

Exact values $F(6) = 14$ and $F(7) = 17$ for a function of Gowers and Long

Lluis Eriksson

21 October 2026

Abstract

For triples $a, b \in \mathbb{Z}^3$ write $a <_2 b$ if $a_i < b_i$ for at least two coordinates i . Let $F(n)$ be the maximal length of a sequence of triples in $\{1, \dots, n\}^3$ with $a^i <_2 a^j$ whenever $i < j$. Gowers and Long [1] conjecture that $F(n) \leq n^{3/2}$ for all n (their Conjecture 1.8) and tabulate $F(n)$ for $n \leq 5$. Using SAT solving with a cube-and-conquer split I find $F(6) = 14$ and $F(7) = 17$. No value violates $n^{3/2}$ ($14 < 14.70$, $17 < 18.52$). The lower bounds are explicit sequences (appendix) checked by independent programs and by Lean 4. The upper bounds come from unsatisfiability. For $F(6) \leq 14$ the evidence is an independent branch-and-bound program, two SAT solvers and a DRUP proof checked by a small checker. For $F(7) \leq 17$ the cubes are cross-checked by two solvers, and 407 of 425 DRUP proofs (plus 46 counting arguments) were checked. This is *not* a formally verified proof, and it neither proves nor refutes the conjecture; it is a small computational datum. See the “Source disclosure” section.

1 The problem

Following Gowers and Long [1], who follow Loh [2], call a triple $a = (a_1, a_2, a_3)$ *2-less* than b , written $a <_2 b$, if $a_i < b_i$ for at least two values of i . A sequence $a^1, \dots, a^m$ of triples is *2-increasing* if $a^r <_2 a^s$ whenever $r < s$. With $[n] = \{1, \dots, n\}$,

$$F(n) = \max\{m : \text{there is a 2-increasing sequence } a^1, \dots, a^m \in [n]^3\}.$$

Gowers and Long prove $F(n) \geq n^{3/2}$ when n is a perfect square and $F(n) \leq n^2 / \exp(\Omega(\log^* n))$ (using Loh’s triangle-removal argument), and state:

Conjecture 1.8 [1]. $F(n) \leq n^{3/2}$ for all n .

They report that a computer search for small examples with exponent above $3/2$ found none, and their Lemma 1.7 shows that a single n with $F(n) = n^\alpha$ would give $F(m) \geq m^\beta$ for every $\beta < \alpha$ and all large m . The formal-conjectures repository of Google DeepMind [3] contains the statement (file `Arxiv/1609.08688/sIncreasingrTuples.lean`, theorem `maximalLength_le_strong`, marked open, read on 21 Oct. 2026), where $F(4) = 8$ appears as a `sorry`. I use the same definition of F .

Since $a <_2 b$ and $b <_2 a$ cannot both hold (two coordinates in each direction among three would overlap), a set S of triples can be ordered into a 2-increasing sequence iff any two members are 2-comparable and the resulting tournament is transitive, and then the order is unique. Distinct triples that agree in two coordinates are not comparable, so each of the three projections to pairs of coordinates is injective on S .

2 Results

Theorem 1. $F(6) = 14$ and $F(7) = 17$.

n	$F(n)$	$n^{3/2}$	Source
1	1	1.00	Gowers–Long, Table 1 [1]; reproduced here
2	2	2.83	same
3	4	5.20	same
4	8	8.00	same (equality with $n^{3/2}$)
5	10	11.18	same; reproduced by brute force, CP-SAT and SAT
6	14	14.70	new here , not in [1]; Section 3
7	17	18.52	new here ; Section 3

Table 1: Values of $F(n)$. Those for $n \leq 5$ are from [1] and were reproduced; I found no earlier value for $n = 6, 7$ (limited search, see Source disclosure).

