THE UPPER ORIENTED GENERAL POSITION NUMBER: FIRST VALUES AND EXTREMAL GRAPHS

KOREA SUPERINTELLIGENCE LABS

Abstract. For a graph G of order n, the *upper oriented general position number* $\mathrm{ugp}(G)$ is the largest general position number of an orientation of G. Chandran S. V., Di Stefano, Erskine, Haritha S., Thomas and Tuite proved that $\mathrm{ugp}(G)$ equals the largest order of an induced subgraph of G that is a comparability graph, remarked that this quantity "has not been studied before", and published no value of it beyond complete graphs, paths and cycles. We compute the extremal function $\mu(n) = \min\{\mathrm{ugp}(G) : |V(G)| = n\}$ for every $n \leq 15$: it is $2, 3, 4, 4, 5, 5, 6, 6, 6, 7, 7, 7, 7, 7$, and $\mu(16) \in \{7, 8\}$. We classify the extremal graphs at order 10: there are exactly two, one of them the line graph of K_5, so that 10 is the largest order with $\mu(n) = 6$. Every minimiser we exhibit — at orders $6, 8, 9, 10, 12, 15$ — is the line graph of a complete or a complete bipartite graph, and we show by example that this pattern does not extend to all such line graphs. We also give exact values for 43 named graphs, among them $\mathrm{ugp}(\text{Petersen}) = 7$, $\mathrm{ugp}(K(6, 2)) = 9$, $\mathrm{ugp}(K(7, 2)) = 12$, and $\mathrm{ugp}(\overline{C_n}) = n - 1$ for every $5 \leq n \leq 16$. All statements are machine-checked in Lean 4.

1. Introduction

A set S of vertices of a digraph D is in *general position* if for no two $u, v \in S$ does a shortest u, v-path pass through a third vertex of S; the *general position number* $\mathrm{gp}(D)$ is the largest order of such a set. For an undirected graph G of order $n \geq 2$, Chandran S. V., Di Stefano, Erskine, Haritha S., Thomas and Tuite [1] define the *upper* and *lower oriented general position numbers*

$$\mathrm{ugp}(G) = \max_D \mathrm{gp}(D), \qquad \mathrm{lgp}(G) = \min_D \mathrm{gp}(D),$$

the extrema being taken over all orientations D of G. The upper parameter has a purely undirected description. Recall that a graph is a *comparability graph* if its edges admit a *transitive orientation*: an orientation in which $x \to y$ and $y \to z$ force the arc $x \to z$.

Theorem 1.1 ([1, Theorem 4.4])**.** *The upper oriented general position number of G is equal to the largest order of an induced subgraph H of G that is also a comparability graph.*

The authors note explicitly that the induced comparability subgraphs in question may be disconnected — a point that is not decorative; the value $\mu(6) = 5$ below depends on it. From Theorem 1.1 they deduce that $\mathrm{ugp}(G) = n$ if and only if G is itself a comparability graph [1, Corollary 4.5], and then write:

> It appears that the order of a largest induced comparability subgraph has not been studied before.

That sentence, and the fact that [1] publishes no value of the quantity beyond three easy families, is the starting point of this paper. What [1] does prove about the size of ugp is three bounds, which we use throughout as sanity anchors: $\mathrm{ugp}(G) \geq \mathrm{diam}(G)+1$, with equality exactly for paths, whence $\mathrm{ugp}(G) \geq 4$ for every graph of order $n \geq 4$ [1, Corollary 4.7]; $\mathrm{ugp}(G) \geq \max\{\alpha^\omega(G), \alpha_2(G)\}$, where α^ω is the largest order of an induced disjoint union of cliques and α_2 the largest order of an induced bipartite subgraph [1, Corollary 4.6]; and $\mathrm{ugp}(C_n) = n$ for even n, with $\mathrm{ugp}(C_n) \leq n - 1$ for odd n. A further remark of [1] is that a comparability graph contains no induced odd cycle of length at least 5, so that a graph with r disjoint induced odd cycles of length ≥ 5 has $\mathrm{ugp}(G) \leq n - r$. These bounds are usually far from the truth: on the Kneser graph $K(7, 2)$ the last one gives $\mathrm{ugp} \leq 17$, against the correct value 12 (Theorem 4.4).

The extremal function. Theorem 1.1 turns ugp into a Ramsey-type question for a hereditary class. Write

$$\mu(n) \;=\; \min\{\mathrm{ugp}(G) \;:\; G \text{ a graph of order } n\},$$

the least order of a largest induced comparability subgraph over all graphs on n vertices. Equivalently, $\mu(n)$ is the largest m such that *every* graph on n vertices contains an induced comparability subgraph of order m. The source proves $\mu(n) \geq 4$ for $n \geq 4$ and, by a Ramsey argument together with the α^{ω}/α_2 bound, that $\mu(n) \to \infty$; it evaluates μ nowhere. Section 3 computes it:

n	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
$\mu(n)$	2	3	4	4	5	5	6	6	6	7	7	7	7	7	$\{7, 8\}$

Two features of the table are results in their own right. First, the flat stretches: μ is constant on $\{4, 5\}$, on $\{6, 7\}$, on $\{8, 9, 10\}$ and on $\{11, \ldots, 15\}$. Since adding an isolated vertex raises ugp by exactly one, $\mu(n) \leq \mu(n+1) \leq \mu(n) + 1$ always (Proposition 2.4), so the content of each entry is whether the step happens; the answer is "no" more and more often. Second, the end of a flat stretch is an extremal question of its own. We answer the one for the value 6 completely: 10 is the largest order N with $\mu(N) = 6$, and exactly two graphs of order 10 realise it, one of them the line graph of K_5 (Theorems 3.9 and 3.10).

We warn against fitting a formula. The value $\lfloor n/2 \rfloor + 2$ agrees with $\mu(n)$ for $n = 4, \ldots, 8$; it is wrong at $n = 10$, where it predicts 7 and the value is 6 (Theorem 3.7). A first-moment estimate on $G(n, 1/2)$ suggests the true growth is $\Theta(\log n)$, and nothing here proves a growth rate at all.

Named graphs, and a pattern in the minimisers. Section 4 records exact values of ugp for 43 named graphs. A few restate [1] — complete graphs, paths, complete multipartite graphs and even cycles are comparability graphs, so their value is the order — and are included as anchors. The rest are new. Odd cycles *attain* the source's bound $n - 1$; the complements of cycles form a second family one below the maximum, $\mathrm{ugp}(\overline{C_n}) = n - 1$ for every $5 \leq n \leq 16$, and are not mentioned in [1] at all; the Petersen graph has ugp $= 7$ and its complement $T(5) = L(K_5)$ has ugp $= 6$; the Kneser graphs $K(6, 2)$ and $K(7, 2)$ have ugp $= 9$ and ugp $= 12$.

The minimisers we exhibit are more structured than a randomised search has any right to produce. Writing $K_m \,\square\, K_n$ for the rook's graph and $T(n) = L(K_n)$ for the triangular graph, the graphs attaining μ at orders $6, 8, 9, 10, 12, 15$ are

$$L(K_{2,3}), \quad L(K_{2,4}), \quad L(K_{3,3}), \quad L(K_5), \quad L(K_{3,4}), \quad L(K_6),$$

all line graphs of complete or complete bipartite graphs (Theorem 4.7). The pattern is a list of cases and not a theorem about the family: $\mathrm{ugp}(K_2 \,\square\, K_5) = 7$ at order 10, where $\mu(10) = 6$, so being such a line graph does not make a graph a minimiser.

What is not addressed. The lower parameter lgp, and with it the main conjecture of [1] — that $\mathrm{ugp}(G) > \mathrm{lgp}(G)$ for every connected graph of order at least 3, which its authors verify only to $n = 6$ — is untouched. Unlike ugp, lgp has no comparability reformulation: it is a genuine minimum over all $2^{|E|}$ orientations with shortest paths computed inside each, and none of the machinery here transfers to it.

