An Exact Fixed-Parameter Algorithm for Extremal Unions of Residue Classes

Alper Ferudun*

September 2, 2026

Abstract

Let $A = \{n_1 < \cdots < n_r\}$ be a finite set of positive integers. Choose one residue class a_i mod n_i for each modulus and maximize the natural density of their union. This is the unsettled maximum-density half of Erdős Problem 278. We give an exact uniform characterization whose state space depends on r, not on the magnitudes of the moduli. To a residue tuple we attach the graph in which ij is an edge precisely when $a_i \equiv a_j \pmod{\gcd(n_i, n_j)}$. Inclusion–exclusion makes the covered density a clique-weighted function of this graph. For every forced edge set, a finite-abelian-group kernel calculation counts compatible tuples by a Smith-normal-form lattice index. Boolean Möbius inversion then counts the tuples with each exact graph. Maximizing over the graphs with positive count gives the exact extremal density.

The resulting factoring-free algorithm uses $2^{O(r^2)}\,\mathrm{poly}(B)$ bit operations, where B is the binary input length. We also give an independent prime-power layer formula for the kernel counts, a reduction to gcd kernels, and a factorization over the connected components of the non-coprimality graph. The construction is compatible with known arithmetic-coloring and abelian-arrangement machinery; the contribution is its exact-stratum composition with the Erdős–Graham density objective.

1 Introduction

For a finite set $A = \{n_1 < \cdots < n_r\}$ of positive integers, choose residues a_i mod n_i and consider

$$U_A(a) = \bigcup_{i=1}^{r} (a_i + n_i\mathbb{Z}).$$

Erdős and Graham asked for the minimum density left uncovered by a suitable choice of residues, and whether the opposite extremum is obtained when all residues are aligned [EG80, p. 28]. In the complementary language used here, the first question asks for the maximum covered density

$$M(A) = \max_a \operatorname{dens} U_A(a).$$

The aligned-residue assertion for the minimum covered density was proved by Simpson; see also Sun's formulation [Sim86, Sun91]. The maximum half is still listed as open by the maintained problem record [Blo26].

The problem is finite for each input: every union is periodic modulo $L = \operatorname{lcm}(n_1, \ldots, n_r)$. Exhausting all $\prod_i n_i$ residue tuples, however, is not an informative uniform answer. Recent work of

*Mercury Software GmbH. Email: alper@mercurycodelab.com. GitHub: `https://github.com/AlperTheKing`.

Cambie gives exact formulae for structured common-core families and shows hardness phenomena, including NP-hardness for the list version in which moduli may repeat [Cam25]. Dahl studied the pairwise-disjoint special case and integer-programming formulations [Dah09]. Strong asymptotic results in restricted ranges of distinct large moduli were obtained by Filaseta, Ford, Konyagin, Pomerance, and Yu [FFK$^+$07].

This note gives a uniform exact characterization for every finite set of distinct moduli. Its key finite state is the pairwise Chinese-remainder compatibility graph. There are $2^{\binom{r}{2}}$ possible graphs, regardless of the sizes of the n_i. Two ingredients then suffice:

(i) inclusion–exclusion expresses the density solely in terms of the cliques of a compatibility graph;

(ii) finite-group kernel counts and Boolean Möbius inversion determine exactly which graphs occur.

The kernel count can be computed without factoring by a Smith normal form. It also has a prime-by-prime component formula, furnishing an independent derivation and a useful implementation check. The resulting algorithm is fixed-parameter tractable in r, with running time $2^{O(r^2)}\,\mathrm{poly}(B)$.

Arithmetic graph colorings and abelian arrangements already contain the general point-counting machinery behind the forced-edge calculation [DM13, LTY21]. We therefore make no novelty claim for Smith indices or arrangement Möbius inversion in isolation. The exact formula below is the composition of that machinery with the clique-weighted objective specific to this covering-density problem.

