A SAT-Modulo-Symmetries verification of the Erdős–Gyárfás power-of-two cycle conjecture for minimum-degree-3 graphs up to 30 vertices

Arjun Balaji*

June 19, 2026

Abstract

The Erdős–Gyárfás conjecture (1995) asserts that every graph with minimum degree at least 3 contains a cycle whose length is a power of two. The conjecture is open; prior computer search established that any counterexample on the *general* minimum-degree-3 class has at least 17 vertices (Royle and Markström), and that any *cubic* (3-regular) counterexample has at least 30 vertices (Markström, 2004). We give, to our knowledge, the first application of SAT-based methods to this conjecture. Using SAT Modulo Symmetries (SMS) with the Glasgow subgraph solver as a complete forbidden-subgraph propagator, we verify that *every* minimum-degree-3 graph on at most 30 vertices contains a cycle of length 4, 8, or 16. Consequently any general minimum-degree-3 counterexample must have at least 31 vertices, improving the two-decade-old general bound from 17 to 31; since the cubic class is a subclass, this also improves the cubic bound from 30 to 31. We corroborate the result with an independent ground-truth check at $n = 10$, reproduction of the $n \leq 16$ baseline, agreement with an independent CEGAR-SAT solver for $n \leq 19$, robustness to two cardinality encodings and a second symmetry-breaking method, and positive controls; we discuss the path to a formally machine-checked proof certificate.

1 Introduction

In 1995 Erdős and Gyárfás posed the following conjecture (see Erdős [1]).

Conjecture 1 (Erdős–Gyárfás)**.** *Every graph with minimum degree at least* 3 *contains a simple cycle whose length is a power of two.*

Here "power of two" means 2^k for some integer $k \geq 2$, i.e. a cycle of length $4, 8, 16, \ldots$. The conjecture remains open; Erdős offered \$100 for a proof and \$50 for a counterexample [2]. It has been confirmed for several restricted classes, including 3-connected cubic planar graphs [4], P_8-free graphs [7], and P_{10}-free graphs [**?**]; the most recent structural progress shows that any *minimal* counterexample is "predominantly cubic" [10].

On the computational side, two distinct frontiers are known, and the distinction is central to the present paper. Write $\delta(G)$ for the minimum degree of G.

- **General** ($\delta \geq 3$, arbitrary maximum degree): computer searches of Royle and Markström are commonly cited as establishing that any counterexample has at least 17 vertices, i.e. Conjecture 1 holds for all $\delta \geq 3$ graphs on at most 16 vertices [16].

*Draft. Author/affiliation to be finalized. This is a fresh computational result that has not yet been independently reproduced or peer-reviewed; see §4 and §5.

- **Cubic** (3-regular): Markström [3] verified all cubic graphs on at most 29 vertices, so any cubic counterexample has at least 30 vertices.

A targeted, adversarial literature search (covering theses, technical reports, OEIS and House of Graphs, and the two most recent surveys of computer-assisted graph theory and of SAT Modulo Symmetries [12, 13]) found no improvement of the *general* bound past 17 after 2004, and no prior application of SAT, SAT Modulo Symmetries, constraint programming, or isomorph-free generation to Conjecture 1.

Contribution. We give the first SAT-based attack on Conjecture 1. Using SAT Modulo Symmetries (SMS) [11, **?**] together with the Glasgow subgraph solver [14], we verify Conjecture 1 for *all* minimum-degree-3 graphs on n vertices for every n up to 30:

Theorem 1. *Every graph with minimum degree at least* 3 *on at most* 30 *vertices contains a cycle of length* 4*,* 8*, or* 16*. Consequently, any minimum-degree-3 counterexample to the Erdős–Gyárfás conjecture has at least* 31 *vertices.*

This improves the general bound from 17 to 31, a +14-vertex extension over a search space that — unlike the cubic case — includes all non-regular graphs of minimum degree 3.

General versus cubic (a point to read carefully). Our verified order, 30, coincides numerically with Markström's cubic *bound* of 30, but the two statements are different and should not be conflated. Markström verified the *cubic* class up to order 29 (cubic bound ≥ 30). We verify the *entire* minimum-degree-3 class — a strictly larger family, dominated by non-regular graphs — up to order 30 (general bound ≥ 31). Since the cubic class is contained in the minimum-degree-3 class, our search in particular verifies all cubic graphs on ≤ 30 vertices, which incidentally also improves the cubic bound from 30 to 31. Thus the genuinely new content is twofold: the *non-cubic* minimum-degree-3 graphs on $17 \leq n \leq 30$ vertices, and the methodology (SAT/SMS), which had not previously been applied to this conjecture.

