THE MINIMUM NUMBER OF MONOCHROMATIC PERMUTED COPIES IN A TWO-COLOURING OF THE REGULAR (2^k-1)-GON

KOREA SUPERINTELLIGENCE LABS

Abstract. Put $n = 2^k - 1$ points evenly on a circle. A *permuted copy* is a k-subset whose cyclic gaps are a permutation of $1, 2, 4, \dots, 2^{k-1}$; Stromquist conjectured that every two-colouring of the n points contains a monochromatic permuted copy, and Damásdi, Frankl, Pach and Pálvölgyi, who verified the conjecture for $k \le 7$ by computer, ask for the *minimum number* $\mu(k)$ of monochromatic permuted copies over all two-colourings. A one-page problem-session note of Stromquist announces, with no proof and no method, the three values $\mu(3) = 2$, $\mu(4) = 4$, $\mu(5) = 10$; no other value of μ, and no proof of these, appears to be on record. We prove $\mu(k) \le (k-1)!$ for every $k \ge 2$, witnessed by the parity colouring $\chi(x) = x \bmod 2$ on the representatives $0, \dots, n-1$; in particular $\mu(8) \le 5040$ unconditionally, at the first case left open by the computer searches. We give proofs of Stromquist's three values, the last of them by an exhaustive branch-and-bound over all 2^{31} colourings whose pruning rule is proved sound, and we give explicit colourings showing $\mu(6) \le 20$ and $\mu(7) \le 120$. The structural tool is an exact restatement: a permuted copy is precisely a translate of the σ-image of a maximal chain of the Boolean lattice $2^{[k]}$ truncated to its bottom k levels, where $\sigma(S) = \sum_{i \in S} 2^i$. Every theorem, proposition and lemma below is machine-checked in Lean 4.

1. Introduction

Fix an integer $k \ge 2$, put

$$n \;=\; 2^k - 1,$$

and place n points evenly on a circle, so that the point set is the cyclic group $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$. The gaps $1, 2, 4, \dots, 2^{k-1}$ sum to $2^k - 1 = n$, so any k points whose cyclic gap sequence is a permutation of those numbers wrap the circle exactly once. Such a k-subset of $\mathbb{Z}_n$ is called a *permuted copy*, or here simply a *copy*, and $\mathcal{H}_k$ denotes the set of all of them: a k-uniform hypergraph on $\mathbb{Z}_n$.

Conjecture 1.1 (Stromquist [2]; [1, Conjecture 1.1])**.** For every $k \ge 3$, every two-colouring of $\mathbb{Z}_n$ contains a monochromatic copy.

The conjecture is a Ramsey-like statement about dividing a circle, and it has resisted proof. The case $k = 3$, where $\mathcal{H}_3$ consists of the seven Fano lines $\{0, 1, 3\} + c$ together with their mirror images $\{0, 1, 5\} + c$, is attributed by [1] to Bialostocki and Nielsen [3]. Damásdi, Frankl, Pach and Pálvölgyi [1, §6] record that "Stromquist confirmed his conjecture by a computer search for $k \le 6$. Using a SAT solver, we confirmed it for $k \le 7$"; they add that they "also ran the $k = 8$ case for approximately 100 days, but the solver did not come to a conclusion". So the conjecture is open from $k = 8$ on, and the obstruction is not conceptual but computational: at $k = 8$ the search space has 2^{255} colourings and the enumeration of copies has 1 285 200 entries.

The same paper raises a second question, in one sentence, and does not answer it:

> "We mention that if Conjecture 1.1 holds, then it would be also interesting to determine the minimal number of monochromatic permuted copies in a two-colouring of a regular $(2^k - 1)$-gon." [1, §6]

Write

$$\mu(k) \;=\; \min_{\chi\colon \mathbb{Z}_n \to \{0,1\}} \#\{\, Q \in \mathcal{H}_k : \chi \text{ is constant on } Q \,\}$$

for that minimum. Conjecture 1.1 is exactly the assertion $\mu(k) \geq 1$; the authors of [1] prove in addition that for every $k \geq 3$ the number of monochromatic copies of any colouring is even ([1, Proposition 6.2], stated there for an arbitrary side-length tuple), which upgrades the conjecture to $\mu(k) \geq 2$ and says that μ takes only even values. The same parity fact, and three values of μ, are announced in Stromquist's one-page problem-session note [2], which calls a monochromatic copy a *winner*: "If $k \geq 3$ then for any coloring, the number of winners is even. The minimum number of winners (over all colorings) is 2 (for 7 points), 4 (for 15 points), and 10 (for 31 points)." That is, $\mu(3) = 2$, $\mu(4) = 4$ and $\mu(5) = 10$, stated with no proof, no method and no further value. Those two sources appear to be everything that is on record about μ. This paper is about μ: an upper bound valid at every k, proofs of the three announced values, explicit colourings at $k = 6, 7$, and the structural restatement that organises all of it.

