LOWER BOUNDS FOR THE MINIMAL LEAST COMMON MULTIPLE OF A DISTINCT COVERING SYSTEM WITH MINIMUM MODULUS $m \geq 8$

KOREA SUPERINTELLIGENCE LABS

ABSTRACT. A covering system is *distinct* when its moduli are pairwise different, and $L_{\min}(m)$ denotes the least value of the least common multiple of the moduli over all distinct covering systems of $\mathbb{Z}$ whose smallest modulus is m. The values $L_{\min}(3) = 120$, $L_{\min}(4) = 360$, $L_{\min}(5) = 1440$ and $L_{\min}(6) = 5040$ are known, and $L_{\min}(7) = 10080$ has recently been claimed; above $m = 7$ the literature records only upper bounds, and the case $m = 8$ has been posed as a problem for which integer programming is "currently not computationally feasible". We prove

$$L_{\min}(8) \geq 7200, \quad L_{\min}(9) \geq 7920, \quad L_{\min}(10) \geq 10080, \quad L_{\min}(11) \geq 27720, \quad L_{\min}(12) \geq 25200,$$

which together with Krukenberg's construction sandwiches $L_{\min}(8)$ between 7200 and 50400. Every published value that a bound consumes is carried as an explicit hypothesis. The engine is an exclusion certificate in exact integer arithmetic: a recursive sharpening of the reciprocal-sum bound along coprime splittings of the period, with no solver anywhere. The same certificate excludes twelve of the eighteen candidates for $L_{\min}(7)$ without an integer program, and shows that every covering of $\mathbb{Z}$ by distinct odd moduli greater than 1 has least common multiple at least 45045, closing three cells left open by the previously certified bound of 10^4. All statements are machine-checked in Lean 4.

1. INTRODUCTION

A *covering system* is a finite set of congruence classes

$$a_1 \bmod n_1, \ \ldots, \ a_r \bmod n_r, \qquad n_i \geq 2,$$

whose union is $\mathbb{Z}$; it is *distinct* if the moduli $n_1, \ldots, n_r$ are pairwise different. Erdős introduced covering systems in 1950 [3], and the two quantities that organise the subject are the smallest modulus and the least common multiple L of the moduli. For $m \geq 2$ let

$$L_{\min}(m) \ = \ \min\big\{ \operatorname{lcm}(n_1, \ldots, n_r) \ : \ \{a_i \bmod n_i\} \text{ is a distinct covering system with } \min_i n_i = m \big\}.$$

Dalton and Trifonov [2] proved $L_{\min}(3) = 120$ and $L_{\min}(4) = 360$; Klein [6] proved $L_{\min}(5) = 1440$ and $L_{\min}(6) = 5040$, and conjectured that $L_{\min}(7) \geq 15120$, in these words: "If the smallest modulus of a covering system with distinct moduli is 7, then the least common multiple of the moduli is at least 15120" ([6, Conjecture 2], verbatim). Zhang and Zhang [12] refuted that conjecture and claimed the value $L_{\min}(7) = 10080$: an explicit 66-class system for the upper bound, and for the lower bound a reciprocal-sum filter followed by three rounds of integer programming, ending in two complete Gurobi runs of 13789.26 and 13159.68 seconds.

Above $m = 7$ nothing of the kind is available. The only published data are *upper* bounds: Krukenberg's constructions [7] give $L_{\min}(8) \leq 2^5 \cdot 3^2 \cdot 5^2 \cdot 7 = 50400$ and $L_{\min}(9) \leq 2^5 \cdot 3^3 \cdot 5^2 \cdot 7 = 151200$, as tabulated in the introduction of [2], and Harrington, Klein, Lowrance and Trifonov [5, Theorem 1.11] exhibit a distinct covering system with minimum modulus 8 and $L = 2^8 \cdot 3^3 \cdot 5^2 = 172800$. The same paper states the situation as an open problem; the Theorem 1.2 it refers to is Klein's theorem above, restated there ([5, Problem 3]; theorem numbers are those of v1):

> "Find all triples of positive integers (a, b, c) with $a \geq b \geq c$ such that there exists a distinct covering system with minimum modulus $m = 8$ and an LCM of its moduli $L = 2^a 3^b 5^c$. … there is no analogue of Theorem 1.2 in the case $m = 8$. Finally, solving Problem 3 using integer programming is currently not computationally feasible."

The elision drops the opening of the paragraph that follows the problem, which lists the difficulties one at a time, and in particular the sentence "most of the covering systems involved will contain more than one hundred congruences".

A second, older question is reached by the same machinery. Erdős and Selfridge asked whether a covering system exists all of whose moduli are odd, distinct and greater than 1; this is still open. Mian and Siddique [11] recently gave a machine-checked exclusion — any such system has $L > 10000$ — and recorded exactly where their method stops. The underlying arithmetic is settled much further by McNew and Setty [10], who determine every primitive covering number below 773500, and all but one below 10^6, by an integer-programming pipeline; that decides every value we treat in that part of this paper, where what is at stake is therefore the reach of a certificate, not the truth of the statement.

Results. Write $\mathrm{LB}(m, B)$ for the assertion $L_{\min}(m) \geq B$, in the precise form given in Definition 2.1. Our main results are the following two theorems, proved in Section 7. To our knowledge no lower bound for $L_{\min}(m)$ with $m \geq 8$ has appeared before; see Remark 7.6 for the search behind that statement.

- *From Klein's theorem alone* (Theorem 7.1): $L_{\min}(8) \geq 5040$, $L_{\min}(9) \geq 5040$ and $L_{\min}(10) \geq 10080$, assuming only the peer-reviewed value $L_{\min}(6) = 5040$ of [6].
- *From the full published chain* (Theorem 7.2), which in addition consumes the claim $L_{\min}(7) = 10080$ of [12]: $L_{\min}(8) \geq 7200$, $L_{\min}(9) \geq 7920$, $L_{\min}(11) \geq 27720$ and $L_{\min}(12) \geq 25200$.

