The Exact $\mathbb{F}_2$-Rank of 2×3 by 3×4 Matrix Multiplication

Ian D'Ambrosio
Nth Research Collective
ian@nthresearch.org

11 August 2026

Abstract

Hopcroft and Kerr gave a 20-multiplication algorithm for multiplying a 2×3 matrix by a 3×4 matrix, and conjectured that 20 multiplications are necessary. Over $\mathbb{F}_2$, Wang recently proved the lower bound 19 with a machine-checkable substitution certificate, leaving the exact rank in $\{19, 20\}$. We prove that the rank is 20. Wang's checked first-factor capacities imply that every hypothetical 19-term decomposition has pairwise distinct first factors and at most one first factor of rank two. A checked CNF/DRAT/LRAT certificate and exhaustive enumeration reduce all such supports to 252 profiles in four symmetry orbits. Row restrictions then exclude the four orbits. Tight 12-term restrictions force output factors into a four-dimensional row space; 13-term restrictions have a unique relation whose quotient factors are all equal. These facts immediately exclude the all-rank-one profiles of types $(7, 7, 5)$ and $(7, 6, 6)$. In the remaining profile, with type $(6, 6, 6)$ and one rank-two factor, the three relations would need at least 18 support incidences but can contain at most 9. An explicit 20-term decomposition is checked against all 576 tensor coefficients. The profile proof, restriction certificate, upper witness, and replay commands are provided as exact artifacts.

Keywords. bilinear complexity; tensor rank; matrix multiplication; finite fields; computer-assisted proof

1 Introduction

Let $R_{\mathbb{F}}(\langle m, n, p\rangle)$ denote the tensor rank over $\mathbb{F}$ of multiplying an $m \times n$ matrix by an $n \times p$ matrix. Equivalently, it is the least number of scalar multiplications in a bilinear algorithm for this map. Exact ranks for small rectangular formats are basic finite problems in algebraic complexity, but even tiny cases can depend on the ground field and resist both general lower bounds and constructive search [1, 4].

For $\langle 2, 3, 4\rangle$, Hopcroft and Kerr gave an algorithm using 20 multiplications and conjectured that 20 are necessary [3]. Smirnov's later table distinguishes an exact upper bound 20 from an approximate bound 19 [4]. AlphaTensor also found an exact rank-20 decomposition over $\mathbb{F}_2$ [2]. Wang's automated finite-field framework recently supplied the first checked lower bound

$$R_{\mathbb{F}_2}(\langle 2, 3, 4\rangle) \geq 19,$$

while retaining 20 as the upper bound [5]. Thus only one exact case remained.

Theorem 1. *The bilinear complexity over $\mathbb{F}_2$ of multiplying a 2×3 matrix by a 3×4 matrix is*

$$R_{\mathbb{F}_2}(\langle 2, 3, 4\rangle) = 20.$$

The lower-bound proof has a certified finite front end and a short structural back end. Wang's substitution bounds reduce any hypothetical rank-19 decomposition to four orbits of possible first-factor supports. We then use the three nonzero row restrictions of $\mathbb{F}_2^2$. A restriction has flattening rank 12. With 12 surviving terms it is tight and rigid; with 13 terms it has exactly one relation. The resulting contradictions depend only on the orbit's row multiplicities, so no search over the remaining two factors is needed.

Contribution and scope. The result closes this exact $\mathbb{F}_2$ format and supplies a new restriction-based exclusion of rank 19. The argument uses the special three-direction geometry of $\mathbb{F}_2^2$ and does not claim a general rank formula for $\langle 2, 3, p\rangle$ or for other fields. Searches over flip graphs and fixed-profile SAT instances were useful diagnostics but play no role in the proof.

2 Tensor conventions and row restrictions

Put

$$A = \mathbb{F}_2^2, \qquad B = \mathbb{F}_2^3, \qquad C = \mathbb{F}_2^4.$$

Using the standard bases to identify coordinate spaces with their duals, the matrix-multiplication tensor is

$$M_{2,3,4} = \sum_{i=1}^{2}\sum_{j=1}^{3}\sum_{k=1}^{4} a_{ij} \otimes b_{jk} \otimes c_{ik} \in (A \otimes B) \otimes (B \otimes C) \otimes (A \otimes C). \tag{1}$$

Suppose, toward a contradiction, that

$$M_{2,3,4} = \sum_{s=1}^{19} U_s \otimes V_s \otimes W_s, \tag{2}$$

where every factor is nonzero. Regard U_s as a 2×3 binary matrix.