Results. *A bound at every k* (Section 4). Let $\chi(x) = x \bmod 2$, computed on the representatives $x \in \{0, 1, \ldots, n-1\}$. Then χ has at most $(k-1)!$ monochromatic copies, so

$$\mu(k) \;\leq\; (k-1)! \qquad \text{for every } k \geq 2.$$

The proof is two elementary steps and no computation. In particular $\mu(8) \leq 5040$, at the case that ran for a hundred days without a verdict; we know of no earlier bound on μ there. At $k = 3, 4, 5, 6$ we verify that the parity colouring has *exactly* $(k-1)!$ monochromatic copies, namely $2, 6, 24, 120$.

Exact values (Section 6). We prove

$$\mu(3) = 2, \qquad \mu(4) = 4, \qquad \mu(5) = 10,$$

the three values [2] announces. We know of no published proof of any of them. The first two follow from an exhaustive sweep of all 2^7 and all 2^{15} colourings. The third does not: 2^{31} colourings tested naively against 744 copies is out of reach of a proof assistant's evaluator. We prove instead the soundness of a branch-and-bound whose prune rests on a single observation — a copy all of whose points lie below the current prefix has its monochromaticity already decided — and run it, over a colour-swap-halved tree, to completion. As a by-product Conjecture 1.1 is reproved at $k = 3, 4, 5$, quantitatively: at $k = 5$ every colouring of the regular 31-gon has at least ten monochromatic copies.

Upper bounds beyond (Section 5). Explicit colourings, found by local search and then counted in the kernel, give $\mu(6) \leq 20$ and $\mu(7) \leq 120$. Neither is known to be optimal; Section 7 says exactly what stands in the way.

The lattice picture (Section 3). With $\sigma(S) = \sum_{i \in S} 2^i$ for $S \subseteq [k] = \{0, 1, \ldots, k-1\}$, a subset $Q \subseteq \mathbb{Z}_n$ is a copy if and only if

$$Q \;=\; \{\, p + \sigma(C_j) : 0 \leq j < k \,\}$$

for some $p \in \mathbb{Z}_n$ and some maximal chain $\emptyset = C_0 \subset C_1 \subset \cdots \subset C_{k-1} \subseteq [k]$ of the Boolean lattice, truncated to its bottom k levels. This is an exact restatement, in both directions and at every k; it makes the three symmetries of $\mathcal{H}_k$ definitional rather than computational, and it turns "χ has no monochromatic copy" into "for every p, both colour classes of the induced colouring of $2^{[k]}$ are cutsets". The source uses a different reformulation — an $O(n^2 k)$ reachability dynamic program — and does not mention this one.

The state of the table is then:

k	n	$\lvert\mathcal{H}_k\rvert$	$\mu(k)$
2	3	3	1
3	7	14	2
4	15	90	4
5	31	744	10
6	63	7 560	≤ 20
7	127	91 440	≤ 120
8	255	1 285 200	≤ 5040

For $3 \le k \le 6$ the third column is the cardinality of $\mathcal{H}_k$ as verified in Section 2; for $k = 2, 7, 8$ it is the length $n\,(k-1)!$ of the anchored enumeration, which equals $|\mathcal{H}_k|$ as soon as that enumeration is duplicate-free. The entry $\mu(2) = 1$ is a degenerate check rather than a result: at $k = 2$ the copies are the three adjacent pairs of $\mathbb{Z}_3$ and the pigeonhole principle gives the value. Note that it is odd; consistently, both statements of the parity fact quoted above are made only for $k \ge 3$.

On novelty. The arguments here are elementary and the claim is not that they are difficult. Against the record described above — the parity proposition of [1] and the three announced values of [2] — what is offered here is: the bound $\mu(k) \le (k-1)!$ at every k, and with it $\mu(8) \le 5040$; the upper bounds $\mu(6) \le 20$ and $\mu(7) \le 120$; the chain restatement of Section 3; and proofs, machine-checked ones, of the three values that [2] announces without any.

The searches behind "the record" were run on 2026-08-29. The journal version of [1] had appeared nine days earlier and had no citations recorded by Crossref, OpenAlex or Semantic Scholar. A full-text search of a local mirror of arXiv (881 145 papers, newest identifier `2608.27451`) for `Stromquist`, `dividing a circle` and `permuted cop` returned 432 papers, of which [1] is the only one on this problem; the rest are other work of Stromquist's — cake-cutting, four-colour reducibility, Banach–Mazur distance, toroidal colouring — or coincidences of the other two search phrases. OEIS searches returned ten sequences for $2, 4, 10, 20$, none of them combinatorially related to this problem, and nothing at all for $2, 4, 10, 20, 120$.