Method, and how to read the computations. Every statement below is machine-checked in Lean 4 over Mathlib (Section 8). Many of them are exhaustive searches, and we state the finite search performed each time, because the searches are the content. Two asymmetries organise the work and are worth stating once.

Lower and upper halves of a single value. A bound $k \leq \mathrm{ugp}(G)$ is a witness: one k-subset and one linear order on it, cheap at any order. A bound $\mathrm{ugp}(G) < k + 1$ is an exhaustion of $\binom{n}{k+1}(k+1)!$ orderings by brute force, and that factorial is the obstruction that Section 5 removes, first by replacing the ordering search with a polynomial comparability test (Gallai's and Golumbic's implication-class criterion) and then by pruning the subset search with heredity.

Lower and upper halves of a value of μ. The bound $\mu(n) \geq m$ says that *every* graph of order n *accepts* some comparability m-subset; the bound $\mu(n) \leq m$ is one graph. Acceptance and refutation fail in opposite directions, and only one of them is dangerous: a refutation test that misses can cost a value but cannot manufacture one, whereas an over-permissive acceptance test yields a theorem that is false. Every acceptance route below is therefore stated as a proved sufficient condition and comes with negative controls (Section 6) refuting each of its weakenings.

2. Preliminaries

All graphs are finite, simple and undirected. We identify a graph on the vertex set $V = \{0, 1, \ldots, n-1\}$ with its adjacency function, and write $x \sim y$ for adjacency. For $S \subseteq V$ we write $G[S]$ for the induced subgraph.

Definition 2.1. A relation F on V is a *transitive orientation of* $G[S]$ if: $F(x, y)$ implies $x, y \in S$ and $x \sim y$; every edge of $G[S]$ is oriented, i.e. $x, y \in S$ and $x \sim y$ imply $F(x, y)$ or $F(y, x)$; F is asymmetric; and F is transitive. We call S a *comparability set* of G if such an F exists, i.e. if $G[S]$ is a comparability graph.

Definition 2.2. $\operatorname{ugp}(G)$ is the largest cardinality of a comparability set of G. We write $\mu(n)$ for the minimum of $\operatorname{ugp}(G)$ over all graphs G of order n.

By Theorem 1.1 this agrees with the maximum of gp over orientations of G, and we use Definition 2.2 throughout. We stress that nothing below re-proves Theorem 1.1: each value we compute is literally a value of "largest order of an induced comparability subgraph", and it is a value of ugp on the authority of [1].

The empty set is a comparability set, and a comparability set has all its subsets as comparability sets — restrict the orientation — so "largest order" is well defined and is not merely "maximal order". The first proposition is what makes ugp computable at all.

Proposition 2.3. *Let S be a set of vertices of G. Then $G[S]$ is a comparability graph if and only if there is an enumeration $v_1, v_2, \ldots, v_k$ of S without repetitions such that*

$$v_i \sim v_j \ \text{ and } \ v_j \sim v_l \implies v_i \sim v_l \qquad \text{for all } 1 \leq i < j < l \leq k.$$

Proof sketch. Given such an enumeration, orient every edge of $G[S]$ forward along it. Asymmetry and totality are immediate from the absence of repetitions, and transitivity is exactly the displayed triple condition.

Conversely, let F be a transitive orientation of $G[S]$ and rank each vertex by its number of F-predecessors inside S:

$$r(x) = \#\{y \in S : F(y, x)\}.$$

If $F(x, y)$ then transitivity gives $\{z : F(z, x)\} \subseteq \{z : F(z, y)\}$, and asymmetry puts x in the second set and not in the first, so the inclusion is strict and $r(x) < r(y)$. Enumerating S in increasing order of r therefore lists every F-arc forwards, and the triple condition follows from transitivity of F. □

The rank argument replaces a topological sort and is three lines; we record it because it is the step that converts an existential over relations into an existential over linear orders, which is a finite search.

Proposition 2.4. *Let $e : V \hookrightarrow W$ be injective and let G be a graph on W. Then* $\operatorname{ugp}(G[e(V)]) \leq \operatorname{ugp}(G)$*; that is,* ugp *is monotone under induced subgraphs. Consequently μ is non-decreasing. Moreover, if G' is G together with one isolated vertex then* $\operatorname{ugp}(G') = \operatorname{ugp}(G) + 1$*, and hence*

$$\mu(n) \ \leq \ \mu(n+1) \ \leq \ \mu(n) + 1 \qquad (n \geq 1).$$

Proof sketch. For the first claim, push a transitive orientation F of an induced subgraph forward along e: set $\tilde{F}(a, b)$ if $a = e(x)$, $b = e(y)$ and $F(x, y)$. Each of the four clauses of Definition 2.1

transfers in two lines; injectivity is used exactly once, in asymmetry, and is necessary there (Proposition 6.2). The pull-back along e is proved in the same way, so the two directions are one statement.

For the isolated vertex v: adding v to a largest comparability set keeps it one, because v has no edges and totality has nothing to orient, which gives $\geq$; deleting v and pulling back gives $\leq$. The sandwich for μ follows by applying the two halves to a minimiser at order n and at order $n+1$. $\square$

Monotonicity is the lever behind most of Section 3: a lower bound on μ proved at one order is a lower bound at every larger order, for free.

Proposition 2.5. *For every $n \geq 4$ and every graph G of order n, $\mathrm{ugp}(G) \geq 4$; for $n \geq 6$, $\mathrm{ugp}(G) \geq 5$; for $n \geq 8$, $\mathrm{ugp}(G) \geq 6$. In other words μ passes the thresholds $(4, 6, 8) \mapsto (4, 5, 6)$.*

Proof sketch. Restrict G to its first 4 (resp. 6, 8) vertices, take there the comparability set produced by the corresponding exhaustive sweep of Section 3, and read it back along Proposition 2.4. Each of the three bounds thus costs a single sweep, at a single order, and holds at all larger orders with no diameter argument: the first re-derives [1, Corollary 4.7] from the sweep over the $2^6 = 64$ labelled graphs on four vertices, and the second improves it by one at every order $n \geq 6$ from the sweep over the 2^{15} labelled graphs on six vertices. $\square$

Proposition 2.6. $\mathrm{ugp}(G) \geq 7$ *for every graph G of order $N \geq 11$; equivalently the class $\{G : \mathrm{ugp}(G) \leq 6\}$ is empty from order 11 upwards. Consequently $\mu(12) \in \{7, 8\}$.*

Proof sketch. The bound at $N = 11$ is Theorem 3.9; it travels upwards by Proposition 2.4, and the isolated-vertex sandwich caps $\mu(12)$ at $\mu(11) + 1 = 8$. $\square$

3. The extremal function $\mu(n)$

Throughout this section a *sweep at order n* means an exhaustive enumeration of all $2^{\binom{n}{2}}$ labelled graphs on n vertices, each tested for a comparability subset of the required order; the test is a depth-first search whose answer is re-checked against Proposition 2.3, so a bug in the search can lose a value but cannot produce one. The enumeration is by a bijection between labelled graphs and the naturals below $2^{\binom{n}{2}}$, whose surjectivity onto graphs is proved, so that a statement about all codes is a statement about all graphs.