Table 1 gives $F(n)$ for $n \leq 7$. So $F(n) \leq n^{3/2}$ holds for $n \leq 7$, with equality only at $n = 4$. This is modest evidence for Conjecture 1.8 and no proof of it. $F(7) = 17$ rules out $F(7) = 19$, the smallest value that would have refuted the conjecture at $n = 7$.

n	Claim	Evidence	Strength
6	sequence of length 14	Python checker; Lean <code>native_decide</code> and <code>decide +kernel</code>	rigorous
7	sequence of length 17	same	rigorous
6	no sequence of length 15	Go branch-and-bound (no SAT); CP-SAT; CaDiCaL; Glucose DRUP proof (452 559 lemmas) checked by Go RUP checker	computational
7	no sequence of length 18	CaDiCaL and MapleChrono on 160 cubes; 407 of 425 DRUP proofs checked	partly proof-checked

Table 2: Status of the claims.

3 Method and verification

Encoding. Variables x_t ($t \in [n]^3$) say that t belongs to the set. Clauses: $\neg x_s \vee \neg x_t$ for every pair that is not 2-comparable, and a sequential-counter encoding of $\sum x_t \geq m$. A satisfying assignment gives a pairwise comparable set, which is a sequence iff it has no cyclic triple $a <_2 b <_2 c <_2 a$. Cyclic triples are excluded lazily: after each model, violated cyclic triples give clauses $\neg x_a \vee \neg x_b \vee \neg x_c$. These are sound (a cyclic triple cannot occur in a 2-increasing sequence) but are not consequences of the pair clauses alone; where they were used they are part of the CNF files that the proofs refer to. For $n = 6$ none were needed; for $n = 7$ only a few cubes needed any (at most 3296 in one cube).

Symmetry and cubes. The relation $<_2$ is invariant under permuting the three coordinates (S_3). Hence the first element of a sequence may be assumed to have non-decreasing coordinates (84 choices for $n = 7$). For $F(7) \geq 18$ the instance is split by the first element. For the first element $(1, 1, 1)$, which S_3 fixes, the second element may also be taken with non-decreasing coordinates (77 subcubes). A second symmetry, $a \mapsto (n+1-a_1, n+1-a_2, n+1-a_3)$ together with reversal of the sequence, exists but is *not* used. The method was validated on $n = 5$ ($m = 10$ satisfiable; $m = 11$ unsatisfiable in all cubes) and $n = 6$ ($m = 14$ satisfiable; $m = 15$ unsatisfiable in all cubes), reproducing CP-SAT and brute force.

Lower bounds. Sequences of length 14 in [6]³ and 17 in [7]³ were found by CP-SAT (OR-Tools) and are listed in the appendix. They are checked by a direct program in Python and in Lean 4 (core, no Mathlib; file `Gowers.lean`, Lean 4.35.0) both by `native_decide` and by `decide +kernel`, which does not trust the compiler. The file also checks that short false witnesses are rejected.

$F(n)$ for $n \leq 6$ by branch and bound, without SAT. A separate program in Go (`bnb.go`) enumerates sequences by choosing the first element (non-decreasing coordinates) and then restricting the candidates to the elements 2-greater than the chosen one, with a bound given by the number of occupied lines in the three pair projections. It returns $F(1), \dots, F(6) = 1, 2, 4, 8, 10, 14$ ($4.7 \cdot 10^8$ nodes and 21 minutes for $n = 6$). It is too slow for $n = 7$ (the node count grew by a factor of about 290 from $n = 5$ to $n = 6$), so it gives no independent check of $F(7) \leq 17$.

Upper bound for $n = 6$. Besides the Go program, CP-SAT proved optimality of 14, and CaDiCaL found $m = 15$ unsatisfiable. Glucose 4 (through PySAT) produced a DRUP proof of 452 559 lemmas for the CNF of $m = 15$ (no cyclic-triple clauses). A forward RUP checker that I wrote in Go (a port of the algorithm of `rupcheck2.py`, two watched literals, deletions ignored) accepts it, reaching a level-0 conflict after 233 894 lemmas (14 minutes). The Python checker `rupcheck2.py`, which I had written for an earlier note [4], did *not* finish on this proof in 38 minutes (stopped by a timeout), so it gives no result there. The two checkers agree on a 2905-lemma proof (both accept) and on five false proofs for the same CNF (both reject): the empty clause alone; the proof with its first half removed; the proof with a false lemma “1 0” prepended; the proof truncated by three lines; and (Go only) the genuine proof of another cube checked against this CNF.