In tabular form, with the entries proved here in bold and the published values that enter as hypotheses in italics:

m	3	4	5	6	7	8	9	10	11	12
lower bound	*120*	*360*	*1440*	*5040*	10080†	**7200**	**7920**	**10080**	**27720**	**25200**
upper bound	120	360	1440	5040	10080	50400	151200	—	—	—

The dagger marks the one hypothesis that is not yet refereed, the claim of [12]: the new entries at $m = 8, 9, 11, 12$ carry it, and the one at $m = 10$ does not. A dash means that no upper bound is quoted in the sources consulted here; at $m \leq 7$ the upper bounds are attained values, and the one at $m = 7$ is Theorem 4.1.

Combining the second of these with Krukenberg's construction sandwiches $L_{\min}(8)$ between 7200 and 50400, a factor of seven. These are bounds and not values: we do not determine $L_{\min}(m)$ for any $m \geq 8$, and Section 7 explains why that should not be attempted by this route.

The engine is a decidable exclusion certificate in exact integer arithmetic (Section 3). Three further consequences of it are worth stating separately.

- *Odd coverings* (Theorem 6.1): every covering system of $\mathbb{Z}$ whose moduli are distinct, greater than 1 and all divisors of an odd L satisfies $L \geq 45045$. This includes the three values 10395, 12285, 17325 that [11] records as beyond its certificate, and it moves the certified bound for the Erdős–Selfridge problem from 10^4 to 45045.
- *Solver-free exclusions at $m = 7$* (Theorem 5.1): twelve of the eighteen candidates that survive the reciprocal-sum filter of [12] are excluded by a closed arithmetic inequality over the divisors of L. In [12] each of these twelve is the output of an integer program whose only content is the word *infeasible*.
- *The residual debt* (Theorem 5.3): exactly six cells stand between the certificate and the lower bound $L_{\min}(7) \geq 10080$, namely $L \in \{5040, 6300, 6720, 7560, 8400, 9240\}$, and the implication "these six plus Klein's theorem give $L_{\min}(7) \geq 10080$" is itself proved. Section 5 reports what the six cost.

What is new, and what is not. The certificate is a multiplicative sharpening of the reciprocal-sum bound along a coprime splitting of the period. In its one-step form this is old, and published at least four times: it is the pairwise-coprime case of [4, Lemma 2.1]; Dalton and Trifonov run it by hand in [2, §4], where at $L = 280$ they write that "the congruences modulo 4, 5, and 7 cover a portion

of the integers with density $1 - \frac{3}{4} \cdot \frac{4}{5} \cdot \frac{6}{7} = \frac{17}{35}$", and their Lemma 15 counts the uncovered residues at $L = 240$ as a product over coprime coordinates; [5, Theorem 1.9] is a Bonferroni truncation of the same idea over pairwise-coprime tuples, conditional on $L = 2^a 3^b 5^c$; and the *capacity certificate* of [11] is a one-step bound of the same shape, charging one class in each member of a pairwise-coprime family of divisors and everything else flat (Remark 3.4). *No claim of novelty is made for it.*

What we add is the recursion — a coprime block is charged by the same bound applied inside it, to any depth, and the best value over all splittings is taken at every node — as a single decidable test, and the reach that this buys. That reach, as Remark 7.5 records, comes almost entirely from charging a coprime *block* as one coordinate, that is from a single splitting, and only marginally from nesting; the recursion is the uniform way of computing the block bounds, and should not be credited with more than that. That the recursion is the natural next step is not our observation either: [11] names it as future work, writing that extending its bound "requires a sharper certificate (natural candidates: charging prime-power towers rather than one prime per coprime slot, …)". Balister, Bollobás, Morris, Sahasrabudhe and Tiba [1] also process the primes of the period one at a time, viewing their measures through the Chinese remainder theorem as product measures, but theirs is a first- and second-moment sieve on a measure of the uncovered set, aimed at asymptotic theorems; it is not a per-L inequality between integers.

Two of our results are new only as routes. The twelve exclusions of Theorem 5.1 are theorems of [12]; the odd-covering bound of Theorem 6.1 is subsumed by [10]. In both cases the mathematics is known and the contribution is that the verdict is now a checkable arithmetic inequality rather than the exit status of a solver — the position [11] takes about its own bound. Only the lower bounds for $L_{\min}(m)$ at $m \geq 8$ are offered as new mathematics.

2. Preliminaries

2.1. **Covering systems, and the two statement forms.** Throughout, a system is given by a finite set $M \subseteq \mathbb{Z}_{\geq 2}$ of moduli together with a residue a_n for each $n \in M$; that the moduli form a *set* is what makes the system distinct, and it is the whole content of the problem (Proposition 7.4(4)). We say M (with residues a) *covers* if every integer x satisfies $x \equiv a_n \pmod{n}$ for some $n \in M$.

Definition 2.1. For $m, B \geq 1$, $\mathrm{LB}(m, B)$ is the statement: whenever M is a finite set of moduli with $m \in M$ and $n \geq m$ for all $n \in M$, the residues a make M cover $\mathbb{Z}$, and L is a positive integer divisible by every $n \in M$, then $L \geq B$.

For $m, L \geq 1$, $\mathrm{NoCov}(m, L)$ is the statement: there is no finite set M of moduli with $n \geq m$ and $n \mid L$ for all $n \in M$, and no choice of residues, covering $\mathbb{Z}$.

Since a positive common multiple of the moduli is at least their least common multiple, $\mathrm{LB}(m, B)$ says exactly $L_{\min}(m) \geq B$, and we use the two interchangeably. We state it with an arbitrary common multiple because that is the form in which it composes: the descent of Proposition 3.7 produces a new system with the *same* L, not with a recomputed least common multiple.