One negative is worth recording as a caution. An earlier round of exactly these searches missed [2] completely. The arXiv version of [1] cites it as a talk at "The VI International AMMCS Interdisciplinary Conference Waterloo, Ontario, Canada, 2023", with no link, and under that description we could not locate it; the journal version cites the same item as the problem session of Permutation Patterns 2023 and prints a URL, and the note at that URL is where the three values are. A negative literature search is worth no more than the bibliographic data it is run against.

Verification. Every theorem, proposition and lemma below, and every numerical value quoted from a finite search, has been checked by the Lean 4 proof assistant against *Mathlib*. Section 8 says precisely which parts rest on exhaustive evaluation, what was searched, and which remarks are outside the formal development.

2. Copies, the anchored enumeration, and μ

Throughout, $k \ge 2$ and $n = 2^k - 1$. We write $[k] = \{0, 1, \dots, k-1\}$ for the exponent set and identify $\mathbb{Z}_n$ with $\{0, 1, \dots, n-1\}$ when representatives matter.

Definition 2.1. Let $a = (a_0, a_1, \dots, a_{k-1})$ be a permutation of $[k]$, let $p \in \mathbb{Z}_n$, and put $s_j = \sum_{i<j} 2^{a_i}$ for $0 \le j \le k$. The *copy* determined by (p, a) is

$$Q(p, a) \;=\; \{\, p + s_j \bmod n \;:\; 0 \le j < k \,\} \subseteq \mathbb{Z}_n.$$

Since $0 = s_0 < s_1 < \cdots < s_{k-1} < s_k = n$, the k residues $p + s_j$ are distinct, and the cyclic gaps of $Q(p, a)$ are $2^{a_0}, 2^{a_1}, \dots, 2^{a_{k-1}}$ in that cyclic order. We write $\mathcal{H}_k = \{Q(p, a)\}$ for the set of all copies, over all n start points and all $k!$ permutations.

The same set is produced by many pairs (p, a), and the redundancy is exactly a rotation.

Lemma 2.2 (Rotation)**.** *Let a be a permutation of $[k]$, let $p \in \mathbb{Z}_n$, and let $0 \le r \le k$. Write $a^{(r)}$ for the cyclic rotation of a by r places. Then*

$$Q\big(p + s_r,\; a^{(r)}\big) \;=\; Q(p, a).$$

Proof sketch. For $0 \le j < k$ one has $s_r + \sum_{i<j} 2^{a^{(r)}_i} \equiv s_{(r+j) \bmod k} \pmod n$: if $r + j < k$ the two sides are equal as integers, and if $r + j \ge k$ the left side exceeds the right by $s_k = n \equiv 0$. Since $j \mapsto (r + j) \bmod k$ permutes $[k]$, the two images coincide. □

Consequently every copy can be written with the gap $1 = 2^0$ used first. Call the pair (p, a) *anchored* when $a_0 = 0$; the anchored pairs number $n\,(k-1)!$, and they still produce every copy, because rotating an arbitrary a until its entry 0 comes first is Lemma 2.2. An anchored copy contains the two adjacent points p and $p+1$; we call $\{p, p+1\}$ its *gap*-1 *edge*, and it is this edge that carries the argument of Section 4.

Proposition 2.3 (The copy count)**.** *The anchored enumeration has exactly $n\,(k-1)!$ entries for every k. It is duplicate-free for $3 \le k \le 6$, so that*

$$|\mathcal{H}_3| = 14, \quad |\mathcal{H}_4| = 90, \quad |\mathcal{H}_5| = 744, \quad |\mathcal{H}_6| = 7\,560,$$

and the enumeration has length 91 440 at $k = 7$ and 1 285 200 at $k = 8$.

Proof sketch. The length is Lemma 2.2 plus the count of anchored pairs, and is proved for all k with no computation. Duplicate-freeness is a finite check, performed by exhaustive evaluation for each k with $3 \le k \le 6$; at $k = 6$ the $O(m^2)$ comparison is run on the integer codes $\sigma(Q)$ of the 7 560 entries rather than on the sets themselves. The cardinality statement follows: distinct entries of a duplicate-free enumeration are distinct copies. □

The count $n\,(k-1)!$ is not ours: it is stated in [1], both in the parity proposition and in the description of their $O(n^2 k)$ checking algorithm, where "there are at most $n \cdot (k-1)!$ possible k-gons that have to be checked". What Proposition 2.3 adds is the duplicate-freeness of the anchored enumeration, which is what licenses replacing a set cardinality by a list length in every count below.