For $x \in A \setminus \{0\}$, restrict the left input to matrices xy^{T} with $y \in B$. Define

$$H_x = x \otimes C \subseteq A \otimes C, \qquad R_x = B \otimes H_x.$$

The restricted target is

$$T_x = \sum_{j=1}^{3}\sum_{k=1}^{4} e_j \otimes E_{jk} \otimes (x \otimes f_k). \tag{3}$$

Its flattening with $B \otimes C$ on one side has rank 12 and image exactly R_x. Moreover,

$$H_x \cap H_y = \{0\} \qquad (x \neq y,\ x, y \in A \setminus \{0\}). \tag{4}$$

If $U_s = a_s b_s^{\mathsf{T}}$ has rank one, then term s survives restriction x exactly when $a_s \cdot x = 1$, and its second flattening factor is

$$q_{s,x} = b_s \otimes W_s. \tag{5}$$

If U_* has rank two, it survives all three nonzero restrictions and

$$q_{*,x} = (x^{\mathsf{T}} U_*) \otimes W_*. \tag{6}$$

The first factor in (6) is nonzero because U_* has full row rank.

3 Two flattening lemmas

Lemma 2 (Tight restriction)**.** *If T_x is expressed by exactly 12 active terms from* (2)*, then the active V_s form a basis of $B \otimes C$ and the active $q_{s,x}$ form a basis of R_x. In particular, every active rank-one term satisfies $W_s \in H_x$.*

Proof. A rank-12 matrix written as a sum of 12 dyads forces both factor lists to have rank 12. Hence the $q_{s,x}$ span, and therefore belong to, the flattening image R_x. If $b_s \otimes W_s \in B \otimes H_x$ with $b_s \neq 0$, apply a functional taking b_s to 1 to obtain $W_s \in H_x$. ☐

Lemma 3 (One-excess restriction)**.** *Suppose that T_x is expressed by 13 active terms, indexed by I. Let*

$$L_x : \mathbb{F}_2^I \longrightarrow B \otimes C, \qquad e_s \longmapsto V_s.$$

Then L_x is surjective and $\ker L_x = \langle \lambda^x \rangle$ *is one-dimensional. If*

$$\pi_x : B \otimes (A \otimes C) \longrightarrow \big(B \otimes (A \otimes C)\big)/R_x$$

is the quotient map, there is a quotient vector z_x such that

$$\pi_x(q_{s,x}) = \lambda_s^x z_x \qquad (s \in I). \tag{7}$$

Proof. The active V_s span 12 dimensions because the restricted flattening has rank 12; their ambient space also has dimension 12. Thus L_x is surjective and has a one-dimensional kernel. Projecting the target outside R_x gives

$$\sum_{s\in I} V_s \otimes \pi_x(q_{s,x}) = 0.$$

Equivalently, the tensor $\sum_{s\in I} e_s \otimes \pi_x(q_{s,x})$ belongs to $\ker L_x \otimes \big((B\otimes(A\otimes C))/R_x\big)$. Since the first factor is spanned by λ^x, equation (7) follows. □

We will also use that over $\mathbb{F}_2$, equality of nonzero pure tensors $b \otimes h = b' \otimes h'$ forces $b = b'$ and $h = h'$. Over a general field the factors are proportional; $\mathbb{F}_2$ has only one nonzero scalar.

4 The certified first-factor profiles

Wang's substitution certificate supplies a checked residual lower bound $L(S)$ for each relevant subspace S of the first tensor factor. The bound applies after restricting the left input to the annihilator $S^\perp$, which kills exactly the terms whose first factor belongs to S. If $m(S)$ terms of a rank-19 decomposition have $U_s \in S$, substitution gives the capacity inequality

$$m(S) + L(S) \leq 19. \tag{8}$$

For every one-dimensional S, the checked residual bound is 18. Since a one-dimensional binary subspace has one nonzero vector, every nonzero first factor therefore occurs at most once.

The profile certificate expands Wang's orbit data to all first-factor subspaces and retains the 1,897 nontrivial inequalities (8). There are 21 nonzero 2×3 matrices of rank one and 42 of rank two. A 62,267-variable CNF asserts all capacities, cardinality 19, and at least two selected rank-two matrices. Its 124,027 clauses are unsatisfiable; both its DRAT and LRAT proofs check. Exhaustive enumeration of the remaining zero- or one-rank-two candidates checks 56,070 supports and retains exactly 252. Quotienting by $\mathrm{GL}(2, 2) \times \mathrm{GL}(3, 2)$, of order 1,008, gives the following complete list.