2.2. **Divisor completion.** Fix m and L and write

$$\mathcal{D}_m(L) \;=\; \{\, d : d \mid L,\ d \geq m \,\}.$$

If a distinct system with all moduli $\geq m$ and dividing L covers $\mathbb{Z}$, then so does the system obtained by giving every $d \in \mathcal{D}_m(L)$ a class — the moduli that were already present keep their residues, and the added ones only cover more. So the moduli set may be taken to be $\mathcal{D}_m(L)$ *known*, and only the residues are free. This is the divisor completion used by [6] and [12], and it is what makes each pair (m, L) a finite question. For a list A of moduli dividing N and residues a put

$$u(N, A, a) \;=\; \#\{\, y \in \{0, 1, \ldots, N-1\} \;:\; y \not\equiv a_d \pmod{d} \text{ for every } d \in A \,\},$$

the number of uncovered residues in one period. Every bound below is a lower bound for $u(N, A, a)$ valid for *all* choices of a, and a covering forces $u(L, \mathcal{D}_m(L), a) = 0$.

2.3. **One period, and the reciprocal-sum bound.**

Proposition 2.2. *Let $L > 0$, let M be a finite set of moduli each dividing L, and let a be residues. If every y with $0 \le y < L$ satisfies $y \equiv a_n \pmod{n}$ for some $n \in M$, then M covers $\mathbb{Z}$.*

Proof. Given $x \in \mathbb{Z}$, apply the hypothesis to $y = x \bmod L$ and add $L \mid x - y$ to $n \mid y - a_n$, using $n \mid L$. □

Proposition 2.3 (Reciprocal-sum bound, integer form)**.** *Let A be a finite list of moduli, each positive and dividing N, and let a be any residues. Then*

$$u(N, A, a) \;\ge\; N - \sum_{d \in A} \frac{N}{d}.$$

In particular, if A covers one period then $\sum_{d \in A} 1/d \ge 1$.

Proof. A class modulo d meets $\{0, \dots, N-1\}$ in exactly N/d points when $d \mid N$; sum over A and subtract from N. □

This is the classical density obstruction, used by Erdős already in [3] and quoted as the reciprocal obstruction in [12, Lemma 1] and as the Mirsky–Newman lemma in [6, Lemma 5]. The form above is the non-strict one, which is what the certificate needs; the strict inequality $\sum 1/d > 1$ available at equality is never used here.

3. The certificate

3.1. **Splitting the period.** The reciprocal-sum bound charges two coprime moduli f_1, f_2 a total of $1/f_1 + 1/f_2$ of the period. But two classes with coprime moduli *must* meet, in density $1/(f_1 f_2)$, so the truth is $1 - (1 - 1/f_1)(1 - 1/f_2)$, which is smaller. Making that precise for a whole coprime splitting is the following pair of statements.

Lemma 3.1 (Chinese remainder product count)**.** *Let $N = N_1 N_2$ with $\gcd(N_1, N_2) = 1$ and $N_1, N_2 > 0$, and let P_1, P_2 be properties of residues modulo N_1 and N_2. Then*

$$\#\{\, y < N : P_1(y \bmod N_1) \text{ and } P_2(y \bmod N_2) \,\} \;=\; \#\{y < N_1 : P_1\} \cdot \#\{y < N_2 : P_2\}.$$

Proof. The map $y \mapsto (y \bmod N_1,\, y \bmod N_2)$ is a bijection from $\{0, \dots, N-1\}$ to the product of the two residue ranges; injectivity is the uniqueness half of the Chinese remainder theorem and surjectivity its existence half. □

Proposition 3.2 (Split bound)**.** *Let $N = N_1 N_2$ with $\gcd(N_1, N_2) = 1$, let A be a finite list of positive moduli dividing N, and split it as*

$$A_1 = \{d \in A : d \mid N_1\}, \qquad A_2 = \{d \in A : d \nmid N_1,\ d \mid N_2\}, \qquad B = A \setminus (A_1 \cup A_2).$$

Then for all residues a,

$$u(N_1, A_1, a) \cdot u(N_2, A_2, a) \;\le\; u(N, A, a) \;+\; \sum_{b \in B} \frac{N}{b}.$$

Proof. Whether a modulus dividing N_1 covers y depends only on $y \bmod N_1$, and likewise for N_2; so by Lemma 3.1 the set of $y < N$ missed by both A_1 and A_2 has exactly $u(N_1, A_1, a) \cdot u(N_2, A_2, a)$ elements. Each such y is either uncovered by all of A, or covered by some $b \in B$, and the classes of B account for at most $\sum_{b \in B} N/b$ points by the count in Proposition 2.3. □

The moduli of B — those straddling the splitting — are the loss: they are charged their full density $1/b$ against the product, exactly as in the reciprocal-sum bound. Everything in A_1 and A_2 is charged multiplicatively instead.

Example 3.3. Take $N = 10395 = 3^3 \cdot 5 \cdot 7 \cdot 11$ and let A be the 31 divisors of N greater than 1. The reciprocal sum is $\sum_{d \in A} 1/d = 12645/10395 > 1$, so Proposition 2.3 says nothing at all here. Split $N = 27 \cdot 385$. The moduli dividing 27 are $3, 9, 27$, and they leave at least $27 - (9 + 3 + 1) = 14$ of the 27 residues modulo 27 uncovered; the moduli dividing 385 but not 27 are $5, 7, 11, 35, 55, 77, 385$, and they leave at least $385 - 191 = 194$ of the 385 residues modulo 385. The remaining 21 moduli — those divisible both by 3 and by a prime of 385 — straddle the splitting, with $\sum_{b \in B} N/b = 2483$. Proposition 3.2 therefore gives

$$u(N, A, a) \;\geq\; 14 \cdot 194 - 2483 \;=\; 233$$

for every choice of residues, so no distinct covering system has all its moduli dividing 10395.