The SAT instances [1] publish corroborate both, outside the formal development. Their DIMACS files for $k = 4, 5, 6, 7$ declare 180, 1 488, 15 120 and 182 880 clauses, that is $2n\,(k-1)!$ — one clause per copy per colour. We downloaded them on 2026-08-29 and checked that the $n\,(k-1)!$ positive clauses of each file are pairwise distinct and are, as vertex sets, exactly the copies enumerated here. That is an independent confirmation of duplicate-freeness at $k = 7$ as well, where our own check does not reach; it is a computation on somebody else's file and is not part of what the kernel verifies.

Definition 2.4. A colouring is a map $\chi\colon \mathbb{Z}_n \to \{0, 1\}$; a copy Q is *monochromatic* for χ when χ is constant on Q. Write $\mathrm{m}_k(\chi)$ for the number of monochromatic copies and

$$\mu(k) \;=\; \min_{\chi}\; \mathrm{m}_k(\chi).$$

Two remarks on how μ is handled in a formal setting, since both matter for reading the proofs below. First, the minimum is taken over *all* colourings, not over a chosen family: in the formalisation χ ranges over all Boolean-valued functions on the natural numbers, and the reduction to the 2^n bit patterns is a lemma, namely that every colouring agrees on $\{0, \dots, n-1\}$ with the bit pattern of some $c < 2^n$, together with the observation that m_k reads only those positions. Second, an upper bound never needs duplicate-freeness: counting monochromatic entries of the anchored enumeration overcounts monochromatic copies if the enumeration repeats itself, so a bound obtained that way is valid regardless.

3. The Boolean-lattice restatement

For $S \subseteq [k]$ put $\sigma(S) = \sum_{i \in S} 2^i$. Then σ maps the 2^k subsets of $[k]$ onto $\mathbb{Z}_n$, injectively as integers, and $\sigma(\emptyset) = 0$, $\sigma([k]) = n \equiv 0$. By a *truncated maximal chain* we mean a family $C_0 \subset C_1 \subset \cdots \subset C_{k-1}$ of subsets of $[k]$ with $|C_j| = j$; equivalently, the bottom k levels of a maximal chain of the Boolean lattice $2^{[k]}$.

Theorem 3.1 (Chain restatement)**.** *Let $k \ge 1$ and $Q \subseteq \mathbb{Z}_n$. Then $Q \in \mathcal{H}_k$ if and only if there are $p \in \mathbb{Z}_n$ and a truncated maximal chain $C_0 \subset \cdots \subset C_{k-1}$ of $2^{[k]}$ with*

$$Q \;=\; \{\, p + \sigma(C_j) \bmod n \;:\; 0 \le j < k \,\}.$$

Proof sketch. Forward: for a permutation a of $[k]$, the prefix sets $C_j = \{a_0, \dots, a_{j-1}\}$ form a truncated maximal chain, and $s_j = \sigma(C_j)$ because the a_i are distinct. Backward, which is the half with content: from a chain, the set $C_{j+1} \setminus C_j$ is a singleton by the cardinality condition, so reading off its element for $j = 0, \dots, k-1$ produces a sequence a; it is repetition-free because the chain is strictly increasing, hence a permutation of $[k]$, and its prefix sets are the C_j by induction. No computation enters; the statement and its proof are uniform in k. □

The point of Theorem 3.1 is that everything structural about $\mathcal{H}_k$ becomes visible in it.

Remark 3.2 (The three symmetries). Extend a truncated chain by $C_k = [k]$, which changes nothing since $\sigma(C_k) = n \equiv 0 = \sigma(C_0)$. Then: *translation* $x \mapsto x + c$ is $p \mapsto p + c$; the *doubling map* $x \mapsto 2x$ is the cyclic shift $i \mapsto i + 1 \bmod k$ of the ground set $[k]$, because $2 \cdot 2^{k-1} = 2^k \equiv 1$, and a cyclic shift carries chains to chains; and *negation* $x \mapsto -x$ is complementation $S \mapsto S^c$, because $\sigma(S) + \sigma(S^c) = n \equiv 0$, which reverses a maximal chain (the start point becoming $-p$). So the group generated by the translations, the doubling map and negation, of order dividing $2kn$, acts on $\mathcal{H}_k$. This remark is elementary and is stated here for orientation; it is not among the machine-checked statements of this paper.