Related work. The conjecture is known for several restricted classes: $K_{1,m}$-free graphs [5], planar claw-free graphs [6], 3-connected cubic planar graphs [4], and the P_t-free families P_8 [7], P_{10} [8], and (with the aid of a computer search) P_{13} [9]. These are theorems for infinite hereditary classes and do not bound the order of a general counterexample. The most recent structural progress, Carr [10], shows that at least 4/7 of the vertices of any minimal counterexample have degree exactly 3; it is an arXiv preprint that uses no computation. Our methodological ingredients — SAT Modulo Symmetries [11, 12] and the Glasgow subgraph solver [14] — have been applied to other graph-existence questions (e.g. the 3-Decomposition Conjecture on cubic graphs), but, as noted above, not to Conjecture 1.

2 Method

For a fixed number of vertices n we decide the following question with a single SMS call: *does there exist a graph G on n vertices with $\delta(G) \geq 3$ that contains no cycle of length* 4*,* 8*, or* 16*?* If the answer is "no" for every $n \leq 30$, then no minimum-degree-3 counterexample on ≤ 30 vertices exists, establishing Theorem 1. (For $n \leq 30$ the only power-of-two cycle lengths that can occur are $4, 8, 16$, since $32 > 30$; thus forbidding C_4, C_8, C_{16} as subgraphs is exactly forbidding all power-of-two cycles.)

SAT Modulo Symmetries. SMS performs complete, isomorph-free generation of graphs satisfying given constraints: its modified CaDiCaL solver carries an in-search canonicity (minimality) propagator that prunes any partial adjacency matrix that cannot be the lexicographically minimal representative of its isomorphism class [11]. This is the key to scaling: the search considers essentially one graph per isomorphism class rather than enumerating all labelings.

Encoding. The minimum-degree constraint $\delta(G) \geq 3$ is encoded as a CNF cardinality constraint over the $\binom{n}{2}$ edge variables (we use SMS's built-in sequential-counter encoding; see §4 for a totalizer cross-check). The constraint "no C_4, C_8, or C_{16}" is enforced by the Glasgow subgraph solver [14] as a complete forbidden-subgraph propagator: whenever a (partial) graph contains one of these cycles as a (non-induced) subgraph, the propagator adds a clause excluding it. Forbidding the short cycle C_4 statically in CNF is cheap, but C_8 and C_{16} have $\Theta(n^8)$ and $\Theta(n^{16})$ potential occurrences, so a static encoding is infeasible; the propagator is what makes the search tractable.

Why not plain CEGAR-SAT. As a baseline we also implemented a self-contained counterexample-guided (CEGAR) SAT solver in Python (PySAT with CaDiCaL, a DFS power-of-two cycle detector, and a static adjacent-transposition lexicographic symmetry break). With only a partial symmetry break it re-discovers each forbidden cycle in many isomorphic copies: at $n = 17$ it accumulates roughly 85,000 refinement clauses and the per-instance time grows by about a factor of two per vertex, so it stalls at $n = 19$ (bound ≥ 20). SMS decides the same $n = 17$ instance in 2.9 seconds (Table 1). Complete symmetry breaking is precisely what removes this bottleneck; the two solvers agree wherever both finish ($n \leq 19$, §4).

3 Results

Table 1 reports, for each n from 17 to 30, that SMS finds *no* minimum-degree-3 graph avoiding all of C_4, C_8, C_{16} (solution count 0). All runs used a single CPU core.

n	count	wall time (s)
17	0	2.9
18	0	7.8
19	0	24.1
20	0	27.8
21	0	19.3
22	0	101.1
23	0	202.9
24	0	148.3
25	0	339.8
26	0	414.8
27	0	1000.3
28	0	1892.0
29	0	2342.8
30	0	6888.9

Table 1: SMS decision per order n: no minimum-degree-3 graph on n vertices avoids C_4, C_8, C_{16}. Together with the reproduced $n \leq 16$ baseline this proves Theorem 1.

