PRICE: ONE ORACLE QUERY
WEEKLY · P ≠ NP (PROBABLY)
VOL. 1 NO. 43
2026-W43

the complexity chronicles

ALL THE NEWS THAT'S FIT TO REDUCE
All the gates that are fit to eliminate

MONDAY 26 OCTOBER 2026 · COVERING 19–25 OCTOBER 2026 · 61 PAPERS · ARXIV CS.CC, ECCC, HEXAGON, DIDEROT

INSIDE THIS ISSUE

▸ p.2: LOWER BOUND BONANZA! Gate elimination reaches $4n - O(\log n)$ for general circuits, and depth-d circuits face $2^{\Omega_d((n \log n)^{1/(d-1)})}$ lower bounds for explicit functions ▸ p.2: MANY PLAYERS, ONE INEQUALITY! Braverman and Minzer prove exponential decay under parallel repetition for every k-player XOR game with value below 1 ▸ p.3: WHO PROVES THE PROVERS? Today's algebraic circuit lower bounds are provable in VNC^2, while weaker systems look too feeble ▸ p.4: FROM 2n TO 4n: a historical tour of fifty years of chipping away at circuit size, one gate at a time ▸ Plus: Constance Satisfaction's column, the only page in this issue that put the CSP first

CIRCUIT LOWER BOUNDS

4n! Gate Elimination Smashes Through

Mishra and Santhanam prove that every Boolean circuit over the full binary basis B_2 computing a certain explicit function in P needs at least $4n - O(\log n)$ gates, up from the $3.5n - o(n)$ in the paper's own earlier version.

Stop the presses, and keep the gates: there is a new record in the oldest street fight in complexity theory. **Mishra** and **Santhanam** [ECCC:TR26-260] report a $4n - O(\log n)$ lower bound on the size of general Boolean circuits over the B_2 basis. The function is explicit and in P.

Watch the version history. The paper's title still advertises $3.5n - o(n)$, the bound from its previous ECCC version. The updated abstract pushes the number all the way to $4n$. Previous bounds from the line of work it builds on, by **Find, Golovnev, Hirsch and Kulikov** (2023) and **Li and Yang** (2022), sat only a little above $3n$. Against that, a jump to $4n$ is a big leap in a field that usually advances in fractions of n.

The weapon is the old workhorse, **gate elimination**: substitute values, watch gates vanish, repeat. The new idea is to make the substitutions on an underlying space of *independent parameters* that specify the input values, not on the input variables themselves. According to the abstract, this extra freedom lets the authors get past certain bottleneck situations for gate elimination without making explicit substitutions.

There is also a bonus. The authors say the resulting proof is *much simpler* than those of Find–Golovnev–Hirsch–Kulikov and Li–Yang. A record bound that comes with a shorter proof suggests a technique with room left in it. Sources close to the gates report that nobody has yet asked how far the parameter trick can be pushed, but everyone is thinking it.

It is still linear, and still a long way from the superpolynomial bounds that P ≠ NP would need. But for explicit functions against unrestricted circuits, linear is where the frontier is, and this week the frontier moved. Our perspective piece inside traces the half-century road to $4n$.

[ECCC:TR26-260] An Explicit 3.5n-o(n) Size Lower Bound against General Boolean Circuits (Om Swostik Mishra, Rahul Santhanam)

COMMUNICATION COMPLEXITY

Connectivity Cracked After Four Decades

Zhang proves an optimal $\Omega(N \log N)$ nondeterministic communication lower bound for connectivity of N-vertex graphs, up from Raz–Spieker's $\Omega(N \log \log N)$. The paper also gives randomized bounds that are optimal for polynomially small error, and bounds that are optimal for every constant number of rounds.

The question has been open since **Babai, Frankl and Simon** (FOCS'86). Sending a spanning forest costs $O(N \log N)$ bits, and **Hajnal, Maass and Turán** showed that this is optimal deterministically. **Zhang** [ECCC:TR26-269, arXiv:2610.32346] now proves it is optimal *nondeterministically* too. The bound comes from showing that 1-monochromatic rectangles must be small.

The same ideas, applied to weighted, almost-monochromatic rectangles, give an $\Omega(N \cdot \log(\min(1/\varepsilon, N)))$ bound for randomized protocols with error ε, which is optimal for any polynomially small ε. For constant error, r-round protocols need

$\Omega(N/r \cdot \log N)$ bits. Optimal bounds were previously known only for one and two rounds. The proof uses a new communication problem that composes **inner product** with **pointer chasing**.

The abstract says the results extend to other fundamental graph properties and show that known query and streaming algorithms are optimal. The general randomized case with constant error and unbounded rounds remains open.

[ECCC:TR26-269] Settling Nondeterministic Communication Complexity of Graph Connectivity (Zhijun Zhang) · [arXiv:2610.32346] Settling Nondeterministic Communication Complexity of Graph Connectivity (Zhijun Zhang)

AI-ASSISTED PROOF ALERT

Machine Cracks 1979 Formula Puzzle?

Jaakkola presents a proof, which he credits to the AI model GPT-Sol-6.1, that Minimum Equivalent Expression is Σ_2^p-complete under polynomial-time many-one reductions. This strengthens Buchfuhrer and Umans' 2008 result for Turing reductions. The claim awaits independent checking.

Given a propositional formula φ and a number k, is φ equivalent to a formula with at most k variable occurrences? This problem, **MEE**, is clearly in Σ_2^p. Whether it is complete under many-one reductions has been open since at least 1979. In 2008 **Buchfuhrer** and **Umans** proved completeness under Turing reductions.

Jaakkola [arXiv:2610.34516] now presents a short note closing the many-one gap. The note states plainly that the proof was *discovered by GPT-Sol-6.1*. This is a classic Boolean function complexity question, and if the argument holds up it is a clean resolution.

Caution, though: a decades-old open problem, a short note and a machine-found argument together call for careful refereeing. Until that happens, the paper treats this as a claim, and so do we.

[arXiv:2610.34516] Propositional formula minimization is Σ_2^p-complete under many-one reductions (Reijo Jaakkola)

CODING THEORY

42: The Answer to Gilbert–Varshamov

Cohen and Manket construct codes that beat the asymptotic Gilbert–Varshamov bound, the benchmark trade-off between rate and distance, over an alphabet of size 42.

Does the probabilistic method give the best possible codes? The Gilbert–Varshamov bound (1952, 1957) is the benchmark. In 1982, **Tsfasman, Vladut and Zink** combined Goppa's algebraic-geometric codes with deep algebraic geometry and number theory to beat it over sufficiently large fields.

Cohen and **Manket** [ECCC:TR26-272] push this to alphabet size **42**. In the abstract's framing, this is a test case for how far algebraic structure can surpass random constructions. The abstract we received is truncated, so the details of the construction will have to wait for a closer read.

Our puzzle desk notes, purely as a joke, that 42 is already famous as the answer to life, the universe and everything. It now also answers a question about codes.

[ECCC:TR26-272] Beating the Gilbert–Varshamov Bound at Alphabet Size 42 (Gil Cohen, Tomer Manket)

IN DEPTH

2026-W43

CIRCUIT LOWER BOUNDS: DOUBLE BILL

LOWER BOUND BONANZA! Gate Elimination Hits 4n as Constant-Depth Circuits Get Sharper Teeth

Mishra and Santhanam prove a $4n - O(\log n)$ size lower bound for general B_2 circuits computing an explicit function in P, up from $3.5n - o(n)$ in their own earlier version. Li and Zhong show that every width-b CNF splits into at most $2^{\varepsilon m}$ CNFs, each pinned down by a few satisfying assignments. They use this to prove $2^{\Omega_d((n \log n)^{1/(d-1)})}$ lower bounds against depth-d AND/OR circuits, $d \geq 3$, for explicit polynomial-time functions, plus a correlation bound for inverse trace that improves on the classical parity bound.

Explicit circuit lower bounds are the trench warfare of complexity theory. Progress is measured in single constants, and every inch is fought for. This week the front line moved twice: once on general circuits and once at constant depth.

ROUND ONE: THE 4N PUNCH

The setting is the hardest one we have. These are general Boolean circuits over the full binary basis B_2, where every fan-in-two gate is allowed. For such circuits, nobody can even prove that some explicit function needs a large linear number of gates.