Remark 3.3 (Cutsets, and why a single level of the picture says nothing). For a colouring χ and $p \in \mathbb{Z}_n$ define $\chi_p \colon 2^{[k]} \to \{0,1\}$ by $\chi_p(S) = \chi(p + \sigma(S))$. By Theorem 3.1, χ has no monochromatic copy if and only if, for every p, no truncated maximal chain of $2^{[k]}$ is monochromatic for χ_p — equivalently, both colour classes of χ_p are *cutsets*, families meeting every maximal chain. Stromquist's conjecture is thus a statement about n coupled two-colourings of the Boolean lattice.

The coupling is the whole content: for a *single* p the condition is vacuous. Colour the singletons of $2^{[k]}$ with one colour and every other subset with the other; every truncated maximal chain contains exactly one singleton and at least one other set, so no chain is monochromatic. Hence no bound on μ can be obtained from one p at a time. (Counting the same copies through all n values of p gives $k \cdot \mathrm{m}_k(\chi) = \sum_p M(\chi_p)$, with $M(\varphi)$ the number of monochromatic truncated maximal chains, provided the anchored enumeration is duplicate-free — which Proposition 2.3 establishes for $k \le 6$.) This remark, too, is outside the formal development.

4. An upper bound at every k

Theorem 4.1 (The parity colouring). *Let $k \ge 2$ and let $\chi(x) = x \bmod 2$, computed on the representatives $x \in \{0, 1, \dots, n-1\}$. Then χ has at most $(k-1)!$ monochromatic copies. Consequently*

$$\mu(k) \;\le\; (k-1)! \qquad (k \ge 2).$$

Proof. By Lemma 2.2 every monochromatic copy is $Q(p, a)$ for an anchored pair, and such a copy contains its start p and the next point $p + 1 \bmod n$. If $p < n - 1$ then p and $p + 1$ are consecutive integers below n, so $\chi(p) \ne \chi(p+1)$ and the copy is not monochromatic. Hence every monochromatic copy is anchored at $p = n - 1$. There are $(k-1)!$ anchored pairs with that start point, so at most $(k-1)!$ copies among them, and $\mu(k) \le \mathrm{m}_k(\chi) \le (k-1)!$. □

The bound is uniform in k and its proof uses no computation; in particular it applies where search does not.

Corollary 4.2. $\mu(8) \le 5040$.

We know of no earlier bound on $\mu(8)$; that is the case whose SAT instance ran for approximately one hundred days without a verdict [1]. Note that the bound says nothing about Conjecture 1.1 itself, which asserts a *lower* bound.

Proposition 4.3 (The witness is counted exactly, for small k). *The parity colouring has exactly $2, 6, 24, 120$ monochromatic copies at $k = 3, 4, 5, 6$ respectively, that is, exactly $(k-1)!$.*

Proof sketch. Four finite evaluations: the monochromatic entries of the anchored enumeration are counted for each k, the enumerations having $14, 90, 744, 7560$ entries, and Proposition 2.3 identifies those counts with counts of copies. □

Remark 4.4. An elementary argument gives exactness at every k. A copy anchored at $n-1$ is $\{n-1\} \cup \{\sigma(S)\}$, with S running over the prefix sets of an ordering of the remaining gap exponents $1, \ldots, k-1$; every such $\sigma(S)$ is an even integer at most $2^k - 2 = n-1$, and $n-1$ is even as well, so all $(k-1)!$ of these copies are in fact monochromatic, and distinct orderings give distinct point sets because σ is injective on subsets. We have not formalised this, and Theorem 4.1 does not need it: an upper bound needs only that the $(k-1)!$ anchored pairs at $n-1$ cover the monochromatic copies, not that they do so injectively. Proposition 4.3 is the machine-checked shadow of the general statement.

Remark 4.5 (Defects). Call $v \in \mathbb{Z}_n$ a *defect* of χ when $\chi(v) = \chi(v+1)$, and let $D(\chi)$ be the number of defects; $D(\chi) \geq 1$ always, since n is odd. The proof of Theorem 4.1 used only that the parity colouring has the single defect $n-1$, and it gives in general $\mathrm{m}_k(\chi) \leq (k-1)!\, D(\chi)$. This is far from tight at the extremal colourings: the best colouring we know at $k = 6$ has $D = 21$ and 20 monochromatic copies, against a bound of 2520. Colourings that minimise μ therefore sit at the opposite end of the defect scale from the clean construction, which is the tension any determination of μ has to resolve. Neither the inequality nor the measurement is part of the formal development.