3.1. **Orders 2 to 7.**

Theorem 3.1. *$\mu(4) = \mu(5) = 4$ and $\mu(6) = \mu(7) = 5$.*

Proof sketch. Each value has two halves. The lower halves are sweeps: all 2^6 labelled graphs of order 4, all 2^{10} of order 5, all 2^{15} of order 6 and all 2^{21} of order 7; in each case every single graph is found to carry a comparability subset of the stated order, and the subset and its ordering are re-checked for distinctness, cardinality and the triple condition of Proposition 2.3. The upper halves are one graph each: a 5-cycle at orders 5 and (with an isolated vertex added) 6, and a graph M_7 of order 7 and its isolated-vertex extension M_8 at orders 7 and 8. The sweep at order 7 is the most expensive of the four, at 6 min 38 s of machine time, about 190 μs per graph. $\square$

Remark 3.2. Two of those four sweeps are redundant. By Proposition 2.4 the bound $4 \leq \mathrm{ugp}$ at order 5 follows from the bound at order 4, and $5 \leq \mathrm{ugp}$ at order 7 from the bound at order 6, so the table $\mu(4), \ldots, \mu(9)$ needs sweeps at three orders and not at five — and the two that survive are exactly the two orders at which μ steps. Monotonicity transports a bound only upwards, so a step must always be paid for. The re-derived statements are checked to be literally the original ones and not lookalikes.

Remark 3.3. An exhaustive isomorph-classification, external to the machine-checked development, records that $\mu(5) = 4$ is attained by C_5 and, up to isomorphism, by nothing else; that $\mu(6) = 5$ is attained by exactly 12 isomorphism classes; and that $\mu(7) = 5$ is attained by exactly two classes, which are complementary to each other. The same computation shows that no graph on eight vertices has ugp $= 5$, although 1 680 labelled graphs on seven vertices do: the parameter skips a value.

3.2. Order 8, and the free ride to orders 9 and 10.

Theorem 3.4. $\mu(8) = 6$*: every graph on eight vertices has an induced comparability subgraph of order* 6*, and some graph on eight vertices has none of order* 7*.*

Proof sketch. The upper half is the graph M_8 of Theorem 3.1, with $\mathrm{ugp}(M_8) = 6$. The lower half is a sweep over all $2^{28} = 268\,435\,456$ labelled graphs of order 8, split into sixteen intervals (about one CPU-hour in total) whose union is reassembled by a case split on the code. Two proved sufficient conditions, tried before the search, settle about three quarters of the graphs outright (Lemmas 5.6 and 5.7), and the backtracking search runs on the remaining quarter. Section 5 gives a second, independent proof of the same lower bound from 415 564 partial adjacency matrices instead of 2^{28} complete ones. □

Proposition 3.5. $\mu(4) = 4$*,* $\mu(5) = 4$*,* $\mu(6) = 5$*,* $\mu(7) = 5$ *and* $\mu(8) = 6$*, as a single statement.*

Past order 8 the lower halves stop costing anything, and each further value of μ becomes a question about a single graph.

Theorem 3.6. $\mu(9) = 6$*.*

Proof sketch. The lower half is free: by Proposition 2.5, $\mathrm{ugp}(G) \geq 6$ for every graph of order $N \geq 8$, with no new evaluation of any kind. The upper half is one graph M_9 on nine vertices with 21 edges and degree sequence $4^4 5^4 6$, for which $\mathrm{ugp}(M_9) = 6$: a comparability 6-set is exhibited, and the exhaustion is a search over the 36 seven-element subsets of $V(M_9)$, each refuted. The graph was found by hill-climbing on the number of comparability 7-subsets, and both independent implementations of the comparability test confirm its value. For comparison, a direct sweep at order 9 would be 2^{36} labelled graphs, about 229 CPU-hours at the measured rate.

It is worth noting that M_9 is *not* the order-8 minimiser plus an isolated vertex, and cannot be: that graph has ugp $= 7$. The recipe that produced minimisers at orders 6 and 8 from those at orders 5 and 7 stops working here, and every vertex of M_9 has a neighbour. □

Theorem 3.7. $\mu(10) = 6$*. In particular* $\mu(10) \neq \lfloor 10/2 \rfloor + 2 = 7$*, so the formula* $\lfloor n/2 \rfloor + 2$*, which agrees with* $\mu(n)$ *for* $4 \leq n \leq 8$*, is false.*

Proof sketch. Again the lower half is Proposition 2.5. The upper half is one graph M_{10} on ten vertices with 24 edges and degree sequence $4^3 5^6 6$, with $\mathrm{ugp}(M_{10}) = 6$: the exhaustion is over the 120 seven-element subsets. M_{10} was found by scanning all $2^9 = 512$ ways of attaching a tenth vertex to a nine-vertex witness. A direct sweep at order 10 would be 2^{45} labelled graphs, some 13 years of CPU time.

We emphasise what this does and does not say. It refutes a formula; it proves nothing about the growth rate of μ. □

Proposition 3.8. $\mu(2) = 2$ *and* $\mu(3) = 3$*, and the nine values* $\mu(2), \ldots, \mu(10)$ *hold simultaneously, together with the sandwich* $\mu(n) \leq \mu(n+1) \leq \mu(n) + 1$ *at all eight consecutive pairs: the row reads step, step, no, step, no, step, no, no.*

Proof sketch. The two small values are two-line sweeps over the 2^1 and 2^3 codes, and both say the same thing: every graph on at most three vertices is a comparability graph, the smallest one that is not being C_5. They had been printed in every table of this parameter and never proved, because the source's own bound ugp ≥ 4 starts at $n = 4$. The step statements are Proposition 2.4 instantiated at each consecutive pair. □

3.3. **Order 11: where the value 6 ends.** The class $\mathcal{C}_6 = \{G : \mathrm{ugp}(G) \le 6\}$ — graphs no seven of whose vertices induce a comparability graph — is closed under induced subgraphs. That is the lever: every member of $\mathcal{C}_6$ at order $N+1$ is a one-vertex extension of a member at order N, so the class can be built level by level and either survives to order 11 or does not. It does not.

Theorem 3.9. $\mu(11) = 7$. *Consequently* $\max\{N : \mu(N) = 6\} = 10$.

Proof sketch. The upper half is M_{10} together with an isolated vertex, which has ugp $= 7$ by Proposition 2.4. The lower half is the classification of order 10 (Theorem 3.10) together with Theorem 3.11: an order-11 graph with ugp ≤ 6 restricts to an order-10 graph with ugp ≤ 6, hence to one of two named graphs up to relabelling, and neither of those admits an eleventh vertex. The relabelling is lifted to 11 vertices fixing the new one.

The result was first obtained by an exhaustive enumeration of $\mathcal{C}_6$ up to isomorphism, level by level from its 220 classes at order 7:

$$220 \ \to\ 226 \ \to\ 30 \ \to\ 2 \ \to\ 0 \qquad \text{at orders } 7, 8, 9, 10, 11.$$

The base is the count of order-7 graphs that are not comparability graphs, $1044 - 824 = 220$; the count rises before it collapses, which a hereditary class is free to do. That enumeration is reproduced by the machine-checked route of Section 5 at orders 8, 9 and 10. □

Theorem 3.10. *Exactly two graphs of order* 10*, up to isomorphism, satisfy* $\mathrm{ugp}(G) \le 6$*: the graph* M_{10} *of Theorem 3.7, and the triangular graph* $T(5) = L(K_5)$*, the complement of the Petersen graph. Both have* ugp $= 6$*.*

Proof sketch. "At most two" comes from a single exhaustive search over 896 060 partial adjacency matrices, pruned in two independent ways — by heredity, since a partial matrix already carrying a comparability 7-set needs no completion, and by isomorph rejection (Proposition 5.13) — leaving seven surviving codes, each of which is carried onto M_{10} or onto $T(5)$ by an explicitly exhibited permutation that the verification re-checks. "Exactly two" needs both values to be attained, which they are, and the two graphs to be non-isomorphic, which follows from their edge counts, 24 against 30, once the edge count is proved to be an isomorphism invariant. □