2 Compatibility graphs and the density functional

Fix $A = \{n_1 < \cdots < n_r\}$, and write

$$\Omega_A = \prod_{i=1}^{r} \mathbb{Z}/n_i\mathbb{Z}, \qquad \mathcal{E} = \binom{[r]}{2}, \qquad d_{ij} = \gcd(n_i, n_j).$$

Definition 2.1. For $a = (a_1, \ldots, a_r) \in \Omega_A$, its *compatibility graph* $\Gamma_A(a) \subseteq \mathcal{E}$ is defined by

$$ij \in \Gamma_A(a) \quad \Longleftrightarrow \quad a_i \equiv a_j \pmod{d_{ij}}.$$

For a graph $H \subseteq \mathcal{E}$, put

$$\Phi_A(H) = \sum_{\substack{\varnothing \neq S \subseteq [r] \\ \binom{S}{2} \subseteq H}} (-1)^{|S|+1} \frac{1}{\mathrm{lcm}(n_i : i \in S)}. \tag{1}$$

Singletons are included in this sum.

Lemma 2.2 (Clique criterion)**.** *For a nonempty $S \subseteq [r]$, the classes $a_i \bmod n_i$, $i \in S$, have a common integer if and only if S is a clique of $\Gamma_A(a)$. If the intersection is nonempty, it is one residue class modulo* $\mathrm{lcm}(n_i : i \in S)$.

Proof. The generalized Chinese remainder theorem says that $x \equiv a_i \pmod{n_i}$, $i \in S$, is soluble if and only if

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

These are precisely the clique conditions. The usual CRT description of the solution set gives one class modulo the least common multiple. □

Proposition 2.3 (Graph-only density)**.** *For every $a \in \Omega_A$,*

$$\operatorname{dens} U_A(a) = \Phi_A(\Gamma_A(a)).$$

Proof. Apply finite inclusion–exclusion to the selected residue classes. By lemma 2.2, a term indexed by S is nonzero exactly when S is a clique, and then its density is the reciprocal of the corresponding least common multiple. □

3 Counting exact compatibility graphs

For $T \subseteq \mathcal{E}$, orient the edges of T arbitrarily and set

$$q = |T|, \qquad D_T = \operatorname{diag}(d_e : e \in T) \in \mathbb{Z}^{q \times q}.$$

Let $B_T \in \mathbb{Z}^{q \times r}$ be the signed edge–vertex incidence matrix. Thus the row of an oriented edge $e = (i, j)$ has $+1$ at i, -1 at j, and zero elsewhere. Define

$$h_A(T) = [\mathbb{Z}^q : D_T\mathbb{Z}^q + B_T\mathbb{Z}^r], \tag{2}$$

with $h_A(\varnothing) = 1$. The index is finite because $D_T\mathbb{Z}^q$ has full rank.

Equivalently, $h_A(T)$ is the product of the q nonzero Smith invariant factors of the matrix $[D_T \mid B_T]$, or the gcd of its $q \times q$ minors. Reversing an edge changes one row by a sign and does not change the index.

Let

$$N_A(T) = \#\{a \in \Omega_A : T \subseteq \Gamma_A(a)\}.$$

Lemma 3.1 (Forced-edge count)**.** *For every $T \subseteq \mathcal{E}$,*

$$N_A(T) = \left(\prod_{i=1}^{r} n_i\right) \frac{h_A(T)}{\prod_{e \in T} d_e}. \tag{3}$$

Proof. Consider the homomorphism

$$\psi_T : \Omega_A \longrightarrow \bigoplus_{e \in T} \mathbb{Z}/d_e\mathbb{Z}, \qquad a \longmapsto B_T a \pmod{D_T}.$$

It is well defined: if $e = ij$, then $d_e \mid n_i, n_j$, so changing a representative a_i by n_i changes every incident target coordinate by zero. Its kernel consists exactly of the tuples compatible on all edges of T, hence has size $N_A(T)$.

The image is generated by the columns of B_T. Consequently

$$\operatorname{coker} \psi_T \cong \mathbb{Z}^q/(D_T\mathbb{Z}^q + B_T\mathbb{Z}^r),$$

whose order is $h_A(T)$. Since the target has order $\prod_{e \in T} d_e$,

$$|\operatorname{im} \psi_T| = \frac{\prod_{e \in T} d_e}{h_A(T)}.$$

The finite-group first isomorphism theorem now yields equation (3). □

For $H \subseteq \mathcal{E}$, let

$$X_A(H) = \#\{a \in \Omega_A : \Gamma_A(a) = H\}.$$

Every tuple counted by $N_A(T)$ has a unique exact graph containing T, so

$$N_A(T) = \sum_{H \supseteq T} X_A(H). \tag{4}$$

Boolean-lattice Möbius inversion gives

$$X_A(H) = \sum_{T \supseteq H} (-1)^{|T \setminus H|} N_A(T). \tag{5}$$

In particular,

$$H \text{ is realizable} \quad \Longleftrightarrow \quad X_A(H) > 0. \tag{6}$$

Although equation (5) is alternating, its result is nonnegative because it counts an exact finite stratum.