5. Upper bounds from explicit colourings

Theorem 5.1. *$\mu(5) \leq 10$, $\mu(6) \leq 20$ and $\mu(7) \leq 120$.*

Proof sketch. Three explicit colourings, written below as bit strings with position 0 first and 1 meaning red. For each, the monochromatic entries of the anchored enumeration — 744, 7 560 and 91 440 of them — are counted by exhaustive evaluation, giving 10, 20 and 120. As noted after Definition 2.4, that count dominates the number of monochromatic copies whatever duplicates the enumeration may contain, so each bound is valid without any appeal to Proposition 2.3. □

$k = 5$, 10 copies:	`0001011010000001111010100101111`
$k = 6$, 20 copies:	`00100110100101100100101100110100101110100101101110100101110110101`
$k = 7$, 120 copies:	`10110000111100100110000111101001111001011110000110010011110000110100101100001111001011100011111010111110001110100111100001101000`

(The $k = 7$ string has 127 characters and is broken across two lines.) The colourings were produced by a min-conflicts local search — repeatedly pick a monochromatic copy at random and flip the vertex of it that helps most, with 12% noise — calibrated by the fact that it returns the proved optima $2, 4, 10$ at $k = 3, 4, 5$ within 100 000 iterations. At $k = 5$ and $k = 6$ the values 10 and 20 are the best found by every search we ran; at $k = 5$ the value is optimal by Theorem 6.3. At $k = 7$ we made fourteen long runs, eight of 5 million moves and six of 12 million; nine of them returned exactly 120, the other five returned 122 or 126, and none went below 120. Two distinct colourings attaining 120 were each found twice from different random starts. The search itself is not on the trust path: only the count of the final colouring is, and that count is exhaustive.

6. Exact values

Theorem 6.1. *$\mu(3) = 2$ and $\mu(4) = 4$.*

Proof sketch. Upper bounds: $\mu(3) \leq 2! = 2$ is Theorem 4.1, and $\mu(4) \leq 4$ is the colouring `110110110100000` counted against the 90 copies. Lower bounds: a full sweep. Every colouring agrees on $\{0, \ldots, n-1\}$ with the bit pattern of some $c < 2^n$ and m_k depends only on those positions, so it suffices to test the $2^7 = 128$ and $2^{15} = 32\,768$ bit patterns, each against the 14 resp. 90 entries of the anchored enumeration, and to check that no pattern falls below the claimed value. The $k = 4$ sweep is 2.95 million copy tests. Duplicate-freeness (Proposition 2.3) is what makes

the enumeration count equal to the copy count, and is needed here because the claim is a lower bound. □

At $k = 5$ the same method would need 2^{31} patterns against 744 copies, some 1.6×10^{12} tests, which no evaluator inside a proof assistant will do. We replace the sweep by a search whose pruning rule is itself proved.

Lemma 6.2 (Soundness of the prune)**.** *Fix k and a prefix length t. Assign the positions $0, 1, \ldots, n-1$ in this order, and for an assignment of the positions below t let N be the number of copies that are monochromatic and contained in $\{0, \ldots, t-1\}$. Then every completion of that prefix has at least N monochromatic copies.*

Proof. A copy contained in $\{0, \ldots, t-1\}$ has all its colours fixed by the prefix, so its monochromaticity is decided; monochromatic copies of the prefix stay monochromatic in every completion, and the remaining copies only add. □

Theorem 6.3. $\mu(5) = 10$.

Proof sketch. The upper bound is Theorem 5.1. For the lower bound, run the depth-first search over prefixes of the 31 positions of $\mathbb{Z}_{31}$, maintaining the count N of Lemma 6.2 and cutting a branch as soon as $N \geq 10$; if every branch is cut or reaches depth 31 with $N \geq 10$, then every colouring has at least ten monochromatic copies. Two reductions make the search feasible. *The colour swap*: replacing χ by its complement preserves the monochromatic copies, so only the colourings with $\chi(0) = 0$ need be searched, halving the tree. *Bit masks*: a copy Q is represented by the integer $\sigma(Q)$ of Section 3, and monochromaticity under the pattern c is the two-operation test $c \,\&\, \sigma(Q) \in \{0, \sigma(Q)\}$; the equivalence with the set-theoretic definition rests on $\mathrm{bit}_x(\sigma(Q)) = [x \in Q]$. The search visits 45 938 422 nodes and performs 2.54×10^9 copy tests; it takes 6.0 seconds as a C program, and the file containing it checks in 22 minutes on a loaded machine, of which some 19 minutes is the sweep itself, extrapolated from shorter runs of the same procedure. Its verdict — that no colouring of $\mathbb{Z}_{31}$ has nine or fewer monochromatic copies — is the one exhaustive computation on which Theorem 6.3 rests. It is corroborated by an independent computation: a SAT instance asserting "at most nine monochromatic copies" at $k = 5$ is unsatisfiable in 32 seconds. □

The search is also checked against the brute-force sweep where both are affordable: at $k = 3$ and $k = 4$ it returns *true* at the true value of μ and *false* one step above it, agreeing with Theorem 6.1, which is two independent decision procedures for the same numbers (Proposition 8.1).