Proposition 4 (Checked profile classification)**.** *The first factors in any rank-19 decomposition* (2) *are pairwise distinct, at most one has matrix rank two, and their support belongs to one of the four orbits in Table 1.*

Table 1: The four certified first-factor support orbits. The type records the multiplicities of the three nonzero row directions among rank-one factors.

orbit	size	type	rank-two factors
0	63	$(7, 7, 5)$	none
1	21	$(7, 6, 6)$	none
2	126	$(7, 6, 6)$	none
3	42	$(6, 6, 6)$	one

The CNF, proof traces, enumeration, group action, and orbit representatives are reconstructed by the verifier described in Section 6. Proposition 4 is the only computer-assisted classification used below.

5 Excluding the four profiles

5.1 Type $(7, 7, 5)$

Two row directions occur seven times. Restricting against the annihilator of either direction removes those seven terms and leaves exactly 12. By Lemma 2, every active output factor lies in the corresponding H_x. The two active sets overlap in the five terms from the deficient row direction. Their nonzero output factors would lie in $H_x \cap H_y = \{0\}$, a contradiction. Orbit 0 is impossible.

5.2 Type $(7, 6, 6)$

Let a_0 be the multiplicity-seven direction and let a_1, a_2 be the other two. Restriction against the annihilator x_0 of a_0 leaves the 12 terms in the a_1 and a_2 groups. Lemma 2 puts all their output factors in H_{x_0}.

Now restrict against the annihilator x_1 of a_1. The seven a_0 terms and six a_2 terms survive, so Lemma 3 applies. For each of the six a_2 terms, $W_s \in H_{x_0}$, and projection modulo H_{x_1} is injective on H_{x_0} by (4). Its quotient pure tensor is therefore nonzero. Equation (7) forces all six relevant relation coefficients to be 1 and all six quotient pure tensors to equal the same nonzero z_{x_1}. The b_s within one row-direction group are distinct because the U_s are distinct. Uniqueness of nonzero pure-tensor factorization gives a contradiction. Orbits 1 and 2 are impossible.

5.3 Type $(6, 6, 6)$ with one rank-two factor

Every nonzero row restriction has 12 active rank-one terms plus the rank-two term. Lemma 3 applies to all three restrictions.

First, $z_x \neq 0$ for every nonzero x. Otherwise (7) places every output factor active under x in H_x. Choose $y \neq x$. The two active sets overlap in one complete six-term rank-one row group and the rank-two term. The seven overlap output factors are nonzero and lie in H_x, so none lies in H_y. Their quotient tensors under restriction y are nonzero. Equation (7) would make the six rank-one quotient tensors equal, contradicting their distinct b_s factors.

Since $z_x \neq 0$, every term in $\operatorname{supp}(\lambda^x)$ has the same nonzero quotient pure tensor. At most one term can come from either active rank-one row group, and at most one is the rank-two term. Hence

$$|\operatorname{supp}(\lambda^x)| \leq 3 \tag{9}$$

for each of the three restrictions.

Each of the 18 rank-one terms is active under exactly two restrictions. If its relation coefficient vanished in both, equation (7) would put its nonzero output factor in two distinct spaces H_x and H_y, contradicting (4). Thus every one of the 18 terms occurs in at least one relation support. The three supports therefore require at least 18 incidences. Equation (9) permits at most

$$3 \text{ restrictions} \times 3 \text{ terms} = 9,$$

a contradiction. Orbit 3 is impossible.

Proof of Theorem 1. Proposition 4 classifies every first-factor support of a hypothetical rank-19 decomposition. The three arguments above exclude all four orbits, so rank 19 is impossible. Wang's checked bound excludes ranks at most 18, while the 20-term witness in Appendix A proves rank at most 20. Therefore the rank is exactly 20. □

6 Exact verification

All computations use binary or integer arithmetic. The final certificate links four independently replayable components:

1. Wang's source commit and the n324 certificate/backtracking archive; a clean Linux build of the pinned verifier reports unconstrained lower bound 19 and `OK`;
2. the profile CNF and its checked DRAT and LRAT traces, followed by direct enumeration of 56,070 candidates and the four-orbit quotient;
3. the row-restriction packet, regenerated from the 252 profiles and checked to exclude every one; and
4. the 20-term witness, expanded into all $6 \cdot 12 \cdot 8 = 576$ tensor coefficients and compared with (1).