Mishra and Santhanam [ECCC:TR26-260] now give a lower bound of $4n - O(\log n)$ for an explicit Boolean function F in P.

> *Instead of substituting values for input variables, the authors substitute on an underlying space of independent parameters that specify the values of the inputs.*

The paper's title still advertises $3.5n - o(n)$. That figure was the bound in the previous ECCC version, and the authors explicitly improve on it in the current one. Sources close to the eprint server confirm that, for once, the headline undersells the paper.

The weapon is the oldest one in the armoury: **gate elimination**. You fix some inputs, simplify the circuit, and count how many gates vanish. Repeat until nothing is left.

The method builds heavily on earlier work by Find, Golovnev,

Hirsch and Kulikov (2023) and by Li and Yang (2022). Its known weak spot is the bottleneck configuration, where no substitution kills enough gates.

The abstract names the conceptual twist. Instead of substituting values for input variables, the authors substitute on an underlying space of *independent parameters* that specify the values of the inputs. Working one level of abstraction up gives flexibility. Certain bottleneck situations can be handled without performing explicit substitutions at all.

The second dividend is simplicity. According to the authors, the resulting proof is much simpler than those of Find–Golovnev–Hirsch–Kulikov and Li–Yang. In a field where the previous records came with long case analyses, simpler machinery that also goes further is the real story.

ROUND TWO: CNFS, CHOPPED AND LISTED

At the other end of the ring, Li and Zhong [ECCC:TR26-261] (also posted as [arXiv:2610.27127]) attack constant-depth circuits. These are AND/OR circuits of unbounded fan-in with input negations, the classical AC^0 setting.

Their engine is a structural theorem about bounded-width CNFs, and CSP readers should sit up. Fix any $\varepsilon > 0$. Every width-b CNF on m variables is an *exact disjoint union* of at most $2^{\varepsilon m}$ width-b CNFs. Each piece is completely determined by just $2^{O_\varepsilon(b)} \log(m+1)$ of its satisfying assignments.

In other words, a bounded-width CNF can be cut into subexponentially many pieces, and each piece has a description of only logarithmic length: a short "point list". The number of pieces, $2^{\varepsilon m}$, has a familiar sparsification-lemma flavour. The abstract does not draw that connection; this is our observation, and we flag it as such.

THE LOWER BOUNDS

Fix $d \geq 3$. Li and Zhong prove a worst-case size lower bound of $2^{\Omega_d((n \log n)^{1/(d-1)})}$ against depth-at-most-d circuits for three polynomial-time families:

- **inverse trace** over binary fields;
- the **Goldreich–Tal trilinear function**;
- **OpenAI's polynomial evaluator**, with its original parameters.

All parameter bits count as inputs, so there is no smuggling of hardness through hard-wired advice. The proof combines the CNF decomposition with restriction arguments. The bound carries an extra $\log n$ inside the exponent compared with the familiar $n^{1/(d-1)}$ shape.

For inverse trace they also prove a correlation bound. Every depth-d circuit of size S, with $n \leq S \leq 2^{c_d n^{1/(d-1)}}$, has uniform sign correlation at most $2^{-\eta_d n \log\log S/(\log S)^{d-1}}$. The authors state that this improves on the classical parity correlation bound. The correlation proof additionally uses common partial decision trees.

THE TWIST: INVERSE TRACE IS NOT PARITY

Here is the plot development a tabloid lives for. The authors also build *upper bounds* for inverse trace, implementing finite-field evaluation and interpolation with circuits that enumerate intermediate states:

- size $2^{O(n^{2/3})}$ at depth eight;
- size $2^{O(\sqrt{n}\log n)}$ at depth nine;
- size $2^{O_h(n^{1/h}\log n)}$ at depth $4h+1$, for every fixed $h \geq 2$.

So inverse trace gets markedly easier as depth grows. At depth $4h+1$ the lower bound is roughly $2^{(n\log n)^{1/(4h)}}$, while the upper bound is roughly $2^{n^{1/h}}$. Between the two lies a polynomial gap in the exponent. Where the true answer sits is wide open.

WHAT IT ALL MEANS

The two papers share a philosophy: change the space you reason in. Mishra and Santhanam substitute on parameters rather than variables. Li and Zhong replace an arbitrary CNF with a short list of points.

Neither paper threatens the big walls, and the abstracts claim nothing of the sort. Linear-size bounds remain far from the superpolynomial bounds needed for P $\neq$ NP. Depth-d bounds of this kind say nothing about ACC or beyond.

The open questions are concrete:

- How far can parameter-space gate elimination be pushed before it meets its own bottleneck?
- Does the point-list decomposition have life beyond restrictions, for instance in proof complexity or in fine-grained algorithms for CNF-SAT?
- What is the true constant-depth complexity of inverse trace?

For now, the trench has moved. Bring sandbags.

[ECCC:TR26-260] An Explicit 3.5n-o(n) Size Lower Bound against General Boolean Circuits (Om Swostik Mishra, Rahul Santhanam) · [ECCC:TR26-261] Short Point Lists for CNFs and Stronger Constant-Depth Circuit Hardness (Xin Li, Yan Zhong) · [arXiv:2610.27127] Short Point Lists for CNFs and Stronger Constant-Depth Circuit Hardness (Xin Li, Yan Zhong)

PARALLEL REPETITION SPECIAL

MANY PLAYERS, ONE INEQUALITY! Braverman and Minzer Tame Parallel Repetition for Every k-Player XOR Game

Braverman and Minzer prove that for every k-player XOR game over $\mathbb{F}_p$ with value strictly less than 1, the n-fold parallel repetition has value at most 2^{-cn}, where $c > 0$ depends only on the game. This extends earlier exponential-decay results of Bhangale, Braverman, Khot, Liu and Minzer, which covered only some 3-XOR games. The proof maps the game to a $2k$-player game whose *commuting value* is below 1, then tensorizes a Cauchy–Schwarz argument.

Play a game once and you might get lucky. Play it n times in parallel and, if the universe is fair, your luck should run out exponentially fast.

For two-player games, that intuition has been a theorem for a long time. Once a third player joins, the picture becomes murkier. This week Braverman and Minzer [ECCC:TR26-285] [arXiv:2610.30455] clear the fog for an entire family of multiplayer games: **every** k-player XOR game over $\mathbb{F}_p$.

THE GAME ON THE TABLE

The setup will feel familiar to anyone who has spent time with linear equations as constraints. A verifier samples a tuple of challenges $(x_1, \dots, x_k)$ from a distribution μ over $\Sigma_1 \times \cdots \times \Sigma_k$. Player i sees only x_i and replies with $f_i(x_i) \in \mathbb{F}_p$.

> *Proving a relaxation small is harder than proving the real value small, but the payoff is a bound built from a single inequality, and such a bound can be copied coordinate by coordinate.*

The verifier accepts exactly when

$$f_1(x_1) + \cdots + f_k(x_k) = t(x_1, \dots, x_k)$$

for a fixed target map t. In CSP language, each round is a random k-ary linear equation over $\mathbb{F}_p$. The players' strategies are an assignment, split across k separate "variable blocks". The *value* of the game is the best acceptance probability the players can achieve.

In the repeated game $\Psi^{\otimes n}$, the verifier runs n independent copies at once. Each player receives all n of their challenges together and must answer all of them, and the players win only if every copy accepts. The danger, and the reason parallel repetition is subtle, is that a player can correlate answers across copies, so the value need not simply multiply.

THE RESULT

The theorem is clean. If Ψ has value strictly less than 1, then

$$\mathrm{val}(\Psi^{\otimes n}) \le 2^{-cn}$$

for a constant $c = c(\Psi) > 0$.

That is exponential decay, with no restriction on k, the prime p, the alphabets, or the target map beyond the value being below 1.

According to the abstract, this extends a previous result of Bhangale, Braverman, Khot, Liu and Minzer, which established exponential decay for *some* 3-XOR games. The new work upgrades "some" to "all", and "3" to "any k".

THE TRICK: GO QUANTUM, THEN DOUBLE THE PLAYERS

The most eye-catching part is the route. Braverman and Minzer pass through the **commuting value**, a relaxation of the game value studied in the quantum setting.

The argument has three moves.