4 The exact extremal formula

Theorem 4.1 (Exact compatibility-graph formula)**.** *For every finite set $A = \{n_1 < \cdots < n_r\}$ of positive integers,*

$$\boxed{M(A) = \max_{\substack{H \subseteq \mathcal{E} \\ X_A(H) > 0}} \Phi_A(H),} \tag{7}$$

where Φ_A is given by equation (1)*, X_A by equation* (5)*, N_A by equation* (3)*, and $h_A(T)$ by equation* (2)*. Equivalently, the least uncovered density is*

$$1 - M(A).$$

Proof. Every residue tuple a has an exact graph $H = \Gamma_A(a)$, and proposition 2.3 gives density $\Phi_A(H)$. Conversely, equation (6) says that each graph retained on the right side of equation (7) occurs as $\Gamma_A(a)$ for at least one tuple. Thus the values maximized on the two sides are exactly the same finite set. □

Remark 4.2 (Compulsory coprime edges)*.* If $d_{ij} = 1$, then ij belongs to every compatibility graph. Formula equation (5) automatically gives zero to graphs omitting such an edge; no separate deletion convention is needed.

Remark 4.3 (Degenerate cases)*.* For $r = 1$, $\mathcal{E} = \varnothing$ and equation (7) gives $M(A) = 1/n_1$. If $1 \in A$, one selected class is all of $\mathbb{Z}$, and the formula simplifies to $M(A) = 1$. Repeated moduli are outside the literal set formulation; they should be retained as labelled vertices for the list variant.

5 A prime-power formula

The Smith-index proof is factoring-free. A second expression, which uses the prime factorizations of the n_i, exposes the combinatorial structure and independently checks lemma 3.1.

For a prime p, put $v_i = v_p(n_i)$. For $k \geq 1$, define

$$V_{p,k} = \{i : v_i \geq k\}, \qquad T_{p,k} = \{ij \in T : i, j \in V_{p,k}\}.$$

Let $c_{p,k}(T)$ be the number of connected components of $(V_{p,k}, T_{p,k})$, counting isolated active vertices.

Proposition 5.1 (Prime-layer count)**.** *For every $T \subseteq \mathcal{E}$,*

$$N_A(T) = \prod_p p^{\sum_{k\geq 1} c_{p,k}(T)}. \tag{8}$$

Consequently,

$$h_A(T) = \prod_p p^{\sum_{k\geq 1}(|T_{p,k}|-|V_{p,k}|+c_{p,k}(T))}. \tag{9}$$

Proof. The CRT splits each coordinate into its p-primary coordinates. At the k-th base-p digit, a forced edge requires equality precisely when both endpoint valuations are at least k. One digit may be chosen independently for each connected component of $(V_{p,k}, T_{p,k})$, giving $p^{c_{p,k}(T)}$ choices. Multiplication over levels and primes proves equation (8).

The exponent of p in $\prod_i n_i$ is $\sum_k |V_{p,k}|$, while its exponent in $\prod_{e\in T} d_e$ is $\sum_k |T_{p,k}|$. Comparing equation (8) with equation (3) gives equation (9). The exponent at each layer is the cycle rank of its graph. □

6 Algorithm and bit complexity

Let

$$B = \sum_{i=1}^{r} \lceil \log_2(n_i + 1) \rceil, \qquad m = \binom{r}{2}.$$

An exact evaluator proceeds as follows.

1. Compute every d_{ij}. For each $T \subseteq \mathcal{E}$, compute a Smith normal form of $[D_T \mid B_T]$, and then $N_A(T)$ from equation (3).
2. Apply the superset Möbius transform equation (5) to obtain every $X_A(H)$.
3. For each H with $X_A(H) > 0$, evaluate $\Phi_A(H)$ exactly and take the maximum.

Theorem 6.1 (Fixed-parameter exact evaluation)**.** *The value $M(A)$ can be computed without integer factorization in*

$$2^{O(r^2)} \operatorname{poly}(B)$$

bit operations and space. Hence exact evaluation is fixed-parameter tractable in the number r of moduli.