Upper bound for $n = 7$. The 160 cubes (83 first elements other than $(1, 1, 1)$, and 77 subcubes of $(1, 1, 1)$) were solved with CaDiCaL (PySAT, incremental, lazy cyclic triples) and, separately, with MapleChrono; the results agree on all 160 (all unsatisfiable). For DRUP proofs, the first-element cube $(1, 1, 2)$ was too hard for Glucose within 15 minutes and was split by *all* possible second elements (312 subcubes, no symmetry used); $(1, 1, 1)$ was split by the 77 sorted second elements. The proof partition has $82 + 312 + 77 = 471$ jobs: 46 are refuted by counting alone (the candidates plus fixed elements number fewer than 18), and 425 have Glucose DRUP proofs (about $7 \cdot 10^7$ lemmas, 2.7 GB). A bug in my first run (cubes with two fixed elements were encoded contradictorily and so “proved” trivially) was found; those proofs were discarded and recomputed, and the discarded output is kept in the supplement (`bogus/`). Of the 425 proofs, 407 were accepted by the Go RUP checker within the budget (limit 900s per proof, 40 minutes wall time on four processes), covering 18 928 069 lemmas. The other 18 proofs (the largest ones, with up to $9 \cdot 10^6$ lemmas) **were not proof-checked**; for them the evidence is only the agreement of CaDiCaL and MapleChrono on the coarser 160-cube partition and the unchecked Glucose answer.

What is not done. (i) There is no complete DRUP verification of $F(7) \leq 17$, no LRAT or drat-trim certificate, and the RUP checkers are not verified software. (ii) The cube-splitting code is a single implementation (`cube.py`, `proofcubes.py`); its correctness rests on the validation for $n \leq 6$ and on solver agreement, and the S_3 symmetry argument is mine and was not formalised. (iii) The independent branch-and-bound covers only $n \leq 6$. (iv) Lean checks the witnesses only; there is no Lean proof of the upper bounds, and no theorem linking the Lean function `good` to the definition in [3]. (v) $F(8)$ was not attempted. (vi) No human review.

4 Source disclosure

This is exploratory computational work carried out with AI assistance: the scripts, formulas, checkers, Lean file and this text were produced by Claude Sonnet 5.5 agents under the direction of the author. There has been no human refereeing and no complete formal verification. The literature search was limited: the Gowers–Long paper (read for the parts cited), the formal-conjectures file, and two web searches for later computations or an OEIS entry. I found none, but absence of evidence is not evidence of priority, and I make no claim of priority. The numerical claims are reproducible with the supplement, and the reader should rerun them rather than rely on this text.

References

- [1] W. T. Gowers and J. Long, *The length of an s -increasing sequence of r -tuples*, arXiv:1609.08688 (v2, 2 Nov. 2016); *Combin. Probab. Comput.* **30** (2021), 686–721.
- [2] P. Loh, *Directed paths: from Ramsey to Ruzsa and Szemerédi*, arXiv:1505.07312 (2015), as cited in [1]; I did not read it.
- [3] Google DeepMind, *formal-conjectures*, file `FormalConjectures/Arxiv/1609.08688/sIncreasingrTuples.lean`, <https://github.com/google-deepmind/formal-conjectures>.
- [4] L. Eriksson, *Exact values for a finite Ramsey-type problem from Green’s list (Problem 22)*, note of 21 Oct. 2026, unpublished.

A Witness sequences

Triples (a_1, a_2, a_3) in order; each is 2-less than all later ones. They are also in `witness_n6.txt` and `witness_n7.txt`.

$n = 6$, length 14.

(1,1,1) (1,2,2) (1,3,3) (2,1,4) (2,2,5)
(2,3,6) (3,4,1) (4,5,1) (3,6,2) (5,4,3)
(4,6,4) (6,4,5) (5,5,6) (6,6,6)

$n = 7$, length 17.

(1,1,1) (2,2,1) (3,2,2) (1,3,3) (2,4,3)
(3,3,4) (4,1,5) (4,2,6) (4,3,7) (5,5,1)
(6,6,1) (5,7,2) (6,7,3) (7,4,4) (7,5,5)
(7,6,6) (7,7,7)