 1–30. doi:10.1016/j.jalgor.2004.04.009.

[5] T. F. Bloom, *Erdős problem #278*, Erdős Problems, `https://www.erdosproblems.com/278`.

[6] H. L. Bodlaender, *A linear-time algorithm for finding tree-decompositions of small treewidth*, SIAM Journal on Computing **25** (1996), no. 6, 1305–1317. doi:10.1137/S0097539793251219.

[7] S. Cambie, *Proving it is impossible; on Erdős problem #278*, arXiv:2508.18270v2, 2026.

[8] V. Conitzer and T. Sandholm, *Complexity of (iterated) dominance*, Proceedings of the Sixth ACM Conference on Electronic Commerce, 2005, 88–97. doi:10.1145/1064009.1064019.

[9] G. Dahl, *Disjoint congruence classes and a timetabling application*, Discrete Applied Mathematics **157** (2009), no. 8, 1702–1710. doi:10.1016/j.dam.2008.11.015.

[10] R. Dechter, *Bucket elimination: A unifying framework for reasoning*, Artificial Intelligence **113** (1999), 41–85.

[11] F. Eisenbrand, N. Hähnle, M. Niemeier, M. Skutella, J. Verschae and A. Wiese, *Scheduling periodic tasks in a hard real-time environment*, ICALP 2010, LNCS **6198**, Springer, 299–311. doi:10.1007/978-3-642-14165-2_26.

[12] P. Erdős and R. L. Graham, *Old and New Problems and Results in Combinatorial Number Theory*, Monographies de L'Enseignement Mathématique **28**, 1980.

[13] M. Filaseta, K. Ford, S. Konyagin, C. Pomerance and G. Yu, *Sieving by large integers and covering systems of congruences*, Journal of the American Mathematical Society **20** (2007), 495–517.

[14] E. C. Freuder, *Complexity of k-tree structured constraint satisfaction problems*, Proceedings of AAAI-90, 1990, 4–9.

[15] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979. Problem AN2, p. 249.

[16] C. Greenhill, *The complexity of counting colourings and independent sets in sparse graphs and hypergraphs*, Computational Complexity **9** (2000), no. 1, 52–72. doi:10.1007/PL00001601. Theorem 3.1.

[17] H. Halberstam and K. F. Roth, *Sequences*, second edition, Springer-Verlag, New York, 1983. Section 5.3, pp. 242–244; theorem attributed to C. A. Rogers.

[18] J. Harrington, J. Klein, J. Lowrance and O. Trifonov, *Covering systems where the prime divisors of all moduli are only* 2*,* 3*, or* 5, arXiv:2605.18644v1, 2026.

[19] N. J. A. Harvey and J. Vondrák, *Short proofs for generalizations of the Lovász local lemma: Shearer's condition and cluster expansion*, arXiv:1711.06797, 2017.

[20] R. D. Hough and P. P. Nielsen, *Covering systems with restricted divisibility*, Duke Mathematical Journal **168** (2019), no. 17, 3261–3295. doi:10.1215/00127094-2019-0058.

[21] Z. Jia, H. Li and Y. Liu, *Resolving Adenwalla's conjecture related to a question of Erdős and Graham about covering systems*, arXiv:2504.09579v3, 2025.

[22] R. M. Karp, *Reducibility among combinatorial problems*, Complexity of Computer Computations, Plenum Press, 1972, 85–103. doi:10.1007/978-1-4684-2001-2_9.

[23] A. Kawamura, Y. Kobayashi and Y. Kusano, *Pinwheel covering*, CIAC 2025, LNCS **15680**, Springer, 185–199. doi:10.1007/978-3-031-92935-9_12.

[24] R. Kleinberg and A. Mishra, *NP-Hardness and a PTAS for the Pinwheel Problem*, arXiv:2604.13974v1, 2026; to appear in FOCS 2026.