Proof. There are 2^m forced-edge sets and 2^m exact graphs. Smith and Hermite normal forms of integer matrices are computable in polynomial bit complexity [KB79]. Here all matrix dimensions are bounded by a function of r, and all entries have $O(B)$ bits. The counts and lattice indices have bit lengths polynomial in r and B.

The fast superset Möbius transform uses $O(m2^m)$ exact additions and subtractions. Direct evaluation of $\Phi_A(H)$ uses at most $2^r - 1$ terms per graph. All denominators divide $L = \operatorname{lcm}(n_1, \ldots, n_r)$, whose bit length is at most B, so density comparisons are exact polynomial-bit-length integer comparisons. Combining these bounds proves the claim. □

Remark 6.2. The exponential dependence on r is essential to the present method and is consistent with known hardness results. The theorem is stronger than enumeration over $\prod_i n_i$ tuples when the moduli are large, but it does not assert polynomial time for unrestricted r, a closed elementary expression for the optimizer, or a classification of maximizing tuples.

7 Two structural reductions

Corollary 7.1 (Gcd-kernel reduction)**.** *For $r > 1$, put*

$$g_i = \operatorname{lcm}_{j \neq i} \gcd(n_i, n_j),$$

and put $g_1 = 1$ when $r = 1$. Then $g_i \mid n_i$, and the optimization defining $M(A)$ may be carried out over $\prod_i \mathbb{Z}/g_i\mathbb{Z}$ instead of Ω_A.

Proof. Every compatibility test incident to i reads a_i only modulo a divisor of g_i; hence the full graph, and by proposition 2.3 the density, depends only on $a_i \bmod g_i$. Conversely, each reduced residue lifts modulo n_i because $g_i \mid n_i$. □

Corollary 7.2 (Non-coprimality component factorization)**.** *Form the graph on $[r]$ with edge ij when $\gcd(n_i, n_j) > 1$, and let its connected components be $C_1, \ldots, C_s$. If $A_C = \{n_i : i \in C\}$, then*

$$1 - M(A) = \prod_{k=1}^{s} \bigl(1 - M(A_{C_k})\bigr). \tag{10}$$

Proof. The least common multiples of the moduli in distinct components are pairwise coprime. The CRT therefore identifies a full period with the product of the component periods. For fixed residue choices, the density missed by every component is the product of the component missed densities. Choices in different components are independent, and minimizing this nonnegative product factorizes. □

8 Examples and exact checks

Example 8.1. For $A = \{2, 4\}$, there is one optional edge with $d_{12} = 2$. Equation (3) gives $N_A(\varnothing) = 8$ and $N_A(\{12\}) = 4$, so each exact graph has four realizing tuples. The compatible graph has density $1/2$, whereas the incompatible graph has density $1/2 + 1/4 = 3/4$. Thus

$$M(\{2, 4\}) = \frac{3}{4}.$$

Example 8.2. For $A = \{2, 3\}$, the sole edge has $d_{12} = 1$ and is compulsory. Both forced counts are six, so inversion gives zero tuples to the graph without the edge. The unique realizable graph has density

$$M(\{2, 3\}) = \frac{1}{2} + \frac{1}{3} - \frac{1}{6} = \frac{2}{3}.$$

A reference implementation accompanies this manuscript. Its main deterministic test suite compares three logically distinct routes:

1. direct enumeration of residue tuples;
2. forced-edge lattice indices computed as gcds of maximal minors, followed by Boolean inversion;
3. prime-power layer component counts, followed by the same inversion.

The three routes agree on every distinct modulus set of size at most four from $\{2,\dots,7\}$, and on three additional five-modulus cases: 59 sets, 4,228 forced-edge instances, and 221 realizable exact graphs. Two six-modulus stress cases check 65,536 further forced-edge instances. The suite also verifies the gcd-kernel and component reductions for every nonempty subset of $\{2,\dots,8\}$, and independently recovers

$$M(\{2,3,\dots,10\}) = \frac{69}{70}.$$

These finite checks are falsification attempts; the proof of theorem 4.1 does not depend on them.

The source package also contains a deliberately separate standard-library implementation. It imports no code from the main checker and compares the prime-layer formula with both exact graph counts and literal covered-set sizes from direct residue enumeration. It passes on 101 modulus sets, comprising 24,745 residue tuples and 3,665 exact graph-count comparisons. This includes every set of size at most four from $\{2,\dots,8\}$, together with selected degenerate and five-modulus cases.