Corollary 6.4. *Conjecture 1.1 holds for $k = 3, 4, 5$: every two-colouring of the regular 7-, 15- and 31-gon contains a monochromatic permuted copy. Quantitatively, it contains at least 2, 4 and 10 of them.*

Proof. Immediate from $\mu(k) \geq 1$, which Theorems 6.1 and 6.3 give with room to spare. □

These cases are not new — $k = 3$ is [3], $k \leq 6$ is Stromquist's own computer search, and $k \leq 7$ is [1] — but the route is different, and each of the three statements is strictly stronger than the conjecture there, because it comes out of an exact value of μ. Of the three, $k = 5$ is the smallest that is settled neither by [3] nor by the remark in [1] that "it is not hard to verify the case $k = 4$ either". Stromquist's note likewise calls $k = 3$ and $k = 4$ "easy to prove" and records that "computer proofs exist when $k \leq 7$" [2].

7. What is not settled

Is $\mu(6) = 20$? The upper bound is Theorem 5.1. For the matching lower bound we encoded "at most 18 monochromatic copies at $k = 6$" as a SAT instance — the published clauses of [1] plus one indicator variable per copy, a cardinality constraint, and a lex-leader symmetry break — in five different ways, the largest of them that we could run having 143 442 variables and 365 430 clauses,

and no encoding produced a verdict within 122 CPU-minutes in total. (A sixth, a full totalizer over the 7 560 indicator variables, reached 28.7 million clauses and was abandoned on memory grounds before solving.) The instance is near-threshold, which is the hard regime for unsatisfiability. The branch-and-bound of Theorem 6.3 does not transfer: at $k = 6$ the smallest span of a copy is $n - 2^{k-1} = 31$, against 15 at $k = 5$, so no copy is decided by a prefix of fewer than 32 positions and the search must visit of order 2^{31} nodes before the prune of Lemma 6.2 can fire even once. This is a structural obstruction, not a matter of speed.

Is $\mu(7) = 120$*?* We do not know; the value is an upper bound backed by search.

Remark 7.1 (A candidate formula, refuted in practice). The values $\mu(3), \mu(4), \mu(5) = 2, 4, 10$ and the bound $\mu(6) \le 20$ are matched by $2\lfloor 2^{k-1}/3 \rfloor = 2, 4, 10, 20$, which would predict $\mu(7) = 42$ and $\mu(k)/n \to 1/3$. (The inner sequence is A000975 [6] up to an index shift: that entry has offset 0 and $a(m) = \lfloor 2^{m+1}/3 \rfloor$, so $\lfloor 2^{k-1}/3 \rfloor = a(k-2)$.) The prediction is wrong in practice: 68 independent local searches at $k = 7$ — 54 runs of the min-conflicts search of Section 5, of lengths from 400 000 to 12 million moves, and 14 runs of three other search designs — never returned a colouring with fewer than 120 monochromatic copies, and $120/127 = 0.94$ is nowhere near $1/3$. **This is a measurement, not a theorem**: no lower bound on $\mu(7)$ is proved here, and refuting $\mu(7) \le 42$ rigorously would need an unsatisfiability proof for a cardinality constraint over 91 440 indicator variables, which we did not attempt. The remark is recorded outside the formal development and nothing else in this paper depends on it.

Stromquist's conjecture at $k = 7$ *and* $k = 8$. These stand outside what this paper verifies, and deliberately so. The $k \le 7$ cases are due to [1] by SAT: they report that "the SAT solver Glucose3 solved the $k = 7$ case in approximately one hour" [1, §6]. We reproduced $k = 7$ as an unsatisfiability result in 561 seconds of CDCL solving; that figure is our own measurement on a shared machine carrying a load of about 25, so it is an upper estimate and not a benchmark, and it is not comparable with theirs. The run used a lex-leader symmetry break over the group generated by the translations, the doubling map, negation and the colour swap (of order $4kn$ for the $k \le 8$ we checked; only the subgroup property, Remark 3.2, is needed for soundness). We report that as an external computation and not as a theorem, because there is no route into a proof assistant's kernel for it: the claim is the unsatisfiability of a search over 2^{127} colourings, so there is neither a small enough space to evaluate exhaustively nor a short certificate to check — the CDCL refutation has of order 10^6 conflicts, and the corresponding clausal proof is of estimated size 10^8–10^9 bytes, far beyond what a kernel can replay. The same obstruction, an order of magnitude worse, applies at $k = 8$, where [1] report about one hundred days of solving without a verdict. Everything in Sections 2–6 was chosen to stay on the side of this line where kernel-checked statements exist.