From the repository root, the exact-rank replay is

```
CERT=output/f2-234-tensor-rank/exact-rank20-certificate-v2.json
python3 -m experiments.f2_234_tight_row_restriction \
  verify-exact-rank20 --artifact "$CERT" --timeout 300
```

It returns `EXACT_RANK_VERIFIED`, rank 20, 252 excluded rank-19 profiles, and 20 upper-bound terms. The upper-witness controls and focused tests run with

```
python3 -m evaluators.f2_234_tensor_rank_eval verify-controls
python3 -m pytest tests/test_f2_234_tight_row_restriction.py \
  tests/test_f2_234_profile_orbits.py \
  tests/test_f2_234_tensor_rank_eval.py -q
```

The controls accept the rank-20 witness and Strassen's rank-7 tensor, and reject a one-bit corruption, a zero factor, and a false declared rank. The focused suite passes 35 tests.

The final certificate has SHA-256 digest

`89dbb5b9de7b13464a7bb7c78d6081a884c9ae933be923a7c787b84fcb5bd29b.`

It records the child digests for the profile packet, reviewed exclusion, Wang log and inputs, upper witness, and profile CNF/DRAT/LRAT proof pair. A fresh independent replay re-derived the restriction proof, reproduced the profile coverage and upper witness, and repeated the primary prior-art search before publication review.

Data and code availability. The exact certificate, proof traces, standalone verifier modules, tests, and paper source form the reproducibility package accompanying this manuscript. A permanent archival copy is available at doi:10.5281/zenodo.21895176; no result in this paper depends on access to that deposit.

Research-system disclosure. Beyond, the research system operated by Nth Research Collective, materially assisted with literature retrieval, hypothesis generation, exact computation, proof critique, orchestration, adversarial review, and drafting. Ian D'Ambrosio selected the target, directed the investigation, reconciled the evidence, verified the final claims, and accepts full responsibility for the manuscript. Beyond is not an author.

A An explicit 20-term upper witness

Index U by the six left-input coordinates in row-major order, V by the 12 right-input coordinates in row-major order, and W by the eight output-dual coordinates in row-major order. For a binary vector $r = (r_0, \ldots, r_{d-1})$, write its hexadecimal mask for the integer $\sum_{j=0}^{d-1} r_j 2^j$. Table 2 lists the 20 rank-one terms $u_s \otimes v_s \otimes w_s$. Expanding the masks gives 24 unit coefficients at $((i,j),(j,k),(i,k))$ and zero at the other 552 coordinates.

Table 2: Hexadecimal factor masks for the exact rank-20 decomposition.

s	u_s	v_s	w_s	s	u_s	v_s	w_s
1	20	100	90	11	20	220	60
2	3e	02a	82	12	10	010	11
3	01	008	0a	13	12	9d5	01
4	38	020	22	14	08	022	f0
5	04	880	09	15	16	980	81
6	03	040	44	16	13	045	41
7	01	044	05	17	3f	62a	02
8	08	023	50	18	36	92a	80
9	04	400	06	19	06	080	88
10	3b	620	42	20	1b	625	40

References

[1] M. Bläser, On the complexity of the multiplication of matrices of small formats, *Journal of Complexity* 19 (2003), 43–60. doi:10.1016/S0885-064X(02)00007-9.

[2] A. Fawzi et al., Discovering faster matrix multiplication algorithms with reinforcement learning, *Nature* 610 (2022), 47–53. doi:10.1038/s41586-022-05172-4.

[3] J. E. Hopcroft and L. R. Kerr, On minimizing the number of multiplications necessary for matrix multiplication, *SIAM Journal on Applied Mathematics* 20 (1971), 30–36. doi:10.1137/0120004.

[4] A. V. Smirnov, The bilinear complexity and practical algorithms for matrix multiplication, *Computational Mathematics and Mathematical Physics* 53 (2013), 1781–1795. doi:10.1134/S0965542513120129.

[5] C. Wang, Automated lower bounds for bilinear complexity over finite fields, arXiv:2603.07280v10, 2026. arXiv:2603.07280v10.