Theorem 3.11. *Let G be a graph on* 11 *vertices whose first ten vertices induce* M_{10}*, or induce* $T(5)$*. Then* $\mathrm{ugp}(G) \ge 7$*. Equivalently, none of the* $2 \cdot 2^{10}$ *one-vertex extensions of the two order-*10 *minimisers lies in* $\mathcal{C}_6$*.*

Proof sketch. This is the half of Theorem 3.9 that is a finite statement about named graphs. For each of the 2^{10} possible neighbourhoods of the eleventh vertex, the search *exhibits* a comparability 7-subset, which is then re-checked from scratch by the certified acceptance test of Proposition 5.4. This is the acceptance direction, in which a bug in the search cannot manufacture a theorem. The step that turns "the sweep was started at level 10 with this code" into "every graph on 11 vertices whose first ten vertices induce this graph" is the colexicographic encoding of Proposition 5.11. □

Theorem 3.12. *The triangular graph* $T(5) = L(K_5) = J(5, 2)$*, the complement of the Petersen graph, is one of exactly two extremal graphs at order* 10*.*

Proof sketch. Immediate from Theorem 3.10 together with $\mathrm{ugp}(T(5)) = 6$ (Theorem 4.1). It is the second class of Theorem 3.10 that is striking: a member of $\mathcal{C}_6$ must be tightly Ramsey-constrained, and $T(5)$ is 6-regular with independence number 2 and clique number 4; its complement, the Petersen graph, has ugp $= 7$ — one more, and enough to keep it out of the class. The two extremal graphs have $|\mathrm{Aut}| = 120$ and $|\mathrm{Aut}| = 12$ respectively, and $10!/12 + 10!/120 = 332\,640$ is exactly the number of labelled order-10 graphs in $\mathcal{C}_6$ found by an isomorphism-free enumeration. The degree, independence, clique and automorphism-group data quoted in this proof are computed externally and are not part of the machine-checked development; of the invariants of the two graphs only their edge counts, 24 and 30, are machine-checked, and those are the only ones the classification uses, since the edge count is proved to be an isomorphism invariant and so separates the two graphs. □

3.4. **Orders** 12 **to** 16.

Theorem 3.13. *$\mu(12) = 7$, attained by the 3×4 rook's graph $K_3 \,\square\, K_4 = L(K_{3,4})$, equivalently the circulant $C_{12}(3, 4, 6)$.*

Proof sketch. The lower half is free by Proposition 2.6. The whole content is one graph on twelve vertices with ugp ≤ 7, and the rook's graph is one: a comparability 7-set is exhibited, and the exhaustion — no eight of its twelve vertices induce a comparability graph — is a hereditarily pruned search visiting 877 nodes. The graph is defined from its board rather than transcribed, so there is no data to validate. The complementary route to $\mu(12) = 8$, namely showing that the hereditary class $\{G : \mathrm{ugp}(G) \leq 7\}$ dies at order 12, was priced from a validated model at more than 46 CPU-hours for one of its four remaining levels and was not run; it is unnecessary, since a witness settles the cell. □

Theorem 3.14. *$\mu(N) = 7$ for every N with $11 \leq N \leq 15$.*

Proof sketch. One graph does all five orders. The triangular graph $T(6) = L(K_6) = J(6, 2)$ has fifteen vertices, is 8-regular, and satisfies $\mathrm{ugp}(T(6)) = 7$: a comparability 7-set is exhibited, and the exhaustion over its 8-subsets is a pruned search visiting 7 008 nodes. Since ugp is monotone under induced subgraphs (Proposition 2.4), every induced subgraph of $T(6)$ on $11, 12, 13, 14$ or 15 vertices has ugp ≤ 7; the matching lower bound is free at every order ≥ 11 by Proposition 2.6. The 225 entries of the adjacency matrix of $T(6)$ are checked against a generated definition of the triangular graph, so the value is a value of $L(K_6)$ and not of a transcription.

Thus μ is now flat at 7 across five consecutive orders, after two at 4, two at 5 and three at 6; the whole row $\mu(2), \ldots, \mu(15)$ holds as a single statement. □

Theorem 3.15. *At least two isomorphism classes attain $\mu(12) = 7$: the rook's graph $K_3 \,\square\, K_4$ and the circulant $C_{12}(1, 4, 5, 6)$.*

Proof sketch. Both values are proved as in Theorem 3.13. The two graphs are 5-regular and 7-regular, hence have 30 and 42 edges, and the edge count is an isomorphism invariant, so they are not isomorphic. A scan of all 63 circulants on twelve vertices finds exactly these two with ugp $= 7$ and none with a smaller value. This is the order-12 shadow of Theorem 3.10 minus the classification: *at least* two, where order 10 has exactly two. How many classes attain $\mu(12)$ is open. □

Proposition 3.16. *$\mu(16) \in \{7, 8\}$, and the upper end is attained by a connected 6-regular graph: $\mathrm{ugp}(K_4 \,\square\, K_4) = 8$, so $\mu(16) \leq 8$ without the isolated-vertex construction. In particular $K_4 \,\square\, K_4$ attains $\mu(16)$ exactly when $\mu(16) = 8$.*

Theorem 3.17. *Let G be a graph on 16 vertices whose first fifteen vertices induce $T(6) = L(K_6)$ and whose sixteenth vertex is adjacent to none of the four vertices of $T(6)$ corresponding to the pairs $\{2, 5\}, \{3, 4\}, \{3, 5\}, \{4, 5\}$. Then $\mathrm{ugp}(G) \geq 8$.*

Proof sketch. These are $2^{11} = 2\,048$ of the $2^{15} = 32\,768$ one-vertex extensions of $T(6)$, and each is refuted by exhibiting a comparability 8-subset through the new vertex, re-checked from scratch by the acceptance test of Proposition 5.4. The index is meaningful because the bits of the code above the order-15 block are exactly the edges from the new vertex.

The remaining 30 720 extensions are not machine-checked, and we state that precisely. The corresponding exhaustive search is 32 769 nodes and was measured at 0.33 s per node inside the proof assistant, hence about 2.9 CPU-hours, which exceeded the budget; an independent implementation of a different comparability algorithm finds 0 survivors among all 32 768 extensions. Every statement that needs the full search carries it as an explicit hypothesis, so that no result is stated as proved that is not.

Even with that hypothesis discharged, the conclusion would not be $\mu(16) \geq 8$: it says nothing about order-16 graphs whose first fifteen vertices induce something other than $T(6)$, and no second order-15 graph with ugp ≤ 7 is known. □

Remark 3.18. Several complete but unverified scans, external to the machine-checked development, bear on $\mu(16)$ without settling it. None of the 105 graphs obtained from $T(6)$ by flipping one pair, and none of the 5 460 obtained by flipping two, has $\mathrm{ugp} \le 7$; since $\mathbb{Z}_{15}$ is the only group of order 15, the 128 circulants on fifteen vertices are all the Cayley graphs there and none has $\mathrm{ugp} \le 7$ — and $L(K_6)$ is not among them, since S_6 has no element of order 15; all 36 096 Cayley graphs of the five abelian groups of order 16 have $\mathrm{ugp} \ge 8$. These say that $T(6)$ is an isolated minimiser rather than that it is the only one, and they explain why local search does not find it: a hill-climber has no downhill neighbour to arrive from.