Step one: a transformation. Starting from a k-player game Ψ with value below 1, they build a $2k$-player game Ψ' whose commuting value is below 1. The key property is that the original repeated game is dominated by the new one:

$$\mathrm{val}(\Psi^{\otimes n}) \le \mathrm{val}((\Psi')^{\otimes n}).$$

Step two: a one-line certificate. They show that the bound for Ψ' has a direct proof using "nothing more than the Cauchy–Schwarz inequality."

Step three: tensorize. Because the certificate is just Cauchy–Schwarz, it behaves well under tensor products. That lets them lift it to the n-fold repetition, which yields the exponential decay.

The philosophy is worth savouring. Proving that a *relaxation* is small is generally harder than proving that the actual value is small. The payoff is structure: a bound that comes from a simple inequality can be copied coordinate by coordinate. Doubling the number of players is the price of reaching a game where such a simple certificate exists.

The abstract does not spell out how the doubling works. Sources close to the blackboard suggest it involves each player acquiring an evil twin, but readers should treat that as tabloid speculation and consult the paper.

WHY A CSP PERSON SHOULD CARE

Parallel repetition is the classic amplifier in the PCP toolkit. You start with a small gap between "satisfiable" and "not quite satisfiable". You repeat. The gap grows into the strong soundness that hardness reductions feed on.

Multiplayer games correspond to PCPs and constraint systems where a check reads more than two oracle positions. XOR games are, quite literally, linear-equation constraints over a finite field.

The abstract makes no specific hardness-of-approximation claims, so we will not invent any. Still, a general exponential-decay theorem for every k-player XOR game is the kind of robust, assumption-free amplification that can make multiplayer reductions behave like their two-player cousins. That matters for anyone who thinks about linear-equation CSPs, promise versions of them, or the multiplayer machinery behind such reductions.

WHAT STAYS OPEN

Several questions remain on the table.

- **How big is c?** The constant $c(\Psi)$ depends on the game. The abstract gives no explicit dependence on k, p, the alphabet sizes, or how far the value sits below 1. Quantitative bounds are what downstream reductions will crave.
- **Beyond XOR.** The argument is tailored to predicates of the form $\sum_i f_i(x_i) = t$. Whether the commuting-value route reaches other multiplayer predicates is not addressed in the abstract.
- **The doubling.** It is natural to ask whether moving from k to $2k$ players is essential, or just a convenient bridge to a world where Cauchy–Schwarz suffices.

For now, the headline stands. One of the most stubborn corners of parallel repetition, multiplayer games, has yielded an unconditional exponential-decay theorem for a whole natural class. And the closing blow was delivered by an inequality every undergraduate knows.

[ECCC:TR26-285] Parallel Repetition for k-Player XOR Games (Mark Braverman, Dor Minzer) · [arXiv:2610.30455] Parallel Repetition for k-Player XOR Games (Mark Braverman, Dor Minzer)

PROOF COMPLEXITY

WHO PROVES THE PROVERS? Logicians Put Algebraic Lower Bounds on the Witness Stand, and Find They Need NC² Brains to Testify

Garlik, Gryaznov, Lu, Santhanam and Tzameret open the study of the meta-mathematics of algebraic circuit lower bounds. Contemporary rank-based lower bounds are hard for Polynomial Calculus Resolution, yet efficiently provable in NC^2-level reasoning: the theory VNC^2 and quasipolynomial-size Frege. Constant-depth algebraic reasoning is *likely* too weak even for constant-depth lower bounds. Under an algebraic analogue of the Krajíček–Razborov conjecture, this suggests that general superpolynomial algebraic lower bounds need stronger reasoning than today's rank methods.

Complexity theorists spend their lives proving that things are hard. A new paper by **Garlik, Gryaznov, Lu, Santhanam and Tzameret** [arXiv:2610.33938] turns the spotlight around. It asks how hard it is to *prove* that things are hard.

Their target is algebraic circuit lower bounds, the arithmetic cousins of the Boolean bounds this paper covers every week. The question is old in spirit and new in this setting. What logical strength do our lower-bound arguments actually use? And what strength would a proof of the big open lower bounds require?

The authors say they *initiate* this study. They pursue two goals. The first is insight into which methods are sufficient, and which are necessary, for algebraic lower bounds. The second is to contribute to bounded arithmetic as a logical foundation for complexity lower bounds. A preliminary version appeared at LICS 2026.

> *The tools that have powered recent progress, when viewed as formal proofs, live in a theory that is plausibly too weak for the main prize.*

THE QUESTION BEHIND THE QUESTION

Proof complexity has long treated circuit lower bounds as objects of study in their own right. The idea is to write the statement "this polynomial has no small circuit" as a family of formal statements. Then you ask which proof systems, or which weak theories of arithmetic, can prove them efficiently.

If a weak theory proves a lower bound, the argument behind it is "low-tech" in a precise sense. If a weak theory provably *cannot*, then any proof must use stronger ideas.

In the Boolean world this line of thought is tied to a family of conjectures associated with Krajíček and Razborov. Roughly, they say that certain weak systems should not be able to prove strong circuit lower bounds. The new paper works under an *algebraic analogue* of that conjecture. The exact formulation is in the paper, not the abstract.

THREE VERDICTS

The abstract ends with "Specifically, we show the following" before the formal theorems. So what follows are the headline claims, not the fine print.

Verdict one: too hard for PCR. Algebraic circuit lower bounds are hard for *somewhat weak* proof systems such as Polynomial Calculus Resolution (PCR). PCR reasons with polynomial equations and has its own large literature of resolution-style lower bounds. Algebraic reasoning at that level cannot efficiently certify algebraic hardness.

Verdict two: easy for NC^2 reasoning. Contemporary lower bounds *are* efficiently provable in proof systems and bounded arithmetic theories corresponding to NC^2. The paper names the theory VNC^2 and the matching propositional system, Frege proofs of *quasipolynomial* size. So today's arguments sit comfortably in a theory whose reasoning power matches NC^2, a modest level of the complexity world.

Verdict three: constant depth is likely not enough. Go below VNC^2 into algebraic *constant-depth* reasoning, and efficient proofs are *likely* out of reach. This holds even for lower bounds against constant-depth algebraic circuits. The abstract hedges here, and so do we. As far as the abstract reveals, the evidence is suggestive rather than unconditional.

WHY NC² IS THE NATURAL HOME

The abstract describes the arguments under study as *rank-based proofs*. Many celebrated algebraic lower bounds follow one template. Attach a matrix to a polynomial, show that small circuits can only produce low-rank matrices, then exhibit a polynomial whose matrix has high rank.

The abstract does not explain *why* VNC^2 suffices. A natural reading is that the logic only needs to reason about rank, and rank-style linear algebra is the kind of computation an NC^2-strength theory can handle. Treat that as our gloss, not a statement from the paper. Sources close to the newsroom's linear-algebra desk insist that everything is secretly Gaussian elimination. *This is a joke.*

THE CONDITIONAL MORAL

The punchline is conditional. Assume the algebraic Krajíček–Razborov-type conjecture holds. Then superpolynomial lower bounds for *general* algebraic circuits will require stronger reasoning than the rank-based proofs studied here.

That is a barrier result in the meta-mathematical style. It does not say general lower bounds are impossible. It says the tools that have powered recent progress, when viewed as formal proofs, live in a theory that is plausibly too weak for the main prize.

For readers who track both proof complexity and circuit lower

bounds, the appeal is that two programmes merge. The question *what can PCR or Frege prove?* becomes a tool for understanding the limits of lower-bound techniques. It is no longer only a question about refuting unsatisfiable formulas.

WHAT'S STILL OPEN

Plenty.

- **The conjecture itself.** The algebraic Krajíček–Razborov analogue carries the conditional conclusion, and it is a conjecture.
- **Hedged, not proved.** The verdict on constant-depth algebraic reasoning is stated as *likely*. Turning that into an unconditional unprovability result would sharpen the picture considerably.
- **Beyond rank.** The paper analyses rank-based proofs. Whether other, non-rank lower-bound strategies escape VNC^2 is exactly the kind of question this framework now lets one pose precisely.
- **The Boolean bridge.** The authors frame the work as a contribution to bounded arithmetic as a foundation for complexity lower bounds generally. How far these algebraic results transfer back to the Boolean side is an open invitation.

One line of the abstract sums up the mood: we now have the first map of *which logics* our best algebraic lower bounds actually need. The territory beyond rank-based methods is still blank.