Remark 3.4 (computed outside the formal development). The capacity certificate of [11] excludes N when

$$\sum_{\substack{d \mid N,\ d>1 \\ d \notin T}} \frac{N}{d} \;<\; \frac{N}{\prod_{d \in T} d} \prod_{d \in T} (d-1)$$

for some pairwise-coprime family T of divisors of N greater than 1. At $N = 10395$ no T satisfies it; the best choice is $T = \{3, 5, 7, 11\}$, where the left side is 4671 and the right side 4320. That is the stall [11] reports, and comparing it with Example 3.3 isolates the difference exactly: there the tower $3, 9, 27$ is charged as a single coordinate, in which 14 of 27 residues survive, whereas one prime per coprime slot keeps $2/3$ of the slot 3 and then pays $N/9 + N/27 = 1540$ in full for the moduli 9 and 27.

3.2. **The recursion.** All arithmetic below is in the non-negative integers, with subtraction truncated at 0; this is harmless, since every quantity is a lower bound for a cardinality.

Definition 3.5. For a period N and a list A of moduli dividing N set

$$\Lambda_0(N, A) = N - \sum_{d \in A} \frac{N}{d},$$

and, for $f \geq 0$,

$$\Lambda_{f+1}(N, A) = \max\Big(\Lambda_0(N, A),\ \max_{\substack{N = N_1 N_2,\ \gcd(N_1, N_2) = 1 \\ N_1 > 1}} \Big[\Lambda_f(N_1, A_1) \cdot \Lambda_f(N_2, A_2) - \sum_{b \in B} \frac{N}{b}\Big]\Big),$$

with A_1, A_2, B as in Proposition 3.2.

Theorem 3.6 (Soundness)**.** *For every $f \geq 0$, every $N > 0$, every list A of positive moduli dividing N and every choice of residues a,*

$$\Lambda_f(N, A) \;\leq\; u(N, A, a).$$

Consequently, if $\Lambda_f(L, \mathcal{D}_m(L)) > 0$ for some f, then NoCov(m, L).

Proof. Induction on f. The base case is Proposition 2.3. For the step, the flat term is again Proposition 2.3, and each splitting term is Proposition 3.2 with the two factors bounded by the inductive hypothesis at N_1 and N_2; the maximum of lower bounds is a lower bound. For the consequence: a covering system with moduli $\geq m$ dividing L makes $u(L, \mathcal{D}_m(L), a) = 0$ after divisor completion, while $\Lambda_f > 0$ forces $u > 0$. □

Three features of Definition 3.5 deserve comment, because each was arrived at by measurement.

First, taking the *maximum* at every node rather than the first positive value is load-bearing: a parent splitting needs the largest bound available at each child, not merely a positive one. The cheaper first-positive variant was implemented first and loses $L = 5544$ and three further cells below 5040 in Proposition 5.2.

Second, the candidate splittings need no correctness proof. The test re-checks $N_1 > 1$, $N_1 \mid N$ and $\gcd(N_1, N/N_1) = 1$ and contributes 0 when any of these fails, so the list of candidates may

be generated by the cheapest enumeration available — trial division to $\sqrt{N}$, junk entries included. Only the *strength* of the certificate, never its soundness, depends on that enumeration.

Third, a block may be a prime-power tower or any union of prime-power components, and this, rather than the depth of the nesting, is where most of the reach comes from; see Remark 7.5.

3.3. **Descent and sweeps.** Two more ingredients turn per-cell exclusions into a lower bound.

Proposition 3.7 (Divisor descent)**.** *Let $k < m$ with $k \mid L$, and suppose* $\mathrm{LB}(k, B)$ *holds. If M is a finite set of moduli with $n \geq m$ and $n \mid L$ for all $n \in M$, and residues a make M cover $\mathbb{Z}$, then $L \geq B$.*

Proof. Adjoin the class $0 \bmod k$. Since every modulus of M is $\geq m > k$, the modulus k is new, so the enlarged set is again a set of distinct moduli; its smallest element is k, it still covers, and every one of its moduli still divides L. Now apply $\mathrm{LB}(k, B)$. □

This is the argument by which [12, §2] imports Klein's $L_{\min}(6) = 5040$, stated once for a general pair $k < m$ instead of once for $k = 6$; carrying $\mathrm{LB}(k, B)$ as a hypothesis is what keeps the dependence of our theorems on published values explicit.

Theorem 3.8 (Sweep)**.** *Let $m \geq 1$ and $B, f \geq 0$. Let $(k_1, B_1), \ldots, (k_r, B_r)$ be pairs with $k_i < m$ and* $\mathrm{LB}(k_i, B_i)$ *for each i, and let E be a finite set of positive integers with* $\mathrm{NoCov}(m, L)$ *for every $L \in E$. Suppose every multiple L of m with $0 < L < B$ satisfies at least one of*

1. *$L \in E$;*
2. *$k_i \mid L$ and $L < B_i$, for some i;*
3. *$\Lambda_f\big(L, \mathcal{D}_m(L)\big) > 0$.*

Then $\mathrm{LB}(m, B)$.

Proof. Let M, a and L be as in Definition 2.1 and suppose $L < B$. Then $m \in M$ gives $m \mid L$, so one of the three cases applies to L: the first and third contradict $\mathrm{NoCov}(m, L)$ — directly, or through Theorem 3.6 — and the second contradicts $L < B_i$ by Proposition 3.7. □

In every application below the recursion depth is $f = 6$.

3.4. **Cost.** Evaluating $\Lambda_f(N, A)$ is cheap. The coprime splittings of N are given by its unitary divisors, so a node offers at most $2^{\omega(N)-1} - 1$ of them, and every N in the sweeps of the formal development has at most five distinct prime factors; the divisor list $\mathcal{D}_m(L)$ is produced by trial division to $\sqrt{L}$, pairing each small divisor with its cofactor. The largest single computation in this paper, the sweep of Proposition 6.2 over 100 000 periods, takes about four minutes.