4. Values for named graphs

Theorem 4.1. *The following values of* ugp *hold.*

graph	*order*	ugp
K_n, $3 \le n \le 9$	n	n
P_n, $3 \le n \le 10$	n	n
$K_{2,2}, K_{2,4}, K_{3,3}, K_{3,4}, K_{4,4}$; $K_{2,2,2}, K_{2,2,2,2}, K_{3,3,3}$	n	n
C_n, *n even,* $4 \le n \le 10$; C_3	n	n
C_n, *n odd,* $n = 5, 7, 9$	n	$n-1$
$\overline{C_n}$, $5 \le n \le 10$	n	$n-1$
Petersen graph	10	7
$T(5) = L(K_5) = \overline{\textit{Petersen}}$	10	6
M_5, M_6, M_7, M_8 *(the minimisers of Theorems 3.1 and 3.4)*	$5, 6, 7, 8$	$4, 5, 5, 6$
$K(6, 2)$	15	≥ 9

Proof sketch. The values equal to the order need only a witness: one enumeration of all n vertices passing the triple condition of Proposition 2.3, together with the trivial bound $\mathrm{ugp}(G) \le n$. These are exactly the graphs that are comparability graphs, and for complete graphs, paths and even cycles the values restate [1]; the complete multipartite values follow from [1, Corollary 4.5] and the classical fact that complete multipartite graphs are comparability graphs. They are included as anchors: they are the check that the quantity being computed is the source's.

A value $k < n$ has two halves: a witness of order k, and an exhaustion establishing that no $(k+1)$-subset admits any enumeration passing the triple condition. The exhaustion runs over $\binom{n}{k+1}(k+1)!$ orderings; for the Petersen graph that is $\binom{10}{8} \cdot 8! = 1\,814\,400$ orderings, checked in 69 s, and for $\overline{C_{10}}$ it is $10! = 3\,628\,800$. For odd cycles the upper bound $n-1$ is the source's; that it is *attained* is proved here. The complements of cycles are not treated in [1] at all.

The Petersen graph is given by ten transcribed row bitmasks, so the headline value is about the Petersen graph only if the transcription is right. An explicit bijection onto the Kneser graph $K(5, 2)$, built from the definition of a Kneser graph rather than from the same data, establishes that it is; the check is independent of the transcription it validates. □

Remark 4.2. At the Petersen graph the source's own lower bound is exactly tight: $\alpha_2 = 7$ and $\alpha^\omega = 6$ (both computed externally), so $\max\{\alpha^\omega, \alpha_2\} = 7 = \mathrm{ugp}$, and [1, Corollary 4.6] therefore *determines* that entry. This makes the Petersen value less surprising than the others; $\mathrm{ugp}(T(5)) = 6$, well below its order, shows that the parameter does not simply track α_2 in general.

The factorial in the exhaustion is the obstruction of Theorem 4.1, and the implication-class criterion of Section 5 removes it.

Theorem 4.3. $\mathrm{ugp}(\overline{C_n}) \ge n-1$ *for every n, and* $\mathrm{ugp}(\overline{C_n}) = n-1$ *for every* $5 \le n \le 16$*. Moreover* $\mathrm{ugp}(C_n) = n-1$ *for the odd orders* $n = 11, 13, 15$*, and* $\mathrm{ugp}(K(6, 2)) = 9$.

Proof sketch. The lower bound for complements of cycles is an argument rather than a computation, and it holds at every order: delete the vertex $n-1$ and take the rest in their natural order $0 < 1 < \cdots < n-2$. Inside that set i and j are adjacent exactly when $j \ne i+1$, so a chain $i \to j \to k$ with

$i < j < k$ has $k \geq i+2$, which is precisely what the arc $i \to k$ asks for. Neither hypothesis of the triple condition is used — the conclusion follows from $i < j < k$ alone — so all the content of the value is in the upper half.

The upper halves are exhaustions over subsets, not over orderings: each subset of the critical order is refuted by exhibiting an implication class that contains a directed edge together with its reverse (Section 5). For $K(6,2)$ this is $\binom{15}{10} = 3\,003$ subsets, where the ordering-based route would have been about 1.1×10^{10} orderings; for $\overline{C_{16}}$ it replaces $16! \approx 2.1 \times 10^{13}$ orderings. The whole file of these values checks in 15.9 s. ☐

Theorem 4.4. $\mathrm{ugp}(K(7,2)) = 12$*: twelve of the twenty-one pairs from* $\{0,\dots,6\}$ *induce a transitively orientable subgraph of the Kneser graph, and no thirteen do.*

Proof sketch. The lower half is one exhibited enumeration of twelve pairs passing the triple condition. The upper half is a single depth-first search over increasing subsets, pruned by heredity (Proposition 5.8): a subset already refuted can lie inside no comparability set, so its whole branch dies. The search visits 30 026 nodes, of mean order 8.00, against the $\binom{21}{13} = 203\,490$ flat refutations the unpruned route would need; at level 13 it reaches 42 subsets and refutes all of them, which is the theorem. In wall time this is 25.8 s against a search that did not finish in 25 minutes.

The twenty-one row bitmasks are verified against a generated definition of $K(7,2)$, so the value is a value of the Kneser graph (Proposition 5.10). ☐

Proposition 4.5. $\mathrm{ugp}(K(7,2)) = 12$ *with no witness supplied by hand: a search inside the verified development returns the twelve pairs* $01, 02, 03, 04, 05, 06, 12, 13, 23, 45, 46, 56$ *— a different* 12*-set from the hand-supplied one — and the certified acceptance test of Proposition 5.4 re-checks it.*

Proposition 4.6. $\mathrm{ugp}(\textit{Petersen}) = 7$ *and* $\mathrm{ugp}(K(6,2)) = 9$*, with both halves polynomial in the order of the graph and no hand-supplied witness: the lower halves are produced by a search whose output is re-checked by the acceptance test of Proposition 5.4, and the upper halves by the refutation criterion of Theorem 5.1.*

Proof sketch. Neither value is new here: both are proved by three or four routes that share no code, which is the point of restating them. A comparability test that refuted too much would show up as a *smaller* value, and one that accepted too much as a larger one; agreement of independent routes on both values is the cheapest available check that neither happens. The measured effect of replacing the ordering search by the implication-class search on the same instance is 15 416 ms to 35 ms, and $2\,844\times$ on the exhaustive negative. ☐

Theorem 4.7. *Each of the following is a value of* ugp *together with the assertion that it is the minimum at that order:*

$$\begin{array}{lll} \mathrm{ugp}(K_2 \,\square\, K_3) = 5 = \mu(6), & \mathrm{ugp}(K_2 \,\square\, K_4) = 6 = \mu(8), & \mathrm{ugp}(K_3 \,\square\, K_3) = 6 = \mu(9), \\ \mathrm{ugp}(T(5)) = 6 = \mu(10), & \mathrm{ugp}(K_3 \,\square\, K_4) = 7 = \mu(12), & \mathrm{ugp}(T(6)) = 7 = \mu(15). \end{array}$$

Since $K_m \,\square\, K_n = L(K_{m,n})$ *and* $T(n) = L(K_n)$*, the six minimisers this development exhibits at orders* $6, 8, 9, 10, 12, 15$ *are the line graphs* $L(K_{2,3}), L(K_{2,4}), L(K_{3,3}), L(K_5), L(K_{3,4}), L(K_6)$.

Proof sketch. Each value is a witness together with a hereditarily pruned exhaustion, as in Theorem 3.13; the rook's graphs are defined from their boards, so there are no data to validate. The minimality assertions are Theorems 3.1, 3.4, 3.6, 3.7, 3.13 and 3.14. ☐

Remark 4.8. The pattern of Theorem 4.7 is a list of cases and not a theorem about the family. Two controls bound it. First, $\mathrm{ugp}(K_2 \,\square\, K_5) = 7$ at order 10, where $\mu(10) = 6$: a rook's graph need not be a minimiser. Second, $\mathrm{ugp}(\text{icosahedron}) = 8$ at order 12, where $\mu(12) = 7$: being 12-vertex, regular and vertex-transitive is not what makes $K_3 \,\square\, K_4$ extremal. We have no explanation of the pattern, and finding one — or a counterexample at a larger order — is the question we would most like answered.