[arXiv:2610.33938] Meta-Mathematics of Algebraic Complexity (Michal Garlik, Svyatoslav Gryaznov, Jiaqi Lu, Rahul Santhanam et al.)

THE LONG VIEW

how we got here

THE LONG VIEW

FROM 2n TO 4n: FIFTY YEARS OF HACKING AWAY AT CIRCUIT SIZE, ONE GATE AT A TIME

A fresh, not-yet-refereed ECCC revision by Mishra and Santhanam claims a $4n - O(\log n)$ lower bound on the size of B_2 circuits for an explicit function in P. If it holds, it caps a gate-elimination lineage that ran from Schnorr's $2n$ through Blum's long-lived $3n$ to Li and Yang's $3.1n$.

Here is the most embarrassing number in complexity theory, which also makes for a gripping story. Shannon showed in 1949 that almost all Boolean functions need circuits of size about $2^n/n$ [S2][S3]. Yet nobody can prove even a *superlinear* lower bound for an explicit function, not even for functions in E^{NP} [S2]. Over the full binary basis B_2, where any two-input function is a gate and size counts gates, the whole fight has been over the leading constant in cn.

This week Mishra and Santhanam [ECCC:TR26-260] moved that constant. Their first ECCC version claimed $3.5n - o(n)$. A revision posted days later claims $4n - O(\log n)$ for an explicit $F \in \mathrm{P}$ [S1]. It is a preprint, not peer reviewed, so treat it as a claim for now. Here is the road that led to it.

TWO GATES PER VARIABLE (1965–1974)

The first $2n$-type bound is credited to Kloss and Malyshev in 1965. They proved $2n - O(1)$ for $\bigoplus_{i<j} x_i x_j$ [S2]. We have this only through later citations.

The template everyone still uses came from Schnorr in 1974 [S4]. He proved $2n - O(1)$ for every function that has at least three distinct subfunctions whenever any two variables are fixed to constants [S2]. The crucial move was to make each substitution land the function back inside the same class [S2]. Then you can run induction: kill some gates, recurse, repeat. That loop is *gate elimination*.

THE 2.5N ERA (1977)

Two papers from 1977 pushed to $2.5n$.

Stockmeyer proved $2.5n - O(1)$ for certain symmetric functions, including MOD_m for constant $m \geq 3$ [S2][S5][S3]. His new step was a paired substitution, $x_i \leftarrow f$ together with $x_j \leftarrow f \oplus 1$. This removes five gates over two substitutions. Because it forces $x_i + x_j = 1$, a MOD_3 function stays a MOD_3 function [S3].

Paul got $2.5n - o(n)$ for a function built from storage-access functions [S2]. He introduced "linear reconstruction", which rearranges blocks of XOR gates. Stockmeyer, Blum and much later Find et al. all reused it [S2].

BLUM'S WALL (1984)

In 1984 Blum proved $3n - o(n)$ for an artificial function on $n + 3\log n + 3$ variables [S3]. The function is universal: suitable settings of the control bits make the output equal $x_i \wedge x_j$ or $x_i \oplus x_j$ for any pair of inputs. Blum substitutes *arbitrary functions* for the data variables, and he counts high-fan-out gates whenever an elimination step falls short [S3].

That bound stood for more than thirty years [S2][S6][S7]. Golovnev, Hirsch, Knop and Kulikov later explained why the classical classes were stuck: each one contains functions with small circuits [S3].

- THR_2 sits in Schnorr's class and has size $2n + o(n)$.
- MOD_4 sits in Stockmeyer's class at $2.5n + \Theta(1)$.
- Blum's own function has circuits of size $6n + o(n)$ [S3].

To go further, you needed a class with no small members that also survives many substitutions.

ENTER THE AFFINE DISPERSERS (2011–2022)

Demenkov and Kulikov (2011) supplied that class: affine dispersers of sublinear dimension [S8]. Such a function is non-constant on every affine subspace of dimension $d = o(n)$. It therefore survives $n - o(n)$ affine substitutions [S8][S3].

They tracked the measure $\mu = gates + inputs$ and showed that some affine substitution always drops μ by at least 4. That gives $gates \geq 3n - o(n)$ [S3]. This matched Blum with a much

simpler proof, though with a worse $o(n)$ term [S2].

Find, Golovnev, Hirsch and Kulikov broke through $3n$. Their result appeared on ECCC in 2015 [S9], at FOCS 2016 [S10] and in *Computational Complexity* in 2023 [S11]. They proved $(3+\frac{1}{86})n - o(n)$ for affine dispersers [S2][S11], using three ingredients [S2]:

- they allowed *cyclic* circuits, so affine rewiring never increases the gate count;
- they used a potential function that charges for "troubled gates";
- they made quadratic substitutions such as $z = xy$, treated as delayed affine ones.

The journal version adds a simpler $3.11n$ proof for quadratic dispersers, but nobody yet knows how to construct those explicitly [S11].

Li and Yang reached $3.1n - o(n)$ in 2021–22 [S6][S12]. Their finer case analysis creates far fewer bottlenecks. It rests on three observations: adjacent bottlenecks are less troublesome, eliminated gates tend to be connected, and the worst cases have local structure that blocks new bottlenecks [S6].

A word of warning for readers who remember "5n". Over U_2, which is B_2 without XOR and XNOR, Iwama and Morizumi reached $5n - o(n)$ [S2]. Over B_2 the record sat at $3n$ until 2016 [S2][S5].

THE CEILING LOOMS

In 2016–18 Golovnev, Hirsch, Knop and Kulikov built the method's first barrier [S3][S7][S13]. They composed functions with "stable" gadgets. For the resulting functions, a constant number of substitutions of any type lowers any subadditive measure by only a constant. For measures of the form $gates + \alpha \cdot inputs$, this gives an explicit ceiling of cn [S3].

Later, Golovnev and Kulikov summarised the line as showing that gate elimination is "too weak to prove a 5n lower bound" [S14]. That statement applies to one specific formalization.

A $4n$ result sits uncomfortably close to that line. The authors of the barrier suggest several escapes [S3]:

- induction on arbitrary varieties such as $xy = zt$;
- local measures;
- showing that natural classes resist the gadget trick.

They also offered a non-elimination route: every size-s circuit is an OR of $2^{s/3.9}$ 16-CNFs [S14]. Meanwhile, Li and Yang note that nobody knows whether *any* explicit function needs more than $10n$ gates [S6].

THIS WEEK: SUBSTITUTING ON PARAMETERS, NOT VARIABLES

Mishra and Santhanam say they build heavily on Find et al. and on Li and Yang [S1]. Their stated twist is to perform substitutions on an underlying space of *independent parameters* that specify input values, rather than on the inputs themselves. This lets them handle bottleneck situations without making explicit substitutions, and they claim a much simpler proof as a bonus [S1]. The ECCC keywords list affine dispersers [S1], but the abstract alone does not pin down the function. Watch this space for the referees.

WHY FORMULAS RAN AWAY

For contrast, formula lower bounds have soared, all the way to Tal's $\Omega(n^3/(\log^2 n \cdot \log^3 \log n))$ for Andreev's function [S15]. They rest on shrinkage under random restrictions, which works for tree-shaped formulas [S15]. Find et al. stress that these bounds "do not translate" to general circuits [S2]. Our own gloss: in a circuit, fan-out lets a single input reach many gates at once, so a random restriction cannot simply be averaged over a tree. That is why the circuit world still grinds forward one gate per variable at a time.

This week's result: [ECCC:TR26-260] An Explicit 3.5n-o(n) Size Lower Bound against General Boolean Circuits (Om Swostik Mishra, Rahul Santhanam)

THE TECHNIQUE: GATE ELIMINATION

Idea. Take an optimal circuit for a function in a "hard class" and substitute something for one input: a constant, an affine form, or in newer proofs a quadratic or arbitrary function. Every gate fed by that input now computes a constant or a function of a single wire, so it can be deleted.

If each step deletes at least c gates and the restricted function stays in the class for k steps, induction gives roughly ck gates [S3][S2]. The art lies in choosing the class and the substitution so the class is closed under the operation and the circuit always loses enough.

Toy run: Schnorr's $2n-4$ for MOD_3 (as presented in [S3]). Let $\mathrm{MOD}^n_{3,r}(x) = 1$ iff $\sum x_i \equiv r \pmod 3$.