8. Verification

Every statement of this paper that is numbered as a theorem, proposition or lemma has been formally verified in Lean 4 [4] (toolchain `v4.32.0`) against *Mathlib* [5]; the development contains no `sorry` and introduces no axiom of its own. The repository is private at the time of writing: repository reference to be added at submission. The remarks are explicitly excluded: Remarks 3.2, 3.3, 4.4, 4.5 and 7.1, and the discussion in Section 7, are elementary observations or measurements stated outside the formal development, and are labelled as such where they appear.

The parts that rest on no computation at all — axioms `propext`, `Classical.choice` and `Quot.sound` only — are Lemma 2.2 and the anchoring statement derived from it, Theorem 3.1 in both directions, Theorem 4.1 with Corollary 4.2, the reduction of an arbitrary colouring to a bit pattern, and the soundness of the branch-and-bound (Lemma 6.2, together with the colour swap and the mask reflection used in Theorem 6.3).

What is delegated to compiled evaluation, through Lean's `native_decide`, is exactly the finite searches; each such statement additionally depends on the axiom asserting that the compiler's evaluation of a closed Boolean expression agrees with its meaning in the kernel. The searches are: the

duplicate-freeness of the anchored enumeration at each k with $3 \le k \le 6$, and the resulting cardinalities of Proposition 2.3; the four counts of Proposition 4.3; the count of each explicit colouring in Theorems 5.1 and 6.1 against the 90, 744, 7 560 and 91 440 entries of the corresponding enumeration; the sweep over all 2^7 and all 2^{15} bit patterns for the lower bounds of Theorem 6.1; and the branch-and-bound of Theorem 6.3, one Boolean whose evaluation visits 45 938 422 nodes and performs 2.54×10^9 copy tests over the colourings of $\mathbb{Z}_{31}$ with $\chi(0) = 0$. Nothing else is evaluated.

Every number quoted was produced first by an independent implementation in C, written against the definitions of [1] rather than against ours, and the formal proof was written afterwards: the copy counts $14, 90, 744, 7\,560, 91\,440$, the values $\mu(3) = 2$, $\mu(4) = 4$ and $\mu(5) \ge 10$, and the counts of all four explicit colourings were reproduced that way, and the $k = 5$ lower bound was corroborated a third time by a SAT solver. The following controls are also part of the development, and are stated because a family of theorems about a wrongly defined object can be true and empty.

Proposition 8.1 (Controls)**.** *$\mu(2) = 1$, so the bound of Theorem 4.1 is attained and cannot be improved uniformly. Neither $\mu(4) \le 3$ nor $\mu(4) \ge 5$ holds, so the exhaustive sweep is two-sided. $\mu(4) < 3!$, so Theorem 4.1 is not an equality. $|\mathcal{H}_5| \ne 31 \cdot 5!$, so anchoring really removes the k-fold overcount. At $k = 3$ the parity colouring leaves* 12 *of the* 14 *copies non-monochromatic, so monochromaticity is not vacuous. At $k = 4$ the unanchored enumeration over all $n \cdot k!$ pairs and the anchored one over $n \cdot (k-1)!$ pairs give literally the same set of copies. And the branch-and-bound returns* true *at $\mu(k)$ and* false *at $\mu(k) + 1$ for $k = 3, 4$, agreeing with the brute-force sweep.*

References

[1] G. Damásdi, N. Frankl, J. Pach and D. Pálvölgyi, *Monochromatic configurations on a circle*, Combinatorial Theory **6** (2026), no. 2, #9; doi:10.5070/C6.68671; arXiv:2504.10687. Quotations and numbered references here are to the published version; Conjecture 1.1, Lemma 6.1 and Proposition 6.2 carry the same numbers in arXiv v2, whose wording differs only in small points of grammar and notation.

[2] W. Stromquist, *A Ramsey-like conjecture for dividing a circle (problem session)*, Permutation Patterns 2023, University of Burgundy, Dijon, 3–7 July 2023. One page, https://2023.permutationpatterns.com/p/problem_session/stromquist.pdf.

[3] A. Bialostocki and M. J. Nielsen, *Minimum sets forcing monochromatic triangles*, Ars Combin. **81** (2006), 297–303.

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

[5] The mathlib Community, *The Lean mathematical library*, in: Certified Programs and Proofs (CPP 2020), ACM, 2020, pp. 367–381; arXiv:1910.09336.

[6] OEIS Foundation Inc., *The On-Line Encyclopedia of Integer Sequences*, https://oeis.org; entry A000975, cited in Remark 7.1.