[25] Y. Kobayashi, B. Lin and J. Swernofsky, *Dense Pinwheel Packing Is Strongly NP-Complete*, arXiv:2609.20075v1, 2026.

[26] Y. Kobayashi, B. Lin and J. Swernofsky, *Hardness and fixed parameter tractability for pinwheel scheduling problems*, Theoretical Computer Science **1077** (2026), 115998. doi:10.1016/j.tcs.2026.115998.

[27] Y. Kobayashi and B. Lin, *Hardness and fixed parameter tractability for pinwheel scheduling problems*, ISAAC 2025, LIPIcs **359** (2025), 47:1–47:15. doi:10.4230/LIPIcs.ISAAC.2025.47.

[28] J. H. M. Korst, E. H. L. Aarts, J. K. Lenstra and J. Wessels, *Periodic multiprocessor scheduling*, PARLE 1991, LNCS **505**, 166–178. doi:10.1007/BFb0035103.

[29] J. H. M. Korst, E. H. L. Aarts and J. K. Lenstra, *Scheduling periodic tasks*, INFORMS Journal on Computing **8** (1996), no. 4, 428–435. doi:10.1287/ijoc.8.4.428.

[30] The Lean contributors, *Lean 4*, proof assistant and programming language. `https://lean-lang.org/`.

[31] The mathlib Community, *The Lean mathematical library*, Proceedings of CPP 2020, 367–381. doi:10.1145/3372885.3373824. arXiv:1910.09336.

[32] A. Mok, L. Rosier, I. Tulchinsky and D. Varvel, *Algorithms and complexity of the periodic maintenance problem*, Microprocessing and Microprogramming **27** (1989), nos. 1–5, 657–664. doi:10.1016/0165-6074(89)90128-2.

[33] K. O'Bryant, *On Z.-W. Sun's disjoint congruence classes conjecture*, Combinatorial Number Theory, de Gruyter, 2007, 403–411.
[34] Y. Onishi, *Exact Optimization of Congruence-Class Unions: Fixed-parameter algorithms and a Lean verification*, preprint, version 1.2, 11 September 2026. `https://github.com/yoshiyoyoyo44/erdos278/tree/v1.2`.
[35] The PrimeNumberTheoremAnd contributors, *PrimeNumberTheoremAnd*, Lean formalisation project, commit `a5154676af9aa3095150ee410cdda80555aa0642`. `https://github.com/AlexKontorovich/PrimeNumberTheoremAnd`.
[36] D. H. J. Polymath, *Deterministic methods to find primes*, arXiv:1009.3956, 2010.
[37] A. D. Scott and A. D. Sokal, *The repulsive lattice gas, the independent-set polynomial, and the Lovász local lemma*, Journal of Statistical Physics **118** (2005), nos. 5–6, 1151–1261. doi:10.1007/s10955-004-2055-4.
[38] J. B. Shearer, *On a problem of Spencer*, Combinatorica **5** (1985), no. 3, 241–245.
[39] R. J. Simpson, *Exact coverings of the integers by arithmetic progressions*, Discrete Mathematics **59** (1986), 181–190. doi:10.1016/0012-365X(86)90079-8. Lemma 2.3.
[40] Z.-W. Sun, *On disjoint residue classes*, Discrete Mathematics **104** (1992), 321–326. doi:10.1016/0012-365X(92)90454-N.
[41] T. Trudgian, *Updating the error term in the prime number theorem*, Ramanujan Journal **39** (2016), no. 2, 225–234. doi:10.1007/s11139-014-9656-6.
[42] W. D. Wei and C. L. Liu, *On a periodic maintenance problem*, Operations Research Letters **2** (1983), no. 2, 90–93. doi:10.1016/0167-6377(83)90044-5.
[43] H. S. Wilf, *What is an answer?*, American Mathematical Monthly **89** (1982), no. 5, 289–292. doi:10.1080/00029890.1982.11995435.