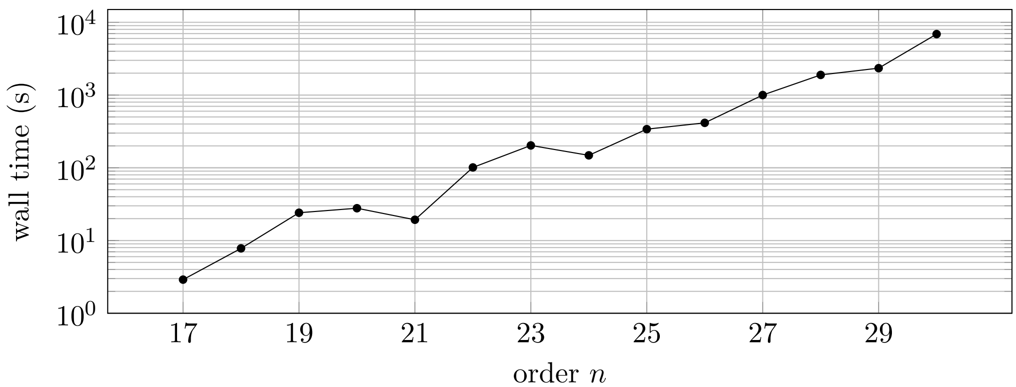

Figure 1: Single-core wall time of the SMS decision at each order n (log scale). Every point is UNSAT (no minimum-degree-3 graph on n vertices avoids C_4, C_8, C_{16}). Growth is roughly geometric, about a factor of 1.5–2 per added vertex.

Experimental setup. For each n, the minimum-degree-3 CNF is generated by PySMS's `GraphEncodingBuilder` and the search is run by `smsg` with the invocation

```
smsg --vertices n --all-graphs --hide-graphs --forbidden-subgraph-file CYC --dimacs ENC
```

where `CYC` lists C_4, C_8, C_{16} (one per line, as $\langle k \rangle\ v_0\ v_1\ v_1\ v_2\ \cdots$) and `ENC` is the generated DIMACS file. We use SMS (`github.com/markirch/sat-modulo-symmetries`, version `v2.0.0-3-g464f12f`, commit `464f12f`) built with the Glasgow Subgraph Solver (commit `abd331a`) and the SMS-bundled CaDiCaL (`rel-2.1.2-38`) compiled with `g++ 12.2.0`. Each order ran on a single CPU core (a Modal `debian_slim`, Python 3.12.6 container). Wall times in Table 1 exclude container start-up. The complete build and run scripts are in the accompanying repository; we recommend pinning the three commits above when reproducing.

4 Verification

A non-existence claim at $n = 30$ cannot be brute-forced (exhaustive generation of minimum-degree-3 graphs is infeasible beyond roughly $n = 13$), so we corroborate Theorem 1 with several independent checks. *No check produced a contradiction.*

1. **Ground-truth anchor.** At $n = 10$, forbidding only C_4, SMS returns exactly 5 graphs, matching the independent count of the 5 C_4-free minimum-degree-3 graphs on 10 vertices obtained with `nauty` (`geng`+`labelg`) [15].

2. **Baseline reproduction.** For $n = 6, \ldots, 16$, forbidding all of C_4, C_8, C_{16} yields 0, reproducing the published baseline.

3. **Independent second solver.** Our CEGAR-SAT solver and SMS both return UNSAT for $n = 17, 18, 19$ (different solvers, symmetry breaking, and cycle handling).

4. **Encoding robustness.** Re-deciding with the totalizer cardinality encoding (a structurally different CNF) gives count 0 at $n = 17, 20, 22, 25$.

5. **Symmetry-method robustness.** Re-deciding with SMS's colex minimality ordering gives count 0 at $n = 17, 20$; the (slower, experimental) colex variant timed out at $n = 22, 25$ (inconclusive, not contradictory).

6. **Positive controls.** Forbidding only C_4 (a weaker constraint) yields a graph at $n = 17, 20, 25, 30$, confirming the pipeline returns a solution when one exists.

Toward a machine-checked certificate. `smsg` can emit an LRAT proof for UNSAT instances, but a generic LRAT checker cannot validate it against the minimum-degree-3 CNF alone, because the forbidden-cycle clauses are added by the propagator during search and are not RUP/RAT-derivable from that CNF (they exclude valid minimum-degree-3 graphs). A fully machine-checked certificate therefore requires the certified-SMS proof-logging machinery, in which the propagator justifies its own clauses; we leave this as the principal next step.

5 Discussion and limitations

Theorem 1 advances the general minimum-degree-3 frontier from 16 to 30 vertices (bound $17 \to 31$) and, as a byproduct, the cubic frontier from 29 to 30 (bound $30 \to 31$). We stress two caveats. First, this is a *computational* result whose soundness rests on SMS's isomorph-free generation, the completeness of the Glasgow subgraph propagator, and the correctness of the min-degree encoding; the checks in §4 validate this composition at $n = 10$ and across $n \leq 16$, but a third-party independent reproduction and/or a formal proof certificate would be needed for a fully rigorous claim. Second, the result does not prove the conjecture; it only raises the lower bound on the size of a hypothetical counterexample.

