Electronic Colloquium on Computational Complexity, Report No. 106 (2023)

Depth-3 Circuits for Inner Product

Mika Göös
EPFL, Switzerland

Ziyi Guan
EPFL, Switzerland

Tiberiu Mosnoi
EPFL, Switzerland

Abstract

What is the Σ_3^2-circuit complexity (depth 3, bottom-fanin 2) of the $2n$-bit inner product function? The complexity is known to be exponential $2^{\alpha_n n}$ for some $\alpha_n = \Omega(1)$. We show that the limiting constant $\alpha := \limsup \alpha_n$ satisfies

$$0.847... \;\leq\; \alpha \;\leq\; 0.965... \,.$$

Determining α is one of the seemingly-simplest open problems about depth-3 circuits. The question was recently raised by Golovnev, Kulikov, and Williams (ITCS 2021) and Frankl, Gryaznov, and Talebanfard (ITCS 2022), who observed that $\alpha \in [0.5, 1]$. To obtain our improved bounds, we analyse a covering LP that captures the Σ_3^2-complexity up to polynomial factors. In particular, our lower bound is proved by constructing a feasible solution to the dual LP.

2012 ACM Subject Classification Theory of computation → Circuit complexity

Keywords and phrases Circuit complexity, inner product

1 Introduction

A Σ_3-*circuit* is an unbounded-fanin depth-3 boolean circuit with an $\vee$-gate at the top. That is, the circuit computes an OR of CNFs. A foremost open problem in circuit complexity is to prove a lower bound of $2^{\omega(\sqrt{n})}$ on the Σ_3-circuit complexity of an explicit n-bit boolean function. Current techniques can prove at best a bound of $2^{\Omega(\sqrt{n})}$ [7, §11].

For the more restricted class of Σ_3^k-*circuits* that have fanin k at the bottom (that is, ORs of k-CNFs), we can hope for improved bounds. For example, the famous satisfiability coding lemma [14] implies that the n-bit parity function has Σ_3^k-circuit complexity at least $2^{n/k}$ and this is tight up to polynomial factors (for constant k). Even stronger, for $k = 2$, Paturi, Saks, and Zane [12] exhibit a function with near-maximal Σ_3^2-complexity $2^{n-o(n)}$. No such near-maximal lower bounds are currently known for $k = 3$.

Inner product. A natural function whose Σ_3^k-complexity remains unknown (up to poly(n) factors) is the *inner product* function IP_n, defined on $2n$-bit inputs $(x, y) \in (\{0,1\}^n)^2$ by

$$\mathrm{IP}_n(x, y) \;:=\; \langle x, y\rangle \bmod 2.$$

Recently, Golovnev, Kulikov, and Williams [2] asked to determine the Σ_3^k-complexity of IP_n in case $k = 3$. Curiously enough, Frankl, Gryaznov, and Talebanfard [1] point out that the problem is nontrivial already in case $k = 2$, and they obtained partial results towards resolving it. It has been known that the Σ_3^2-complexity of IP_n is between $2^{n/2}$ and 2^n [14, 2].

1.1 Our result

Our main result is to prove improved upper and lower bounds for inner product.

ISSN 1433-8092

▶ **Theorem 1** (Main result). *Write the* Σ_3^2*-complexity of* IP_n *as* $2^{\alpha_n n}$ *for some* $\alpha_n \geq 0$*. Then*

$$\alpha \;\coloneqq\; \limsup \alpha_n \;\in\; [0.847..., 0.965...].$$

It remains an intriguing problem to determine α precisely. It is surprising (for us, at least) that neither of the previous bounds $\alpha \in [0.5, 1]$ were tight, especially because the problem is seemingly one of the simplest open questions about depth-3 circuits.

Studying exact exponents of Σ_3^k-circuit complexities is a relatively unexplored research direction, and we believe it could foster the development of new lower bound techniques. In particular, a major motivation for this comes from *depth reduction* results. For example, in case $k = 16$, Golovnev, Kulikov, and Williams [2] have shown that proving near-maximal $2^{n-o(n)}$ bounds for Σ_3^{16}-circuits would already yield new improved lower bounds for *unrestricted* (unbounded depth) circuits. Their result extends classical connections discovered by Valiant [15]; see also the monograph [16, §3].

1.2 Overview of techniques

To obtain our improved bounds on α in Theorem 1—both upper and lower bounds—we study a fractional covering problem, formulated as a linear program (LP), that captures the Σ_3^2-circuit complexity up to $\mathrm{poly}(n)$ factors.