5. The comparability tests, and the searches they make possible

This section collects the tools. None of the mathematics in §5.1 is new — it is Gallai's and Golumbic's — but the searches of Sections 3 and 4 are built on it, and their soundness is what licenses reading the computations as proofs.

5.1. **Refutation and acceptance by implication classes.** For a graph G, the *forcing relation* Γ of Gallai and Golumbic, defined on the directed edges of G, is

$$(a,b)\ \Gamma\ (a',b') \quad \text{iff} \quad \big(a=a' \text{ and } bb' \notin E\big) \ \text{ or } \ \big(b=b' \text{ and } aa' \notin E\big).$$

Γ is symmetric, so its transitive closure partitions the directed edges into *implication classes*. Call a class *consistent* if it does not contain a directed edge together with its reverse. A transitive orientation contains an implication class entirely or not at all: if it contains (a,b) and $(a,b)\,\Gamma\,(a,d)$ with $bd \notin E$, then orienting $d \to a$ would give $d \to b$ by transitivity, so bd would be an edge. Hence a graph with an inconsistent implication class is not a comparability graph. This is the *refutation* half, and it is what an upper bound on ugp consumes. A missing converse can cost a value; it cannot produce a wrong one.

Theorem 5.1 (Gallai)**.** *A consistent implication class is a transitive relation: if (a,b) and (b,c) lie in it, so does (a,c).*

Proof sketch. Fix the tail a of one edge (a,b) of the class and propagate along the class the statement

$$P(p,q): \quad (a,p) \in \text{class} \implies (a,q) \in \text{class}.$$

The base $P(a,b)$ is free, because its conclusion is the hypothesis one is given. Each Γ-step preserves P, using only the *shortcut lemma*: a consistent class containing (x,y) and (y,w) forces xw to be an edge, since otherwise $(x,y)\,\Gamma\,(w,y)$ and the class would contain (y,w) together with (w,y). Reading P at the far end (b,c) gives transitivity. This avoids the Triangle Lemma and the canonical Γ-chains of the textbook route; the statement is Gallai's and Golumbic's, only the route is different. □

Lemma 5.2. *Let C be a consistent implication class of G and let F' be* any *transitive orientation of the graph obtained from G by deleting the undirected edges underlying C. Then $C \cup F'$ is a transitive orientation of G.*

Proof sketch. Four transitivity cases. The case C, C is Theorem 5.1. Each mixed case needs three steps: the shortcut is an edge, since otherwise a Γ-step would put a deleted edge into C; and it is oriented the right way, whether it lies under C (by Theorem 5.1) or not (by transitivity of F'). The word *any* is load-bearing: a transitive orientation of G need *not* restrict to one of the smaller graph. On K_3 with $C = \{(0,1)\}$ and $F = \{(0,1),(0,2),(2,1)\}$, the restriction $\{(0,2),(2,1)\}$ on the path $0-2-1$ is not transitive, although the path is a comparability graph — via the other orientation. □

Theorem 5.3 (Golumbic's TRO theorem)**.** *If the greedy Γ-decomposition of G — pick a directed edge, check that its implication class is consistent, delete that class's undirected edges, repeat — runs to completion, then G is a comparability graph.*

Proof sketch. Induction on the derivation, one application of Lemma 5.2 per step. Termination is not needed for the statement, but is proved anyway: each step strictly decreases the number of edges. The theorem is quoted from [3]; the equivalence stated there is between (i) G is a comparability graph, (ii) every implication class of G is consistent, and (iii) every class of a Γ-decomposition is consistent. □

Proposition 5.4. *There is a polynomial-time* certified *acceptance test for comparability: an untrusted procedure builds an orientation from the implication classes, and the four clauses of Definition 2.1 are then verified directly on the result, in time cubic in the size of the vertex list. A bug in the construction can only fail to produce a certificate; it cannot produce a wrong one.*

Proof sketch. Making the run of the greedy decomposition itself a verified fact would require proving the class-closure algorithm complete. That is avoidable: the algorithm hands back the orientation it built, and verification re-checks it against the definition, never mentioning the search. Theorem 5.3 then plays its proper role — it is the *guarantee* that the untrusted search does not fail on a comparability graph. The measured effect against the ordering search on the same instance is 15 416 ms to 35 ms. □

Remark 5.5. The implication (ii)⇒(iii) of Theorem 5.3 — that consistency survives the deletion of an implication class — is not verified here. It is not automatic: deleting edges increases non-adjacency, so Γ grows and classes merge. The obvious repair is false, by the K_3 example of Lemma 5.2; the route that works is Gallai's modular decomposition and was not attempted. Nothing above depends on it: the refutation and the acceptance tests each carry their own certificate, and a graph on which neither fired would simply yield no theorem. Exhaustively, over all 2^{10} labelled graphs on five vertices and all 2^{15} on six, the acceptance test accepts exactly the graphs the refutation test does not refute.

5.2. **Cheap sufficient conditions for acceptance.** A bound $\mu(n) \geq m$ is an acceptance per graph, so what makes a sweep cheap is a sufficient condition with a one-line orientation, tried before the search.

Lemma 5.6. *A bipartite induced subgraph is a comparability graph.*

Proof sketch. Orient every edge from its first-class end to its second-class end. Asymmetry and totality are immediate, and transitivity is vacuous: a chain $x \to y \to z$ would need y to lie in both classes. □

Lemma 5.7. *Let b label the vertices of S so that, inside S and for $x \neq y$, $x \sim y$ if and only if $b(x) = b(y)$ — that is, $G[S]$ is a disjoint union of cliques. Then S is a comparability set. The same holds when S carries a rank function r with $x \sim y$ implying $r(x) \neq r(y)$ and the induced relation transitive.*

Proof sketch. For the block version, orient every edge by the ambient linear order on the vertices. Transitivity is the only clause with content: $x \sim y$ and $y \sim z$ give $b(x) = b(y) = b(z)$, and $x < z$ with $b(x) = b(z)$ gives $x \sim z$ from the other direction of the equivalence. Both directions of that equivalence are load-bearing (Proposition 6.1). □

Measured exhaustively over all 2^{28} graphs of order 8, a bipartite 6-subset exists in 74.38% of them and a union-of-cliques 6-subset in 10.31%; one or the other in 76.11%. The backtracking search runs on the remaining quarter.

5.3. **Pruning the subset search.**

Proposition 5.8. *Comparability passes to induced subgraphs, so a refuted subset kills every superset. A depth-first search over increasing subsets that refutes at every node and prunes the branch on the first refutation is sound: if it reports that no m-subset extending a given partial set is a comparability set, then none is.*

Proof sketch. Induction on the list of remaining vertices. The pruning step is three lines: the refutation criterion fires on $P \cup \{v\}$, so that set is not a comparability set, and heredity carries the refutation to every superset. Only the affirmative answer carries an obligation, so a bug in the branching returns "false" and produces nothing. □

Proposition 5.9. *On $K(7,2)$ at $m = 13$ the pruned search visits 30 026 nodes, of mean order 8.00, against $\binom{21}{13} = 203\,490$ flat refutations, and finishes in 25.8 s where the flat route timed out past 25 minutes. No compiled shared library, no chunk split and no certificate file were needed.*

Proposition 5.10. *The twenty-one row bitmasks used for $K(7,2)$ define the Kneser graph: all 441 entries agree with the graph generated from a lexicographic list of the 2-subsets of $\{0,\dots,6\}$ and a disjointness test.*

5.4. **Pruning the search over graphs.** For a bound $\mu(n) \geq m$ the object being searched is the set of graphs, and the same hereditary lever applies to it.