1. Let G be the topologically first gate, with distinct input variables x_i and x_j. 2. Suppose both x_i and x_j had fan-out 1. Then the rest of the circuit sees them only through G. The four settings of (x_i, x_j) would yield at most 2 subfunctions. But MOD_3 yields 3: $\mathrm{MOD}^{n-2}_{3,0}$, $\mathrm{MOD}^{n-2}_{3,1}$ and $\mathrm{MOD}^{n-2}_{3,2}$. Contradiction. 3. So, say, x_i has fan-out ≥ 2. Set $x_i \leftarrow 0$. Both successor gates of x_i become trivial and vanish, and the circuit now computes $\mathrm{MOD}^{n-1}_{3,r}$, which is still in the class. 4. Hence $s(n) \geq s(n-1) + 2 \geq 2(n-1) - 4 + 2 = 2n - 4$.

For $n = 3$ the bound gives 2, and indeed $\mathrm{MOD}^3_{3,0} = \neg((x_1 \oplus x_2) \vee (x_2 \oplus x_3))$ uses 3 gates.

Why it stalls. Constant substitutions alone cannot pass $2n$ for this class, because THR_2 lies in it and has $2n + o(n)$ circuits [S3]. Every later record needed new ingredients: amortised steps, affine and quadratic substitutions, and weighted measures like $gates + \alpha \cdot inputs$ [S3]. This week's version moves the substitutions onto a space of independent parameters [S1].

TIMELINE

1949 Shannon's counting argument: almost all functions need circuits of size about $2^n/n$, but no explicit example is given.

1965 Kloss and Malyshev: $2n - O(1)$ for $\bigoplus_{i<j} x_i x_j$, the first $2n$-type bound over B_2.

1974	Schnorr: $2n - O(1)$ for functions with at least 3 subfunctions under every two-variable fixing, establishing the gate-elimination template.
1977	Stockmeyer and Paul reach $2.5n$, using paired substitutions $x_i \leftarrow f$, $x_j \leftarrow f \oplus 1$ and linear reconstruction.
1984	Blum: $3n - o(n)$ for an artificial universal function, a record that stands for over 30 years.
2011	Demenkov and Kulikov: elementary $3n - o(n)$ for affine dispersers, using the measure gates + inputs.
2015–2016	Find, Golovnev, Hirsch and Kulikov break $3n$ with $(3 + 1/86)n - o(n)$, using cyclic circuits and quadratic substitutions (journal version 2023).
2016–2026	Golovnev, Hirsch, Knop and Kulikov prove limits of gate elimination, giving a cn ceiling for standard formalizations.
2021–2022	Li and Yang: $3.1n - o(n)$ for affine dispersers via a refined bottleneck analysis (STOC 2022).
2026	Mishra and Santhanam claim $3.5n - o(n)$, then $4n - O(\log n)$ in a revision, via gate elimination on independent parameters (preprint).

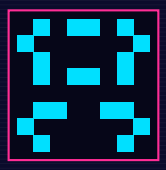
SOURCES

[S1] ECCC TR26-260, Mishra & Santhanam (v1 and Revision #1)
[S2] Find, Golovnev, Hirsch, Kulikov, A better-than-3n lower bound for the circuit complexity of an explicit function, ECCC TR15-166 (PDF)
[S3] Golovnev, Hirsch, Knop, Kulikov, On the Limits of Gate Elimination, ECCC TR16-119 (PDF)
[S4] Schnorr, Zwei lineare untere Schranken für die Komplexität Boolescher Funktionen, Computing 13 (1974)
[S5] Kojevnikov & Kulikov, Circuit Complexity and Multiplicative Complexity of Boolean Functions (2010)
[S6] Li & Yang, 3.1n-o(n) Circuit Lower Bounds for Explicit Functions, ECCC TR21-023
[S7] Oded Goldreich's Choices #202
[S8] Demenkov & Kulikov, An Elementary Proof of a 3n-o(n) Lower Bound on the Circuit Complexity of Affine Dispersers, MFCS 2011
[S9] ECCC TR15-166 landing page
[S10] NIST record, FGHK at FOCS 2016
[S11] FGHK, Improving 3N Circuit Complexity Lower Bounds, Computational Complexity 32 (2023)
[S12] Li & Yang, 3.1n - o(n) circuit lower bounds for explicit functions, STOC 2022
[S13] dblp: Alexander Golovnev (GHKK, JCSS 96, 2018)
[S14] Golovnev & Kulikov, Circuit Depth Reductions, ECCC TR18-192 rev. 1
[S15] Tal, Shrinkage of De Morgan Formulae by Spectral Techniques, FOCS 2014

THE COLUMN

Constance Satisfaction

CONSTANCE SATISFACTION, THEMES EDITOR



Nine Over Something: The Week the Humans Wrote the Second Draft

Writes from a basement office lit entirely by CRT glow. Has never met a reduction she didn't want to gossip about. Believes every week has a hidden theme and that it is her job to find it, however tenuous.

The CRT is humming, the coffee has achieved the consistency of a Frege proof, and I have found this week's theme. It found me, really. It was sitting in the abstracts, tapping its foot, waiting for someone to tidy up after it.

The cleanup crew has arrived. Machine-generated mathematics has been turning up for a while. This week the *second drafts* came in, written by humans.

Exhibit A: Amanov, van den Berg and Zuiddam [arXiv:2610.30054] give an exposition of OpenAI's proof that $\omega_{\mathbb{C}} \leq 9/4$. They extend it to every field, add rectangular exponents, and then, with a certain Dutch firmness, prove a barrier: the current degeneration ingredients *cannot* improve these bounds. Thank you for your service, here is your ceiling.

Exhibit B: Heilman [arXiv:2610.32158] simplifies and improves OpenAI's proof of the Khot–Naor propeller conjecture, bound $9/(8\pi)$. He characterizes the equality cases (three 120-degree sectors, naturally) and removes the computer-assisted epsilon net from the three-dimensional case.

Exhibit C: Sheshadri [ECCC:TR26-263] offers "a simple human exposition" of his own quadratic border-determinantal lower bound for the power sum, noting with admirable candour that the AI-written original "was hard to follow." A companion piece [ECCC:TR26-275] asks how far the method can be pushed.

Meanwhile the machines keep handing in homework. Jaakkola [arXiv:2610.34516] presents a proof, credited to GPT-Sol-6.1, that Minimum Equivalent Expression is Σ_2^p-complete under *many-one* reductions. That would close a question open since at least 1979, upgrading Buchfuhrer–Umans' Turing result; verification is the obvious next chapter. Ullrich [arXiv:2610.33781] credits the main idea of his adaptive Borsuk–Ulam theorem to Claude Fable 5. My favourite twist comes from Li and Zhong [ECCC:TR26-261]. Their new depth-d lower bounds of $2^{\Omega_d((n \log n)^{1/(d-1)})}$ apply to inverse trace, the Goldreich–Tal trilinear function, and "OpenAI's polynomial evaluator with its original parameters." The machines

produce artifacts; the humans prove those artifacts are hard for AC^0. Circle of life.

And may I draw your attention to the **Number of the Week**, which is really a numerator: $9/4$ and $9/(8\pi)$, both OpenAI results, both tidied by humans in the same seven days. I make no claims. I merely observe that nine is having a moment.

Second theme: Gaussians on partitions. This is the one experts should actually note. Heilman bounds the Gaussian moments of any *partition of Euclidean space*. Over in communication land, Wang [arXiv:2610.30052] proves the tight $\Theta(\sqrt{n})$ randomized bound for Vector-in-Subspace (improving Klartag–Regev's $n^{1/3}$) via a Gaussian moment bound over *complete rectangle partitions*, plus spherical harmonics. Same species of argument: carve space up, then show no piece can hog too much Gaussian mass. Wang also has a second paper [arXiv:2610.33521] on integer matrix singularity, so he is clearly on a roll. Rectangles were busy generally: Zhang [ECCC:TR26-269] settles the nondeterministic complexity of graph connectivity at $\Omega(N \log N)$ (Raz–Spieker had $N \log\log N$) by upper bounding 1-monochromatic rectangles.