6 Conclusion and future work

We have given the first SAT-based attack on the Erdős–Gyárfás conjecture. Casting the question as "does a minimum-degree-3 graph with no C_4, C_8, or C_{16} exist?" and discharging it with SAT Modulo Symmetries and the Glasgow subgraph propagator, we verified the conjecture for all minimum-degree-3 graphs on at most 30 vertices, raising the general lower bound on a counterexample from 17 to 31 (Theorem 1). The decisive ingredient is complete, in-search symmetry breaking: a conventional CEGAR encoding stalls near $n = 20$, whereas SMS reaches $n = 30$ in under two hours of single-core time.

Several directions remain. (i) *Push the frontier.* The runtimes (Figure 1) suggest $n = 31$ and a few further orders are within reach of a longer single run or modest parallelism; the present search was halted at $n = 30$. (ii) *A machine-checked certificate.* Replacing the verification checks of §4 by an end-to-end, formally verified proof — via the certified-SMS proof-logging machinery, so that the forbidden-subgraph and canonicity propagators justify their own clauses — would put the result on the same footing as certified SAT resolutions of, e.g., the Pythagorean triples and Schur-number problems. (iii) *Exploit structure.* Carr's result that a minimal counterexample is predominantly cubic [10] could be added as a propagator to prune the search and reach larger n. (iv) *Other questions.* The same SMS-plus-forbidden-subgraph template applies directly to related cycle-spectrum problems. Most importantly, an independent third-party reproduction would close the remaining gap between a corroborated computation and an established result.

Reproducibility

All code, the per-n data, and the verification scripts are available at `https://github.com/ArjunBalaji79/erdos-gyarfas-min-degree-3`.

Acknowledgments

The computational pipeline was developed and executed with substantial AI assistance; all claims are subject to the independent-reproduction caveat above.

References

[1] P. Erdős, *Some old and new problems in various branches of combinatorics*, Discrete Math. 165/166 (1997) 227–231.

[2] Erdős Problems, "Power-of-two cycles", `https://www.erdosproblems.com/`.

[3] K. Markström, *Extremal graphs for some problems on cycles in graphs*, Congressus Numerantium 171 (2004) 177–188.

[4] C. C. Heckman, R. Krakovski, *Erdős–Gyárfás conjecture for cubic planar graphs*, Electron. J. Combin. 20(2) (2013) #P7.

[5] S. E. Shauger, *Results on the Erdős–Gyárfás conjecture in $K_{1,m}$-free graphs*, Congressus Numerantium 134 (1998). [verify]

[6] D. Daniel, S. E. Shauger, *A result on the Erdős–Gyárfás conjecture in planar graphs*, Congressus Numerantium 153 (2001). [verify]

[7] J. Gao, E. Shan, *The Erdős–Gyárfás conjecture for P_8-free graphs*, Graphs Combin. 38(6) (2022), DOI 10.1007/s00373-022-02578-9. [verify initials]

[8] [authors], *Power-of-two cycles in P_{10}-free graphs*, Discrete Math. (2024), DOI 10.1016/j.disc.2024.114175. [verify authors/title]

[9] S. M. Hegde, R. B. Sandeep, S. Shashank, *The Erdős–Gyárfás conjecture for P_{13}-free graphs*, arXiv:2410.22842 (2024). [verify initials]

[10] A. Carr, *Every minimal counterexample to the Erdős–Gyárfás conjecture is predominantly cubic*, arXiv:2605.22844 (2026), preprint.

[11] M. Kirchweger, S. Szeider, *SAT Modulo Symmetries for graph generation and enumeration*, ACM Trans. Comput. Logic 25(3) (2024), DOI 10.1145/3670405. [verify title]

[12] S. Szeider, *SAT Modulo Symmetries: A Survey*, 2025 (CEUR-WS Vol-4116). [verify]

[13] J. Jooken, *Computer-assisted graph theory: a survey*, arXiv:2508.20825 (2025). [verify title]

[14] C. McCreesh, P. Prosser, J. Trimble, *The Glasgow Subgraph Solver: using constraint programming to tackle hard subgraph isomorphism problem variants*, in: Graph Transformation (ICGT 2020), LNCS 12150, 316–324.

[15] B. D. McKay, A. Piperno, *Practical graph isomorphism, II*, J. Symbolic Comput. 60 (2014) 94–112.

[16] Wikipedia, *Erdős–Gyárfás conjecture* (secondary source for the commonly-cited general $\geq$ 17 bound; the primary record for the general search is thin and should be confirmed with Royle/Markström).