4. The record at $m = 7$, re-checked

Theorem 4.1. *The* 66 *congruence classes printed in* [12, §7] *cover $\mathbb{Z}$. Their moduli are pairwise distinct, the smallest is* 7*, the largest is* 10080 *and all of them divide* 10080*. Hence $L_{\min}(7) \leq 10080$; in fact* $\mathrm{LB}(7, 10081)$ *fails, and so does* $\mathrm{LB}(7, 15120)$*, which is* [6, Conjecture 2].

Proof sketch. The system is the one printed in [12]; what is carried out here is its verification. By Proposition 2.2 it suffices to check one period, which is an exhaustive scan of the 10080 residues, each tested against the 66 classes — 665 280 modular comparisons, with zero residues left uncovered. The shape claims (distinctness, 66 moduli, all dividing 10080, smallest 7) are read off the table in the same evaluation. The two negative statements follow because the system itself is a counterexample. □

The two halves of the main claim of [12] are thereby separated: the refutation of Klein's Conjecture 2 needs no solver and is settled outright, while the matching lower bound is the part that rested on integer programming, and is the subject of Section 5.

The lower-bound apparatus of [12] starts from the reciprocal-sum filter over the raw candidate set R_7 of all 720 multiples of 7 in $[5040, 10080)$. We recompute its output from its own definitions, as a check on the reading of the source.

Proposition 4.2. *Write* $S_m(L) = \sum_{d \in \mathcal{D}_m(L)} 1/d$. *Then:*

(1) *the* $L \in R_7$ *with* $S_7(L) > 1$ *are exactly the eighteen values printed in* [12, §4]*, namely* 5040*,* 5544*,* 5880*,* 6048*,* 6300*,* 6552*,* 6720*,* 6930*,* 7056*,* 7392*,* 7560*,* 8064*,* 8400*,* 8568*,* 8820*,* 9072*,* 9240*,* 9576*;*

(2) *the positive multiples* $L < 5040$ *of* 7 *with* $S_7(L) > 1$ *are exactly the nine values printed in* [12, §2]*, namely* 1260*,* 1680*,* 2520*,* 3024*,* 3360*,* 3780*,* 4032*,* 4200*,* 4620*, and every one of them is divisible by* 6*;*

(3) $\mathcal{D}_7(9240)$ *has* 58 *elements, and for the twenty-element subset* $A \subseteq \mathcal{D}_7(9240)$ *of the partial-sum argument of* [12, §6] *one has*

$$\sum_{d \in \mathcal{D}_7(9240) \setminus A} \frac{1}{d} \;=\; \frac{1754}{9240} \;=\; \frac{877}{4620}.$$

Proof sketch. Exhaustive evaluation over the 720 and 719 candidates respectively, with $S_7(L) > 1$ written as the integer inequality $L < \sum_{d \in \mathcal{D}_7(L)} L/d$ so that no rational arithmetic occurs. All three lists agree with [12] element for element. □

5. Solver-free exclusions, and the six cells that remain

Theorem 5.1. NoCov$(7, L)$ *holds for each of*

$$L \in \{5544,\ 5880,\ 6048,\ 6552,\ 6930,\ 7056,\ 7392,\ 8064,\ 8568,\ 8820,\ 9072,\ 9576\}.$$

Proof sketch. Each is one evaluation of $\Lambda_6(L, \mathcal{D}_7(L)) > 0$ followed by Theorem 3.6. The evaluation is a finite arithmetic computation over the divisors of L: for $L = 5544$, for instance, the value is 114, so at least 114 of the 5544 residues escape every admissible system. □

All twelve are theorems of [12], where they are the outcome of integer programs running 5.13 to 17.90 seconds each and reporting infeasibility. The mathematics is therefore not new; what the theorem adds is that the verdict is now a closed inequality between integers, checkable without trusting a solver.

Proposition 5.2 (Reach)**.** *Over the whole raw candidate set* R_7 *— all* 720 *multiples of* 7 *lying in the interval* $[5040, 10080)$*, not merely the eighteen that survive the filter — the certificate* $\Lambda_6(\cdot, \mathcal{D}_7(\cdot)) > 0$ *fails at exactly six values,*

$$5040, \quad 6300, \quad 6720, \quad 7560, \quad 8400, \quad 9240,$$

and over the 719 *positive multiples of* 7 *below* 5040 *it fails at exactly two, namely* 2520 *and* 3360*. Moreover the two dense partial covers printed in* [12] *cover* 4857 *of the* 5040 *residues and* 7416 *of the* 7560 *residues respectively.*

Proof sketch. Two exhaustive sweeps, of 720 and 719 cells, each cell one evaluation of Λ_6; and two exhaustive scans of a period against a printed list of 40 and 58 classes. □

So twelve of the eighteen candidates fall to exact arithmetic, and below 5040 the descent through Klein's theorem — which [12] applies to all nine values of its own list — is needed at two values only. The two near-covers explain why 5040 and 7560 resist: 183 and 144 residues out of a full period is far less than any density certificate, this one included, can detect.

Theorem 5.3. *Assume* LB$(6, 5040)$*, i.e. Klein's* $L_{\min}(6) = 5040$*. Then* LB$(7, 5040)$*. If in addition* NoCov$(7, L)$ *holds for each of the six values* $L \in \{5040, 6300, 6720, 7560, 8400, 9240\}$*, then* LB$(7, 10080)$*, i.e.* $L_{\min}(7) \geq 10080$*.*

Proof sketch. Two applications of Theorem 3.8 with $f = 6$ and the single descent pair $(6, 5040)$: a sweep over the 719 positive multiples of 7 below 5040, and a sweep over the 1439 positive multiples of 7 below 10080 with the six cells supplied as the escape set E. □

The point of stating it this way is that the outstanding debt becomes a hypothesis one can read, rather than a remark. The six cells are expensive, and it is worth recording how expensive, since the shape of the cost decides what a next attempt should look like.