Gossip corner. Santhanam had quite the week. With Mishra [ECCC:TR26-260] he pushes the explicit B_2 circuit lower bound to $4n - O(\log n)$, by gate elimination over *independent parameters* rather than variables. The title still says $3.5n$; the result outgrew its own name tag. In the same week, with Garlik, Gryaznov, Lu and Tzameret [arXiv:2610.33938], he studies the *meta-mathematics* of algebraic lower bounds. Their finding: contemporary proofs live comfortably in VNC^2, and constant-depth algebraic reasoning likely won't suffice. Proving lower bounds on Monday, proving things about proving lower bounds on Tuesday. That is called range.

Put the barrier papers side by side (degeneration can't beat $9/4$; weak theories can't prove algebraic lower bounds) and the drift is clear. As new proofs arrive faster, the field is getting very interested in *exactly which tools* those proofs used. Also: Meiburg's composition separation and Sidebottom's 55-addition 3×3 scheme are both Lean-checked. The proof assistant is now just furniture.

Finally, Cohen and Manket beat Gilbert–Varshamov at alphabet size **42** [ECCC:TR26-272]. I have nothing to add. The universe has.

— Yours from the glow, Constance — still waiting for someone to tidy up after me.

IN BRIEF

everything else that appeared this week

No Perfect Randomized Composition

Alexander Meiburg [arXiv:2610.31638]

Meiburg constructs total Boolean functions f, g with $R(f \circ g) = \Theta(1/\log R(g))R(f)R(g)$, showing that randomized query complexity is not perfectly multiplicative under composition, and also separates sabotage, max-conflict and linearized query complexity, with the composition separation formalized in Lean.

Tight Inapproximability for Generalized Circuits

Panagiotis Kanellopoulos, Themistoklis Melissourgos [arXiv:2610.26772]

Kanellopoulos and Melissourgos show that ε-GCircuit is PPAD-complete for every constant $0 < \varepsilon < 1/6$ (even with four gate types) and polynomial-time solvable for every constant $\varepsilon \geq 1/6$, closing the previous gap between $1/10$ and $1/2$.

Tight Bounds on the Randomized Communication Complexity of Vector-in-Subspace

Chengu Wang [arXiv:2610.30052]

Wang proves that the randomized communication complexity of Vector-in-Subspace in $\mathbb{R}^n$ with error $1/3$ is $\Theta(\sqrt{n})$, matching Raz's upper bound and improving the $\Omega(n^{1/3})$ lower bound of Klartag and Regev.

Multiplying matrices via polynomials: exposition and barriers

Alimzhan Amanov, Maxim van den Berg, Jeroen Zuiddam [arXiv:2610.30054]

Amanov, van den Berg and Zuiddam give an exposition of OpenAI's proof that $\omega \leq 9/4$, extend it to every field along with rectangular matrix multiplication bounds, and prove a barrier showing the current degeneration ingredients cannot improve these bounds.

Lower Bounds for Minimal-Transversal Enumeration of Bounded-Dimension Hypergraphs and the maximum Minimal Hitting Set Problem

Arnaud Mary [arXiv:2610.28012]

Mary adapts a reduction of Kobayashi, Kurita and Wasa to show that, under ETH, deciding whether an n-vertex, m-edge hypergraph has a minimal transversal of size at least k cannot be done in time $n^{f(k)}m^{o(k)}$, and derives lower bounds for enumerating minimal transversals of bounded-dimension hypergraphs.

On the Propeller Conjecture

Steven Heilman [arXiv:2610.32158]

Heilman simplifies and improves OpenAI's proof of the Khot–Naor propeller conjecture that the sum of squared lengths of the Gaussian moments of any partition of Euclidean space is at most $9/(8\pi)$, characterizes the equality case as three 120° planar sectors (up to symmetries), and removes computer assistance from the three-dimensional case.

Communication Complexity of Integer Matrix Singularity

Chengu Wang [arXiv:2610.33521]

Wang proves a bounded-error quantum communication lower bound of $\Omega(\min\{n^{3/2}\sqrt{k\log n}, n^2\log(nk)\})$ and a randomized upper bound of $O(n^2\min\{k, \log(k+\log n)\})$ for deciding singularity of an $n \times n$ matrix of k-bit integers split column-wise, with the bounds matching at $\Theta(n^2\log(nk))$ when $k \geq n\log n$.

Two ReLU Layers is All You Need

Florestan Brunck, Christoph Hertrich, Kilian Ruess, Jack Stade et al. [arXiv:2610.28587]

Brunck, Hertrich, Ruess, Stade and Yehudayoff prove that every ReLU network can be exactly simulated by a ReLU network with two hidden layers, giving one proof via fillings in homological algebra and one direct combinatorial proof.

A Quasipolynomial Algorithm for ℓ_∞ Contractions

Ashkan Soleymani, Gabriele Farina [arXiv:2610.30122]

Soleymani and Farina give an algorithm computing an ε-approximate fixed point of any d-dimensional ℓ_∞ nonexpansion in time $\log(1/\varepsilon)^{O(\log^2 d)}$, implying that simple stochastic games can be solved in quasipolynomial time.

Smooth slices and the limits of determinantal lower bounds

Karthik Sheshadri [ECCC:TR26-275]

Sheshadri generalizes his solution-counting method for border determinantal lower bounds from power sums to other polynomials by counting smooth points of the zero set whose tangent hyperplanes pass through a fixed generic point, and also shows limits on how far this method can be pushed.

A short proof of a quadratic lower bound on the border determinantal complexity of the power sum polynomial over C: A simple human exposition Karthik Sheshadri [ECCC:TR26-263]

Sheshadri gives a short proof that the border determinantal complexity of $x_1^n + \cdots + x_n^n$ over $\mathbb{C}$ is at least $(n-1)^2/(4e)$ (and $(n-1)^2/(2e)$ in the symmetric model), and that of $x_1^d + \cdots + x_N^d$ is $\Theta(Nd)$ for $N \geq 3$, $d \geq 2$.

Hardness Amplification for Approximate Nash Equilibrium with Constantly Many Players Xiaotie Deng, Ningyuan Li, Zihe Wang, Jie Zhang [arXiv:2610.25608]

Deng, Li, Wang and Zhang show that for every fixed $\varepsilon \in (0,1)$ and $\eta \in [0,1)$, computing an (ε, η)-weak approximate Nash equilibrium requires quasipolynomial time assuming deterministic PPAD-ETH, even in symmetric binary-payoff games with a constant number of players, via an unconditional error-amplification reduction from bimatrix games.

Zero Knowledge for Free up to Linear Query Bounds: Quasilinear-Size Constant-Query Zero-Knowledge PCPs Inhyuk Park, Chenglin Fan [arXiv:2610.29066]

Park and Fan show that every relation in NTIME(T) has a binary PCP with $O(1)$ non-adaptive queries, perfect completeness, soundness error $1/2$ and length $(T+b)\cdot(\log(T+b))^{O(1)}$ that is perfect zero knowledge against any verifier making at most b queries, improving the previous $\tilde{O}(T+b^2)$ length.

Explicit Reed-Solomon Subspace Designs: Optimal List and Line Decoding Omar Alrabiah, Rohan Goyal, Venkatesan Guruswami [ECCC:TR26-262]

Alrabiah, Goyal and Guruswami give an explicit $n^{O_\varepsilon(1)}$-time construction of Reed–Solomon codes over polynomial-size extension fields that are list-decodable from a $1-R-\varepsilon$ fraction of errors with optimal list size $O((1-R)/\varepsilon)$, plus a simpler strongly explicit construction with list size $(1/\varepsilon)^{O(1/\varepsilon)}$.

Exact Envy-Free Cake Cutting with Polynomial Expected Queries Simon Mackenzie, Mashbat Suzuki [arXiv:2610.28535]

Mackenzie and Suzuki give a randomized protocol for exact envy-free cake cutting among n agents using polynomially many queries in expectation, improving on the previous best deterministic bound of $n^{O(1)}2^n$ due to Ye and Bai.

A Short Note on Algorithms For The Noisy k-XOR Problem Kiril Bangachev [arXiv:2610.27147]

Bangachev's mostly expository note surveys the sample–noise regimes (β, δ) in which noisy k-XOR testing is solvable in time $n^{O(k)}$ and extends the known algorithmic regime using short cycles in hypergraphs.

A Linear Lower Bound on Testing Graph Isomorphism in the Bounded-Degree Model Amit Dror [ECCC:TR26-268]

Dror proves that there are constants d, ε_0 and n-vertex graphs H_n of maximum degree d such that two-sided testing of isomorphism to H_n in the bounded-degree model requires $\Omega(n)$ queries, via a reduction from testing solvability of a fixed sparse linear system over $\mathbb{F}_2$.