Proposition 5.11. *A partial adjacency matrix — the adjacencies among the first k vertices only — whose own induced subgraph already carries a comparability m-set needs no completion, since $m \leq \mathrm{ugp}$ is monotone in the order (Proposition 2.4). The depth-first search that builds the matrix one vertex at a time and kills such a prefix on the spot is sound: if it reports success, then every graph of order N agreeing with the given prefix satisfies $m \leq \mathrm{ugp}$.*

Proof sketch. Two pieces of bookkeeping make the statement say what it should. The code of a graph is colexicographic: pairs are listed by their larger element, so that the subgraph on the first k vertices occupies the low bits and "prefix of the code" and "induced subgraph on the first k vertices" are literally the same operation. And at each node only the m-subsets through the newest vertex are tested; every other one lives in the parent prefix and was tested at the level of its own largest vertex. Soundness is proved in one direction only, so a branching bug returns "false" and produces nothing. □

Remark 5.12. The prefix search re-proves the lower half of Theorem 3.4 from 415 564 partial matrices instead of 268 435 456 complete ones, a factor of 646, in a single evaluation with no chunk split, no compiled shared library and no certificate file, where the direct sweep needed sixteen chunks and about an hour of CPU. The surviving prefixes per level are $1, 2, 8, 64, 1024, 2604, 1680, 0$; here $2604 = 2^{15} - 30\,164$ is the number of labelled non-comparability graphs on six vertices and 1680 the number of labelled graphs of order 7 with $\mathrm{ugp} = 5$. The re-derived statement is checked to be literally the original one.

Proposition 5.13. *Order the codes of graphs by their list of rows, read left to right (row j being the neighbourhood of vertex j inside $\{0,\dots,j-1\}$). Prune a prefix whenever an* exhibited *permutation of $\{0,\dots,k-1\}$ gives a strictly larger code. This is sound for any set of candidate permutations: the maximal labelling of every isomorphism class survives every such prune, so the search still reaches a representative of every class.*

Proof sketch. Two facts. First, if two codes differ first at row j then they compare the same way at every level $k \geq j$, because the rows below j are equal and the rows above are each too small to catch up: a decision taken at level j is final. Second, a permutation of $\{0,\dots,j-1\}$ extended by the identity touches only the rows $1,\dots,j-1$, so an improving permutation of a prefix is an improving permutation of every completion. Together: if c is the maximal labelling of its class, no prefix of c admits an improving permutation.

Only the firing of the prune is trusted, and what it asserts is exhibited — a specific permutation and a row-by-row comparison, re-run during verification. Enlarging the candidate set can only prune more; missing a permutation can only make the search bigger. This is orderly generation in the sense of Read and Faradžev [6, 7], with the soundness of the prune proved rather than assumed, so that the improving permutation may come from an untrusted search. In particular no maximum over the symmetric group is ever formed: the soundness argument is a downward induction on the code, a pruned prefix handing back a strictly better labelled graph.

Measured at order 10 with $m = 7$: 2 804 129 868 nodes without the prune, 896 060 with it, a factor of 3129; an exact canonical form would give 112 668, eight times fewer again, at a cost per node that makes it slower overall. □

Proposition 5.14. *The same route reproduces known counts it was not fitted to. Its survivor counts at orders $2,\dots,10$ are $2, 4, 11, 45, 307, 1172, 2211, 310, 7$, against isomorphism class counts $2, 4, 11, 34, 156, 220, 226, 30, 2$. At orders ≤ 4 the two agree and the common value $2, 4, 11$ is* [9]*; at* *those orders the comparability half of the search is vacuous, so what is verified there is pure isomorph rejection. At order* 7 *the class count is* $220 = 1044 - 824$*, the difference of* [9] *and* [8]*. At orders* $8, 9, 10$ *the survivor lists reduce to* $226, 30, 2$ *classes — the ladder of Theorem 3.9.*

6. Sharpness, and the controls

Exhaustive computations are only as good as the statement they exhaust, so every definition and every value above is accompanied by refutations of its neighbours. We record them because they are the reason to believe the numbers, not because they are difficult.

Proposition 6.1. *Each clause of Definition 2.1 is load-bearing, and the values of Section 3 are sharp in both directions.*

1. *Dropping totality lets the empty relation through, and every graph would score* n*.*
2. *Dropping transitivity lets the cyclic orientation of* C_5 *through, and* C_5 *would score* 5 *instead of* 4*; that orientation is indeed not transitive.*
3. *Replacing the ordered triples* $i < j < l$ *of Proposition 2.3 by unordered ones — the orientation-free reading — already fails on* C_4*, whose value is* 4*.*
4. *Disconnected induced subgraphs are needed, not merely allowed: the only* 5*-subset of the order-*6 *minimiser avoiding its isolated vertex admits no valid enumeration, so every comparability* 5*-set there is disconnected, and* $\mu(6) = 5$ *depends on it.*
5. *Looplessness is not decoration: the all-adjacent matrix is refused as a graph, and has* $\mathrm{ugp} = 0$*, since a loop forces* $F(x, x)$ *by totality and forbids it by asymmetry.*
6. *Too strong, refuted:* $\mu(5) \neq 5$ *and* $\mu(7) \neq 6$*, each by an exhibited graph;* $\mu(n+1) = \mu(n) + 1$ *fails at* $n = 4$*; the source's bound* $\mathrm{ugp} \geq \mathrm{diam} + 1$ *is not tight at* C_5 *(bound* 3*, value* 4*); and* ugp *is not monotone in the edge set — deleting one edge from the Petersen graph raises* ugp *from* 7 *to at least* 8*.*
7. *Too weak, refuted:* $\mathrm{ugp} \leq n - 1$ *for non-complete graphs fails at* C_6*;* $\mathrm{ugp}(\textit{Petersen}) \leq 6$ *fails;* $\mathrm{ugp} \leq 6$ *at order* 10 *fails at* C_{10}*; and* $\mu(6) \leq 4$ *fails.*

Proposition 6.2. *Controls for the transfer lemma and the ninth value. Injectivity in Proposition 2.4 cannot be dropped: a constant map pulls* C_5 *back to the empty graph on five vertices, whose* ugp *is* 5*, against* $\mathrm{ugp}(C_5) = 4$*. The order hypothesis cannot be dropped:* $6 \leq \mathrm{ugp}$ *fails at order* 7*. An isolated vertex raises* ugp *by neither* 0 *nor* 2*. As a positive control,* every 8*-vertex induced subgraph of* M_9 *and every* 9*-vertex induced subgraph of* M_{10} *has* $\mathrm{ugp} = 6$ *exactly, for every injection and with no new evaluation; this is the sharpest available cross-check between consecutive values, and a wrong witness would fail it.*

Proposition 6.3. *Controls for the searches over graphs. The extension theorem is sharp: at* $m = 8$ *the identical search over the identical* 1024 *extensions of Theorem 3.11 returns "false", so the value* 7 *there is a property of the graphs and not of a search that succeeds at every* m*. The search is not vacuously affirmative: asked for* $\mu(7) \geq 6$*, which is false, it says "false". The prefix code and the list of allowed order-*10 *codes both carry weight — dropping either end of the order-*4 *list makes the same search fail — and a wrong permutation witness is rejected. Switching the isomorph-rejection prune off at order* 5*, where the comparability prune cannot fire at all, isolates its effect exactly:* 1100 *nodes against* 220*.*

Proposition 6.4. *Controls for the values of Sections 3 and 4. For* $\mu(12)$ *and* $\mu(15)$ *both neighbours are refuted. The pruned subset search is not vacuously refuting: at one order less it returns "false" on the same graphs, finding the very set it cannot refute, which is the witness of the lower half; and it declines to refute* K_6*, which is a comparability graph on all six vertices. It does refute the smallest graph that deserves it:* C_5 *has no comparability* 5*-subset and does have a* 4*-subset. The exhibited enumerations are load-bearing — for* $K_3 \,\square\, K_4$ *and for* $T(6)$ *the seven vertices in their natural order do* not *pass the triple condition. Finally, the hereditary pruning does not reduce the number of* *comparability tests at these sizes (*877 *against* 495 *for* $K_3 \square K_4$*,* 7008 *against* 6435 *for* $T(6)$*); it is a win only when the binomial coefficient is large, as at* $K(7,2)$*.*