Remark 5.4 (Prices; measured outside the formal development). In [12] the cheap two of the six, 6300 and 6720, cost 21 and 132 seconds of integer programming; 8400 and 9240 need its partial-sum program; and 5040 and 7560 cost two complete Gurobi runs of 13789.26 and 13159.68 seconds on at most 32 threads — on the order of 120 CPU-hours apiece — which produce no certificate at all. A propositional encoding does not help: the divisor-completed instance at $L = 6300$, with three classes preset by the Chinese-remainder normalisation of [12] and a sequential at-most-one encoding, has 45 046 variables and 73 800 clauses, and a run of `cadical --lrat` had not terminated after 5 minutes 40 seconds, having emitted 5.08 gigabytes of proof log at about 15 MB/s. A cell of this kind will fall to a sharper arithmetic bound or not at all.

Remark 5.5 (How far the certificate misses; measured outside the formal development). Running the recursion of Definition 3.5 without the truncation at 0, so that the value may be negative, measures how much more uncovered density would be needed. As a fraction of L the best value at the six cells is

L	deficit	L	deficit	L	deficit
6300	$-17/1260$	8400	$-353/4200$	7560	$-41/210$
6720	$-69/1120$	9240	$-19/154$	5040	$-187/840$

So 6300 is the nearest miss: 85 residues out of 6300. Its best splitting is $6300 = 900 \cdot 7$, and it would suffice to know that at least 0.3557 of $\mathbb{Z}/900\mathbb{Z}$ escapes the moduli dividing 900, where the recursion proves only 0.34. Replacing the recursive bound at a small child by the exact minimum of u there is therefore the first thing to try.

6. Coverings by odd moduli

At $m = 2$ the bound of Proposition 2.3 is exactly the deficiency test: since $\sum_{d \mid N,\, d \geq 2} N/d = \sigma(N) - N$, it is positive if and only if $\sigma(N) < 2N$. Every deficient N is therefore excluded outright, and the content of the next theorem lies entirely on the odd N with $\sigma(N) \geq 2N$ — of which there are 134 below 60 000, the smallest being 945 (a count made outside the formal development). The recursion excludes 133 of those; the exception is 45045.

Theorem 6.1. *Let L be odd and positive, let M be a finite set of moduli with $n \geq 2$ and $n \mid L$ for every $n \in M$, and let residues a make M cover $\mathbb{Z}$. Then $L \geq 45045$. In particular* $\mathrm{NoCov}(2, N)$ *holds for $N = 10395 = 3^3 \cdot 5 \cdot 7 \cdot 11$, $N = 12285 = 3^3 \cdot 5 \cdot 7 \cdot 13$ and $N = 17325 = 3^2 \cdot 5^2 \cdot 7 \cdot 11$.*

Proof sketch. All moduli divide the odd number L, so they are odd, distinct and > 1. An exhaustive sweep over the 22 522 odd values $N < 45045$ verifies $\Lambda_6(N, \mathcal{D}_2(N)) > 0$ for every one of them; Theorem 3.6 then gives $\mathrm{NoCov}(2, N)$ for each, and no such L can be smaller than 45045. At the three named values the certificate returns 555, 861 and 831 uncovered residues respectively. □

The bound itself is not new: [10] determines every primitive covering number below 773500, which decides every N here. What is new is that the exclusion is checkable. The previous certified bound, $L > 10000$ [11], stopped at exactly the three values named above; its Limitations section states that the certificate "stalls at the four-prime-factor stragglers 10395, 12285, 17325 — verified inside Lean, not merely observed — so extending the certified bound past 10^4 by this route requires a sharper certificate (natural candidates: charging prime-power towers rather than one prime per coprime slot, …)". A block of the certificate above is exactly such a tower, charged as one coordinate, and that is why the three cells close.

Proposition 6.2 (The odd frontier)**.** *Among the 100 000 odd values $N < 200000$ the certificate $\Lambda_6(N, \mathcal{D}_2(N)) > 0$ fails at exactly five:*

$$45045 = 3^2 \cdot 5 \cdot 7 \cdot 11 \cdot 13, \quad 135135 = 3^3 \cdot 5 \cdot 7 \cdot 11 \cdot 13, \quad 155925 = 3^4 \cdot 5^2 \cdot 7 \cdot 11,$$

$$176715 = 3^3 \cdot 5 \cdot 7 \cdot 11 \cdot 17, \quad 197505 = 3^3 \cdot 5 \cdot 7 \cdot 11 \cdot 19.$$

Every survivor has four or five distinct odd prime factors — the shape that stopped [11] one rung earlier. The first of them, 45045, is the closest miss anywhere in this work: the untruncated recursion of Remark 5.5 returns -171 out of 45045, a deficit of 0.38%. It is the named next target, as 10395 was for [11].

Proposition 6.3 (Controls)**.** *The classical system* $\{0 \bmod 2, 0 \bmod 3, 1 \bmod 4, 1 \bmod 6, 11 \bmod 12\}$*, printed in the introduction of* [6]*, covers* $\mathbb{Z}$*. The certificate does not fire there:* $\Lambda_6(12, \mathcal{D}_2(12)) = 0$*. Nor does it fire at the frontier:* $\Lambda_6(45045, \mathcal{D}_2(45045)) = 0$*.*

The first control is the soundness test that a positive certificate at $N = 12$ would fail; the second says that Theorem 6.1 stops where the method stops, rather than where a computation was cut off. These are the two controls [11] runs for its own bound, at the corresponding values.