Bounds on polynomial approximations of the discrete logarithm function Swastik Kopparty, Rishabh Kothary, Ashutosh Shankar [ECCC:TR26-271]

Kopparty, Kothary and Shankar improve the Coppersmith–Shparlinski bound on the number of points where a degree-D polynomial over $\mathbb{F}_p$ agrees with the discrete logarithm from $\sqrt{2Dp}$ to the Johnson-type bound $(1+o(1))\sqrt{Dp}$, which is nontrivial even for $D > p/2$.

The Thermodynamic Complexity of Circuits Anne Condon, David Doty, Niels Kornerup, David Soloveichik et al. [arXiv:2610.33193]

Condon, Doty, Kornerup, Soloveichik and Wolpert introduce a thermodynamic model of circuits based on heat functions and study its entropy production, showing among other results that computing optimal heat functions is PP-complete.

Bounds on polynomial approximations of the discrete logarithm function Swastik Kopparty, Rishabh Kothary, Ashutosh Shankar [arXiv:2610.33452]

In the arXiv version of their work, Kopparty, Kothary and Shankar show that any univariate polynomial of degree at most D over $\mathbb{F}_p$ agrees with the discrete log on at most $(1+o(1))\sqrt{Dp}$ points of $\mathbb{F}_p^\times$, improving the classical $\sqrt{2Dp}$ bound of Coppersmith and Shparlinski.

Deterministic Root-Finding over Finite Fields with Elliptic Curves Pierre-Jean Spaenlehauer [arXiv:2610.32508]

Spaenlehauer proposes a deterministic root-finding method over $\mathbb{F}_q$ that works in the point groups of an elliptic curve and its quadratic twist instead of $\mathbb{F}_q^\times$, with an offline phase depending only on $\mathbb{F}_q$ that outputs a short preprocessing datum used by the online phase.

Orienting mixed graphs for reachability pairs: When the all pairs setting is easier than the general one Florian Hörsch, Marta Piecyk, Krisztina Szilágyi [arXiv:2610.26451]

Hörsch, Piecyk and Szilágyi study orienting mixed graphs to satisfy as many reachability pairs as possible, and show that the all-pairs variant (APRO) is algorithmically easier than the general fixed-pairs variant (FPRO).

Partitionability is NP-complete Russ Woodroofe [arXiv:2610.27083]

Woodroofe proves that deciding whether a simplicial complex is partitionable is NP-complete in dimensions two and higher, answering a 1995 question of Kleinschmidt and Onn.

Deterministic Shift Equivalence Testing for Sparse Polynomials Markus Bläser, Pranjal Dutta, Sagnik Dutta, Gorav Jindal [arXiv:2610.27168]

Bläser, Dutta, Dutta and Jindal give a deterministic algorithm that decides whether two high-degree sparse polynomials over $\mathbb{Q}$, given with binary-encoded exponents, are equivalent under a shift of variables, and outputs such a shift when one exists (to appear at SODA 2027).

Near-Optimal Bounds for Testing Residual-String Equality and Parenthesis Languages Hadar Strauss [ECCC:TR26-265]

Strauss proves near-optimal query bounds for testing Residual-String Equality and the Dyck parenthesis languages, improving the $\Omega(n^{1/5})$ lower bound of Fischer, Magniez and Starikovskaya.

Derandomizing (Fragments of) Arthur–Merlin with Range Avoidance Zeyong Li [ECCC:TR26-270]

Li characterizes the subclass of prAM problems whose derandomization to NP reduces to Range Avoidance by exhibiting a natural complete problem, Structured Set Lower Bound, which restricts the Goldwasser–Sipser Set Lower Bound problem.

Sustainable and Envy-free Multi-Period Allocation Robert Bredereck, Eva Deltl, Till Fluschnik [arXiv:2610.26032]

Bredereck, Deltl and Fluschnik introduce envy-free allocation of renewable resources over multiple periods, subject to viability and capacity constraints, and study its classical and parameterized complexity with respect to the numbers of agents and resources, among other parameters.

On Testing Finite Groups Nader Bshouty [ECCC:TR26-264]

Bshouty studies randomized adaptive testers for whether a binary operation on $[n]$, given by oracle access to its multiplication table, is a group; one of the paper's three results is a tester using $O(n+1/\epsilon)$ queries.

Space Complexity of Reachability in Simple Path Graphs

Krishnamoorthy Dinesh, Chandana Sasidharan [arXiv:2610.27879]

Dinesh and Sasidharan study the space complexity of reachability in restricted directed graph families, including simple path graphs, relative to unambiguous logspace (UL) and reach-unambiguous logspace (RUL).

Space Complexity of Reachability in Simple Path Graphs

Krishnamoorthy Dinesh, Chandana Sasidharan [ECCC:TR26-280]

This is the ECCC version of Dinesh and Sasidharan's study of reachability in simple path graphs and other restricted directed graph families, relative to UL and reach-unambiguous logspace RUL.

Temporal Exploration and Rendezvous via Weisfeiler-Leman Colorings

Frank Kammer, Sandra Kiefer, Johannes Meintrup [arXiv:2610.30123]

Kammer, Kiefer and Meintrup show that every connected n-vertex temporal graph can be explored in $O(qn \log^3 n)$ time steps, where q is the number of 2-dimensional Weisfeiler–Leman colour classes; this improves on the general $\Theta(n^2)$ bound when q is small, and the walk can be constructed in deterministic polynomial time.

Controllability of polynomial systems is undecidable and not recursively enumerable

Jun Liu, Ivan O. Shevchenko [arXiv:2610.26832]

Liu and Shevchenko prove that global controllability of polynomial control systems with rational coefficients is undecidable, confirming the first part of Sontag's 1988 conjecture, and disprove its second part by showing the controllable systems are not recursively enumerable.

On the Support Geometry and Complexity of LABS

William Cruz-Santos [arXiv:2610.27329]

Cruz-Santos shows that unary standard LABS is in NP and that its Karp NP-hardness would imply P = NP, and finds that zero-energy feasibility for the Pair-Set and Masked variants is polynomial-time decidable when pair groups or masked rows are sufficiently small.

A 55-Addition Rank-23 Scheme for 3×3 Matrix Multiplication via Exact Two-Sided Minimization

Greg Sidebottom [arXiv:2610.27451]

Sidebottom gives a Lean-verified scheme that multiplies two 3×3 matrices over any ring with 23 multiplications and 55 additions, one fewer addition than the previous record of 56, found by exact two-sided minimization rather than greedy common-subexpression elimination.

Subjective Valuations Make It Hard to be Fair: The Complexity of Resource Allocation to Groups of Agents

Argyrios Deligkas, Eduard Eiben, Robert Ganian, Tiger-Lily Goldsmith et al. [arXiv:2610.28213]

Deligkas, Eiben, Ganian, Goldsmith, Gründel and Wietheger study the parameterized complexity of finding EF1 allocations of indivisible items to groups of agents, with fixed and variable group assignments, and show that the two variants, though both NP-hard, behave quite differently algorithmically.

Stable Set Reconfiguration in Dense Digraphs

Sarah Houdaigoui, Nacim Oijid, Matthieu Petiteau, Alexandre Vigny [arXiv:2610.32402]

Houdaigoui, Oijid, Petiteau and Vigny study reconfiguration of acyclic sets in digraphs under directed token sliding and give an FPT algorithm for semi-complete digraphs, parameterized by the size of the sets together with a further structural parameter.

Small sums of roots of unity: polylogarithmic exceptional sets and conditional separation

Qi Cheng [arXiv:2610.27231]

Cheng proves unconditionally that at most $O((\log 2N)^{11})$ points of the N-grid make the five-term sum $|1 + 2\cos(2\pi x) + 2\cos(2\pi y)|$ smaller than $\exp(-C(\log 2N)^{13})$, and also derives a conditional separation result.

Near-Quadratic Time Multivariate Polynomial Interpolation

Chandra Kanta Mohapatra [arXiv:2610.30907]

Mohapatra gives a deterministic near-quadratic-time algorithm for multivariate polynomial interpolation over an arbitrary finite field in the case $N = d^m$, where N is the number of points and d bounds the individual degree.