9 Relation to arithmetic colorings

Set $L = \operatorname{lcm}(n_1,\dots,n_r)$. The forced-edge count fits naturally into prior work on arithmetic graph colorings. Give an edge $e = ij \in T$ the label L/d_e, and color every vertex in $\mathbb{Z}/L\mathbb{Z}$. The dotted-edge constraint of D'Adderio and Moci [DM13],

$$\frac{L}{d_e}c_i = \frac{L}{d_e}c_j \pmod{L},$$

is equivalent to $c_i \equiv c_j \pmod{d_e}$. Reducing $c_i \bmod L$ to $a_i \bmod n_i$ has exactly L/n_i lifts in each coordinate and preserves every constraint. Their arithmetic-coloring count is therefore

$$\left(\prod_i \frac{L}{n_i}\right) N_A(T).$$

More generally, G-Tutte polynomials and abelian Lie group arrangements organize point counts of kernel arrangements through finite abelian groups and Smith invariants [LTY21]. From that viewpoint, equation (5) is exact-stratum inversion for a finite subgroup arrangement. What is specific here is the choice of varying edge targets $\mathbb{Z}/d_{ij}\mathbb{Z}$, followed by the clique-weighted density functional Φ_A and optimization over the nonempty exact strata.

10 Limitations and interpretation

Theorem 4.1 gives the exact value for every literal finite input and a uniform algorithm whose state space is controlled by the number of moduli. It does not make the optimization simple in r, produce a closed formula without a finite combinatorial maximum, or classify all maximizing residues. Cambie's hardness results explain why one should not expect an unrestricted polynomial-time formula [Cam25].

The historical phrase *what is the maximum* is open-ended. We therefore state the mathematical output precisely as an exact effective characterization and fixed-parameter algorithm. Whether a problem list chooses to call that a complete resolution is a matter of interpretation, not an additional theorem.

Reproducibility

The source package includes a standard-library Python checker, an independent small reimplementation, deterministic test descriptions, and SHA-256 hashes for the audited source files. Exact integer and rational arithmetic is used throughout.

The discovery, exposition, literature triage, and computational testing were carried out with substantial assistance from an AI system under human direction. AI output was not treated as authority: the argument was rederived obligation by obligation, the main count was obtained in two independent mathematical forms, and the formula was compared against direct enumeration. This manuscript is an unrefereed preprint. Author attribution and final responsibility must be supplied by the submitting human author.

References

[Blo26] Thomas F. Bloom. Erdős problem #278, 2026. Maintained problem record, accessed 29 August 2026.

[Cam25] Stijn Cambie. Proving it is impossible; on Erdős problem #278, 2025. Version 2, revised 17 August 2026.

[Dah09] Geir Dahl. Disjoint congruence classes and a timetabling application. *Discrete Applied Mathematics*, 157(8):1702–1710, 2009.

[DM13] Michele D'Adderio and Luca Moci. Graph colorings, flows and arithmetic Tutte polynomial. *Journal of Combinatorial Theory, Series A*, 120(1):11–27, 2013.

[EG80] Paul Erdős and Ronald L. Graham. *Old and New Problems and Results in Combinatorial Number Theory*, volume 28 of *Monographies de L'Enseignement Mathématique*. L'Enseignement Mathématique, Geneva, 1980.

[FFK+07] Michael Filaseta, Kevin Ford, Sergei Konyagin, Carl Pomerance, and Gang Yu. Sieving by large integers and covering systems of congruences. *Journal of the American Mathematical Society*, 20(2):495–517, 2007.

[KB79] Ravindran Kannan and Achim Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM Journal on Computing*, 8(4):499–507, 1979.

[LTY21] Ye Liu, Tan Nhat Tran, and Masahiko Yoshinaga. G-Tutte polynomials and abelian Lie group arrangements. *International Mathematics Research Notices*, 2021(1):150–188, 2021.

[Sim86] R. J. Simpson. Exact coverings of the integers by arithmetic progressions. *Discrete Mathematics*, 59(1–2):181–190, 1986. MR 87f:11011.

[Sun91] Zhi-Wei Sun. A theorem concerning systems of residue classes. *Acta Mathematica Universitatis Comenianae*, 60(1):123–131, 1991.