7. Lower bounds for $L_{\min}(m)$, $m \geq 8$

Theorem 7.1. *Assume* $\mathrm{LB}(6, 5040)$*, that is, the lower-bound half of Klein's theorem* $L_{\min}(6) = 5040$ [6, Theorem 4]*. Then*

$$L_{\min}(8) \geq 5040, \qquad L_{\min}(9) \geq 5040, \qquad L_{\min}(10) \geq 10080.$$

Proof sketch. Three applications of Theorem 3.8 with $f = 6$, empty escape set and the single descent pair $(6, 5040)$. The sweeps run over the 629 multiples of 8 below 5040, the 559 multiples of 9 below 5040, and the 1007 multiples of 10 below 10080; each cell is excluded either by the descent (when $6 \mid L$ and $L < 5040$) or by an evaluation of Λ_6. In the third sweep the descent disposes of the 167 cells that are divisible by 6 and smaller than 5040, and the certificate of the other 840. □

Theorem 7.2. *Assume the published lower bounds* $\mathrm{LB}(3, 120)$ *and* $\mathrm{LB}(4, 360)$ *of* [2]*, the bounds* $\mathrm{LB}(5, 1440)$ *and* $\mathrm{LB}(6, 5040)$ *of* [6]*, and* $\mathrm{LB}(7, 10080)$ *of* [12]*. Then*

$$L_{\min}(8) \geq 7200, \quad L_{\min}(9) \geq 7920, \quad L_{\min}(11) \geq 27720, \quad L_{\min}(12) \geq 25200,$$

where $7200 = 2^5 \cdot 3^2 \cdot 5^2$*,* $7920 = 2^4 \cdot 3^2 \cdot 5 \cdot 11$*,* $27720 = 2^3 \cdot 3^2 \cdot 5 \cdot 7 \cdot 11$ *and* $25200 = 2^4 \cdot 3^2 \cdot 5^2 \cdot 7$*.*

Proof sketch. Four applications of Theorem 3.8 with $f = 6$ and empty escape set, chained: the descent pairs for $m = 8$ are the five hypotheses; for $m = 9$ they are those together with $(8, 7200)$; for $m = 11$, with $(8, 7200), (9, 7920), (10, 10080)$; and for $m = 12$, with those and $(11, 27720)$, where $L_{\min}(10) \geq 10080$ is Theorem 7.1. The sweeps run over the 899 multiples of 8 below 7200, the 879 multiples of 9 below 7920, the 2519 multiples of 11 below 27720 and the 2099 multiples of 12 below 25200. In the four sweeps the descent disposes of 349, 409, 358 and 780 cells respectively, and the certificate of the remaining 550, 470, 2161 and 1319. No cell is left over, which is the assertion. □

Corollary 7.3. *Under the hypotheses of Theorem 7.2,* $7200 \leq L_{\min}(8) \leq 50400$*; the upper bound is Krukenberg's construction as tabulated in* [2]*, and needs no hypothesis.*

The upper bound is quoted at second hand and it is worth saying exactly what is quoted. The introduction of [2] states that "for m between 3 and 18 except for $m = 7$, Krukenberg constructed distinct covering systems with least modulus m while trying to keep the least common multiple L of all moduli small", and then tabulates L for $3 \leq m \leq 9$; the entries for $m = 8$ and $m = 9$ are $2^5 \cdot 3^2 \cdot 5^2 \cdot 7 = 50400$ and $2^5 \cdot 3^3 \cdot 5^2 \cdot 7 = 151200$. The least modulus is thus exactly m, which is what makes these bounds on $L_{\min}(8)$ and $L_{\min}(9)$. Krukenberg's dissertation [7] itself was not consulted.

The two theorems together read as follows, the middle column being the one that consumes the unrefereed value $L_{\min}(7) = 10080$ and the left column the one that does not.

m	from $L_{\min}(6) = 5040$	with $L_{\min}(7) = 10080$	best upper bound
8	≥ 5040	≥ 7200	50400 [2]
9	≥ 5040	≥ 7920	151200 [2]
10	≥ 10080	≥ 10080	—
11	—	≥ 27720	—
12	—	≥ 25200	—

Proposition 7.4 (Controls)**.** (1) $\mathrm{LB}(7, 10081)$ *is false, so no bound proved here at* $m = 7$ *can be pushed past* 10080*, and the certificate must be blind somewhere in* $[5040, 10080)$*.*

(2) $\Lambda_6(10080, \mathcal{D}_7(10080)) = 0$*: it is blind exactly at the value that carries a covering system.*

(3) $\Lambda_6(7200, \mathcal{D}_8(7200)) = 0$*: the bound* $L_{\min}(8) \geq 7200$ *stops where the method stops, and is not an arbitrary cut-off.*

(4) $\mathrm{NoCov}(3, 3)$ *holds. Three congruence classes modulo* 3 *do cover* $\mathbb{Z}$*; a* set *of moduli holds* 3 *once, and one class modulo* 3 *covers a third of* $\mathbb{Z}$*. So distinctness carries the whole question.*

(5) $\mathrm{LB}(2, 13)$ *is false, since the system of Proposition 6.3 has least common multiple* 12*. A bound of the shape proved here is thus a real constraint and not an empty quantification; in particular* $L_{\min}(2) \leq 12$*.*

Remark 7.5 (Where the recursion is actually needed; measured outside the formal development). Of the 714 cells of R_7 that the certificate excludes, 713 are already excluded at $f = 1$: 701 of them by the flat bound of Proposition 2.3 alone, and twelve more by a single coprime splitting with both children bounded by that same proposition. The sole exception is $L = 5544$, where $\Lambda_1 = 0$ and $\Lambda_2 = 54$. The whole odd sweep below 45045 likewise goes through at $f = 1$. So the reach reported here comes mostly from charging a coprime *block* as one coordinate rather than one prime per slot — the step [11] names — and only marginally from nesting. Deeper nesting still pays where it applies: at the three stragglers of Theorem 6.1 the values at $f = 1, 2, 3$ are $233, 542, 555$ and $511, 848, 861$ and $166, 807, 831$. The formal development runs at $f = 6$ throughout, so these f-by-f values are read off the independent implementation of Definition 3.5 described in Section 8; it reproduces every $f = 6$ value stated in this paper exactly.