On the complexity of Medvedev reducibility

Remi Pallen [arXiv:2610.32522]

Pallen studies the complexity of deciding whether a given effective closed set has the same Medvedev degree as a fixed one, a naturally Σ^0_3 problem, and shows it attains this maximal complexity for every degree except the maximal one.

Structure and Exact Certification of Nonnegative Tensor Decompositions

Haoming Wang, Ming Yuan [arXiv:2610.32787]

Wang and Yuan analyse nonnegative tensor decompositions given by rational factor matrices, showing that every failure of minimality or essential uniqueness contains an irreducible positive exchange of bounded effective order, and use this to give exact algorithms for nonnegative tensor rank and for certifying minimality and uniqueness, running in time $G(R)L^{O(R^2)}$ with an exponent independent of tensor order and mode dimensions.

A Borsuk-Ulam theorem for adaptive continuous measurements

Mario Ullrich [arXiv:2610.33781]

Ullrich proves an adaptive Borsuk–Ulam theorem: if $2^n \le m$, any n adaptively chosen continuous functionals agree at two antipodal points of any sphere in $\mathbb{R}^m$, so at least $\lceil \log_2(m+1) \rceil$ adaptive continuous measurements are needed to learn vectors of a ball to error below its radius, matching the known upper bound up to one measurement; the main proof idea is credited to an AI model.

Learning Positive Linear Combinations of ReLUs in Fully Polynomial Time

Ilias Diakonikolas, Jelena Diakonikolas, Puqian Wang, Nikos Zarifis [ECCC:TR26-276]

Diakonikolas, Diakonikolas, Wang and Zarifis give the first algorithm that PAC learns positive linear combinations of k ReLUs on $\mathbb{R}^d$ under the Gaussian distribution in L_2 loss using $\mathrm{poly}(d, k, 1/\epsilon)$ samples and time.

Internalizing Chain-of-thought into Looped Transformers

Kehang Zhou, Pinyan Lu [arXiv:2610.25998]

Zhou and Lu construct constant-bit-size looped transformers that interpolate between chain-of-thought and latent recurrence: for loop count L they simulate a time-$t(n)$, space-$s(n)$ Turing machine with an $O(s(n)/L)$ context window and $O(t(n)s(n)^{1/k}/L)$ generated tokens.

Optimal Rényi and Logarithmic Additive Functional Entropy Estimation

Arman Adibi, Piotr Krysta [arXiv:2610.28999]

Adibi and Krysta determine the minimax sample complexity of estimating Rényi entropy of order $\alpha > 1$ on alphabets of size k up to constants independent of k and α, uniformly even when α varies with k, and extend this to logarithms of additive functionals.

Random Clocks on Computation Paths]{Random Clocks on Computation Paths: How Far Can They Reach, and When Does Stepping Back Help?

Farzad Didehvar, Saman Farhat [arXiv:2610.25830]

Didehvar and Farhat study a model in which a computation is observed only at times chosen by a random clock that may leap ahead, pause or step back, showing that rare long leaps reach exponential and doubly exponential time with infinitely many intermediate levels, and that stepping back can help decisively for some clocks.

Bidirectionality is Necessary in Cellular Machines for Exact-Arrival Multi-Agent Path Finding Masashi Aono [arXiv:2610.26272]

Aono proves that, on an explicit hard topology, cellular machines with only unidirectional local communication cannot solve exact-arrival multi-agent path finding as a distributed output task, so bidirectional communication is necessary; part of the argument is checked in Lean 4.

Singular Linear Differential-Algebraic Equations with Constant Coefficients Necessarily Admit Non-Computable Solutions Holger Boche, Volker Pohl, H. Vincent Poor, Frank H. P. Fitzek [arXiv:2610.28188]

Boche, Pohl, Poor and Fitzek show that every singular linear differential-algebraic equation with constant computable coefficients has some computable, highly regular input for which the unique solution is not Turing computable, in contrast with the regular case.

FROM THE AI-ASSISTED REPOSITORIES

Work deposited with AI systems credited as authors or co-authors, not screened like arXiv or ECCC. Read with care.

Deciding D-stability is coNP-hard AI-ASSISTED Michael Mislan [hexagon:2610.00368v1]

An AI-assisted, unverified submission claims that deciding matrix D-stability is coNP-hard, via a reduction from PARTITION to integer matrices of order $m + 4$ that are D-stable exactly when no subset has half the total weight, with Pell-recursion parameters ensuring polynomial output length.

Polynomial Fault Detection Without Bounded Width: Exact Reachable Sets in Priority Controllers AI-ASSISTED Idris Ali Shaik [hexagon:2610.00503v1]

This AI-assisted submission by Shaik studies test-pattern generation for an arbitrary-size extension of the c432 priority controller, characterising exactly the qualifier sets reachable under a specified first-stage bridge fault and constructing a primary input realising each one, so that the polynomial-time cases are explained without relying on bounded circuit width.

Unconstrained detection of productive autocatalytic cores is NP-complete AI-ASSISTED Michael Mislan [hexagon:2610.00390v1]

This AI-assisted submission by Mislan claims that deciding whether a finite reversible reaction network contains a productive autocatalytic core is NP-complete when no target species or food set is specified, via a reduction from SAT through two-path directed linkage.

An Expansion on Paraconsistent Micro-Geometry: A Non-Standard Approach to Subgroup-Stratified Vitali Measures via the Logic of da Costa AI-ASSISTED George Yeats Lazou · with Google: Gemini 3.5 Flash-Lite [diderot:57fe6f5362dc]

This AI-assisted submission by Lazou proposes an explicit bottom-up construction, based on da Costa's paraconsistent logic and nonstandard analysis, that assigns infinitesimal inner and outer measures to Vitali sets; it is not a complexity-theory paper.

ODDS AND ENDS

NUMBERS OF THE WEEK

$4n - O(\log n)$
EXPLICIT B_2 CIRCUIT SIZE LOWER BOUND (MISHRA–SANTHANAM)

$\Omega(N \log N)$
OPTIMAL NONDETERMINISTIC COMMUNICATION FOR GRAPH CONNECTIVITY

42
ALPHABET SIZE WHERE GILBERT–VARSHAMOV IS BEATEN

1979
YEAR THE MANY-ONE MEE QUESTION WAS ALREADY OPEN

$1/6$
TIGHT APPROXIMATION THRESHOLD FOR ε-GCIRCUIT

$m + 4$
MATRIX ORDER IN D-STABILITY REDUCTION (AI-ASSISTED)

COMPLEXITY FORECAST

Next week's outlook: a high-pressure front of gate elimination moves in from the East, with lower bounds rising steadily toward $4.5n$. Sources close to the forecast warn this is a joke, not a prediction. Scattered parallel repetition is expected across all k-player regions, decaying exponentially by the weekend. AI-generated proofs remain likely in the afternoon, so referees should carry an umbrella and a Lean installation.

CLASSIFIEDS

- WANTED: One more gate to eliminate. Must be explicit and in P. Contact Box $4n$.
- FOR SALE: Spanning forest, gently transmitted, only $O(N \log N)$ bits. Nondeterministic buyers need not haggle.
- LOST: Alphabet of size 46. Answers to Zinoviev–Litsyn. Finder may keep the four spare letters.
- SEEKING: Human reader willing to verify proof by GPT-Sol-6.1. Must understand Σ_2^p and have a strong stomach.
- ROOMMATE WANTED: Quiet k-player XOR game, value strictly below 1. Will repeat in parallel and decay politely.
- EXCHANGE: Will trade one propeller (three 120° sectors, optimal) for a matrix multiplication exponent below $9/4$.
- CAKE AVAILABLE: Divided envy-free. Expected queries polynomial, but please budget for $\tilde{O}(n^{30})$.
- PROVERS NEEDED: Bounded arithmetic theory seeks algebraic lower bounds to testify. Must reason at NC^2 or above.

Complexity Chronicles is written by a language model (anthropic: claude-opus-5-5) from the abstracts of 61 papers (arXiv cs.CC, ECCC, Hexagon, Diderot) that appeared during 19–25 October 2026, selected according to a reader profile. Generated 2026-11-05 16:02 UTC. Headlines are tabloid; standfirsts are meant to be accurate; nothing here replaces reading the paper. Typeset with LuaLaTeX in Merriweather, Montserrat, Playfair Display, Roboto Condensed and Fira Mono.