Proposition 6.5. *Controls for the order-*16 *search of Theorem 3.17. The same search at the same budget, one order down over a* 14*-vertex induced subgraph of* $T(6)$*, returns "false", because that prefix does extend — to* $T(6)$ *itself — and the accepted extension is exhibited rather than asserted. At budget* 9 *instead of* 8 *the order-*16 *search returns "false", so* 8 *is not an artefact. The search does not succeed at its own root: this is not a computation but a consequence of* $\operatorname{ugp}(T(6)) = 7 < 8$ *and the soundness of the test. The hypothesis is satisfiable, since* $T(6)$ *plus an isolated vertex is such a graph; and over that sub-family the bound* 8 *is attained exactly.*

Proposition 6.6. *Controls for the implication-class machinery. The relation* Γ *is not itself transitive: on the paw (a triangle* $0,1,2$ *with a pendant vertex* 3 *at* 0*) one has* $(0,1)\,\Gamma\,(0,3)$ *and* $(0,3)\,\Gamma\,(0,2)$ *but not* $(0,1)\,\Gamma\,(0,2)$*, because* 1 *and* 2 are *adjacent — so the implication classes are classes of the transitive closure, and the closure does real work. The refutation test fires on* C_5 *and stays silent on* K_6*, on* $K_{3,3}$*, on* C_8 *and on the nine-vertex witness inside* $\overline{C_{10}}$*; the acceptance test accepts those and refuses* C_5*. Both neighbours of each of* $\operatorname{ugp}(\mathit{Petersen}) = 7$ *and* $\operatorname{ugp}(K(6,2)) = 9$ *are refuted.*

7. A validation against known counts

Since $\operatorname{ugp}(G) = n$ if and only if G is a comparability graph [1, Corollary 4.5], the number of graphs of order n with $\operatorname{ugp} = n$ must be the number of comparability graphs of order n, which is [8]. Recomputing that count from scratch — counting labelled graphs and converting by Burnside's lemma over the cycle types of S_n — gives

$$11, \quad 33, \quad 144, \quad 824, \quad 6793 \qquad (n = 4, \ldots, 8),$$

agreeing with [8] at every term. The full distributions of ugp over isomorphism classes have row sums [9]: at $n = 5$, one class with $\operatorname{ugp} = 4$ and 33 with $\operatorname{ugp} = 5$ (34 in total); at $n = 6$, 12 with $\operatorname{ugp} = 5$ and 144 with $\operatorname{ugp} = 6$ (156); at $n = 7$, 2 with $\operatorname{ugp} = 5$, 218 with $\operatorname{ugp} = 6$ and 824 with $\operatorname{ugp} = 7$ (1044). The corresponding labelled counts of comparability graphs,

$$1,\ 2,\ 8,\ 64,\ 1012,\ 30\,164,\ 1\,527\,416,\ 119\,369\,968 \qquad (n = 1, \ldots, 8),$$

are not in the OEIS. These counts are computations, not machine-checked theorems: verifying them inside the proof assistant would require the upper half of ugp for every non-comparability graph and so inherits the factorial obstruction at $n = 7$. Three independent implementations agree wherever they overlap, two of them on structurally different criteria, cross-validated on every induced subgraph of every graph on at most 7 vertices — 268 million comparability tests, no disagreement.

8. Verification

Every statement in Sections 2–6 carrying a theorem, proposition or lemma heading has been formalised and machine-checked in Lean 4 against Mathlib, with no unproved assumptions beyond the ambient logical axioms of that system; the definitions, the bridge of Proposition 2.3, the encoding of graphs by natural numbers, the soundness of every search, and the whole of Section 5.1 are checked by the kernel without any appeal to compiled evaluation, and compiled evaluation enters only where a specific finite search is run — one such evaluation per sweep, per exhaustion and per attaining graph, with the axiom list of each theorem recorded. Two statements are deliberately conditional and are marked as such in the text: the part of Theorem 3.17 covering the remaining 30 720 extensions of $T(6)$, which appears everywhere as an explicit hypothesis, and the counting statements of Section 7, which are computations reproduced by three independent programs rather than theorems. Repository reference to be added at submission.

9. Open questions

Question 9.1. What is $\mu(16)$? Proposition 3.16 brackets it in $\{7, 8\}$. The upper end would follow from showing that the hereditary class $\{G : \mathrm{ugp}(G) \leq 7\}$ dies at order 16; the lower end needs a single graph on sixteen vertices with ugp ≤ 7, and Remark 3.18 explains why local search has not produced one.

Question 9.2. What is the growth rate of μ? The flat stretches lengthen — two orders at 4, two at 5, three at 6, at least five at 7 — which is consistent with $\Theta(\log n)$ and inconsistent with any $\lfloor n/2 \rfloor$-type formula (Theorem 3.7), but nothing here proves a rate.

Question 9.3. Why are the minimisers line graphs (Theorem 4.7), and how far does that go? Equivalently: what is the largest induced comparability subgraph of $L(K_n)$? We do not know whether $T(7) = L(K_7)$ or the rook's graphs beyond order 16 continue the pattern.

Question 9.4. How many isomorphism classes attain $\mu(12)$? Theorem 3.15 gives at least two; order 10 has exactly two (Theorem 3.10), and nothing here classifies order 12.

Remark 9.5. Finally, a caveat on the literature: the searches behind the claims of novelty here cover the arXiv full-text corpus, the forward citations of [1] (of which there are none) and the OEIS, but not the off-arXiv algorithms literature — Golumbic's book [2], the survey of Brandstädt, Le and Spinrad, and the graph-class databases were not consulted — where the maximum induced subgraph problem for a hereditary class is classical and where a table of largest induced comparability subgraphs, should one exist, is most likely to be found.

References

[1] Ullas Chandran S. V., Gabriele Di Stefano, Grahame Erskine, Haritha S., Elias John Thomas and James Tuite, *The general position number of digraphs*, arXiv:2604.15909v1 [math.CO] (17 April 2026), Section 4.

[2] M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press, New York, 1980; chapter 5 for the Γ-implication-class criterion.

[3] M. C. Golumbic, *Recognizing comparability graphs in SETL*, SETL Newsletter 163, 19 March 1976.

[4] M. C. Golumbic, *Comparability graphs and a new matroid*, J. Combin. Theory Ser. B **22** (1977), 68–90.

[5] T. Gallai, *Transitiv orientierbare Graphen*, Acta Math. Acad. Sci. Hungar. **18** (1967), 25–66.

[6] R. C. Read, *Every one a winner, or how to avoid isomorphism search when cataloguing combinatorial configurations*, in: Algorithmic Aspects of Combinatorics, Ann. Discrete Math. **2**, North-Holland, Amsterdam, 1978, pp. 107–120.

[7] I. A. Faradžev, *Constructive enumeration of combinatorial objects*, in: Problèmes combinatoires et théorie des graphes (Orsay, 9–13 juillet 1976), Colloq. Internat. CNRS No. 260, CNRS, Paris, 1978, pp. 131–135.

[8] N. J. A. Sloane, *The On-Line Encyclopedia of Integer Sequences*, sequence A123416: number of comparability graphs on n nodes.

[9] N. J. A. Sloane, *The On-Line Encyclopedia of Integer Sequences*, sequence A000088: number of graphs on n unlabelled nodes.