Remark 7.6. We have not found a lower bound for $L_{\min}(m)$ at any $m \geq 8$ in the literature. The search behind that statement, run on 2026-08-28: a local full-text arXiv corpus, in which 455 distinct covering-system papers were extracted and searched for the phrases "minimal / minimum / smallest / least possible least common multiple" (seven hits, three of them in [12], the only paper using the notation $L_{\min}(m)$, and four unrelated) and "least common multiple of the moduli … at least" (four hits, all in [6], for 1440, 5040 and 15120); the forward citation graphs of [6] and [2]; and the On-Line Encyclopedia of Integer Sequences, where the sequences $120, 360, 1440, 5040, 10080$ and $12, 36, 60, 108, 168$ both return nothing. The integers 7200, 7920, 27720 and 25200 occur in no covering-system paper of the corpus as a bound on $L_{\min}$. On that date [6] had exactly two citing papers, [12] and [5], neither of which states a lower bound at any $m \geq 8$, and [12] itself was still at v1 with no recorded citations — which is also why its value $L_{\min}(7) = 10080$ is carried here as a hypothesis rather than used silently. The positive evidence that the question is open is [5, Problem 3], quoted in Section 1.

What should not be attempted. Theorem 7.2 gives bounds, not values, and the gap is wide: at $m = 8$ it is a factor of seven. Closing it means an existence search over a divisor-completed system with more than a hundred congruences, which is what [5, Problem 3] calls not computationally feasible; Remark 5.4 shows what that costs already at $m = 7$, where the answer is known. The productive direction is the opposite one: sharpen the bound inside a small coprime block, as Remark 5.5 indicates, and the frontier moves at $L = 6300$ and at $N = 45045$ simultaneously. Extending the table in m is by contrast cheap: continuing the same chain one and two steps further gives $L_{\min}(13) \geq 65520$

and $L_{\min}(14) \geq 55440$, at sweep sizes of 5039 and 3959 cells, and at both values the certificate stops where the sweep does. These last two bounds are computed with the independent implementation of Section 8 only; they are not part of the machine-checked development, they are quoted here as an indication of cost, and nothing else in this paper uses them.

8. Verification

Every statement in this paper has been formally verified in Lean 4 (toolchain `v4.32.0`) against *Mathlib* [8, 9], in a development of five files and 67 theorems that is free of `sorry` and of added axioms; the repository is private at the time of writing, and a repository reference is to be added at submission. The reasoning layer is kernel-checked with no appeal to compiled evaluation: this covers Propositions 2.2, 2.3, 3.2 and 3.7, Lemma 3.1, the soundness Theorem 3.6 of the certificate, and the sweep Theorem 3.8, all of which depend on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`. What is delegated to compiled evaluation, through Lean's `native_decide`, is exactly the finite tables: the period scan of Theorem 4.1; the filter lists of Proposition 4.2; one Boolean per cell in Theorem 5.1; one per sweep in Proposition 5.2, Theorems 5.3, 6.1, 7.1, 7.2 and Proposition 6.2; and one per control in Propositions 6.3 and 7.4. Each of these is a decidable Boolean whose truth is the exhaustive search named in the corresponding proof sketch; the soundness bridge from the Boolean to the mathematical statement is in the kernel-checked layer in every case. Elapsed times, in import order, are 12.3, 8.3, 22.5, 260.1 and 21.9 seconds, the fourth being dominated by the sweep of Proposition 6.2. Every table was produced first by an independent Python implementation of the definitions of Section 3, and the formal statement was written against it and then checked to agree element for element; this applies in particular to the six lists and the exact rational of Proposition 4.2, which are the reading of [12] on which Sections 4 and 5 depend.

References

[1] P. Balister, B. Bollobás, R. Morris, J. Sahasrabudhe and M. Tiba, *On the Erdős covering problem: the density of the uncovered set*, Invent. Math. **228** (2022), 377–414; arXiv:1811.03547.

[2] J. Dalton and O. Trifonov, *Extreme covering systems*, J. Integer Seq. **25** (2022), Article 22.9.1.

[3] P. Erdős, *On integers of the form $2^k + p$ and some related problems*, Summa Brasil. Math. **2** (1950), 113–123.

[4] M. Filaseta, K. Ford, S. Konyagin, C. Pomerance and G. Yu, *Sieving by large integers and covering systems of congruences*, J. Amer. Math. Soc. **20** (2007), 495–517; arXiv:math/0507374.

[5] 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).

[6] J. Klein, *On a conjecture of Krukenberg and a problem of Dalton and Trifonov*, Integers **26** (2026), #A38, doi:10.5281/zenodo.19402698; arXiv:2508.18062.

[7] C. E. Krukenberg, *Covering sets of the integers*, Ph.D. dissertation, University of Illinois, Urbana-Champaign, 1971. (Not consulted; every value attributed to it here is quoted from [2].)

[8] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction — CADE 28, Lecture Notes in Computer Science 12699, Springer, 2021, pp. 625–635.

[9] The mathlib Community, *The Lean mathematical library*, in: CPP 2020, ACM, 2020, pp. 367–381.

[10] N. McNew and J. Setty, *On the densities of covering numbers and abundant numbers*, Math. Comp., published electronically 22 June 2026, doi:10.1090/mcom/4209; arXiv:2507.23041.

[11] I. Mian and S. Siddique, *Kernel-checked exclusions for the Erdős–Selfridge odd covering problem: any odd covering of $\mathbb{Z}$ has lcm exceeding 10000*, arXiv:2607.25628v1 (2026).

[12] S. Zhang and J. Zhang, *A distinct covering system with minimum modulus 7 and minimal least common multiple 10080*, arXiv:2607.19029v1 (2026).