To our knowledge, LPs have not been widely employed in analysing depth-3 circuits. They are, however, routinely used to prove strong lower bounds in the related area of *communication complexity* [9]. Many such LP-based methods are catalogued by Jain and Klauck [6]. Moreover, Lee and Shraibman [10] give a monograph-length treatment on how to use LP duality to prove communication lower bounds. In one of the earliest examples, Karchmer, Kushilevitz, and Nisan [8] characterised nondeterministic communication complexity via a fractional covering problem. The formulation we use is a straightforward adaptation of this for depth-3 circuits. A similar formulation also appeared in the work of Hirahara [4] that connects depth-3 complexity with one-sided CNF approximations.

Covering LP. The size of a Σ_3^2-circuit is determined (up to $O(n^2)$ factors) by the fanin of the top $\vee$-gate. Suppose a circuit with top-fanin m computes a function $f\colon \{0,1\}^n \to \{0,1\}$. We can view the circuit as expressing the set of 1-inputs $f^{-1}(1)$ as a union of m sets,

$$f^{-1}(1) \;=\; \bigcup_{i \in [m]} \phi_i^{-1}(1), \tag{1}$$

where each $\phi_i^{-1}(1)$ is the set of inputs accepted by a 2-CNF formula ϕ_i. The least top-fanin needed to compute f is then captured by the optimal *integer solutions* to the following covering LP. In this LP, we assign a fractional weight $w_\phi \in [0,1]$ for each 2-CNF ϕ that is *consistent* with f, meaning that $\phi(x) \leq f(x)$ for every input $x \in \{0,1\}^n$. We let Φ denote the set of all 2-CNFs consistent with f.

$$\begin{array}{rlll} \min & \sum_{\phi\in\Phi} w_\phi & & \\ \textit{subject to} & \sum_{\phi\in\Phi} w_\phi \phi(x) \;\geq\; 1, & \forall x \in f^{-1}(1) & \\ & w_\phi \in [0,1], & \forall \phi \in \Phi & \end{array} \tag{LP}$$

A classic result of Lovász [11] says that the integrality gap of a covering LP is small.

▶ **Lemma 2** (Lovász [11]). *Let* Opt *and* $\mathsf{Opt}^{\mathbb{Z}}$ *denote the value of* (LP) *optimised over fractional solutions* ($w_\phi \in [0,1]$) *and integral solutions* ($w_\phi \in \{0,1\}$)*, respectively. Then*

$$\mathsf{Opt} \;\le\; \mathsf{Opt}^{\mathbb{Z}} \;\le\; O(n)\cdot \mathsf{Opt}.$$

Consequently, to determine the Σ_3^2-complexity of $f = \mathrm{IP}_n$ we only need to solve the fractional (LP). We will use the (LP) in Section 2 to construct circuits for IP_n that witness the upper bound $\alpha \le 0.965\ldots$.

Dual LP. A common method to prove a depth-3 lower bound is to estimate the number of accepting inputs for any consistent CNF, say, by $\max_{\phi\in\Phi} |\phi^{-1}(1)| \le C$, and then conclude that the top-fanin must be at least $|f^{-1}(1)|/C$. Such arguments are standard in the *top-down* circuit lower bound literature [3, 14, 12, 13, 5].

An important generalisation of this method is to first choose a hard distribution $\mathcal{D}$ over the 1-inputs $f^{-1}(1)$ and then measure the size of $\phi^{-1}(1)$ relative to $\mathcal{D}$. If we can show $\max_{\phi\in\Phi} \Pr_{x\sim\mathcal{D}}[\phi(x) = 1] \le p$, then the top-fanin must be at least $1/p$. Indeed, the following optimisation problem captures the best lower bound provable with this method.

$$\begin{array}{rlll} \max & 1/p & & \\ \textit{subject to} & \sum_{x\in f^{-1}(1)} \mathcal{D}(x)\phi(x) \;\le\; p, & \forall \phi \in \Phi & \\ & \sum_{x\in f^{-1}(1)} \mathcal{D}(x) \;=\; 1, & & \text{(Dual LP)} \\ & \mathcal{D}(x) \in [0,1], & \forall x \in f^{-1}(1) & \end{array}$$

This program is not written in standard LP format as we are seemingly optimising a nonlinear function. However, it is equivalent[1] to $\max \sum_x A(x)$ s.t. $\sum_x A(x)\phi(x) \le 1$ and $A(x) \ge 0$, which is the canonical dual to (LP). Hence, by strong duality, we can always prove a tight lower bound (up to polynomial factors) on depth-3 complexity by finding the right hard distribution $\mathcal{D}$.

Hard distribution for IP. What hard distribution $\mathcal{D}$ should we choose to prove a strong lower bound for IP_n? If we choose $\mathcal{D}$ to be the uniform distribution over $\mathrm{IP}_n^{-1}(1)$, then prior work [1, Thm 28] showed that this only yields the bound $\alpha \ge \log \frac{4}{3} = 0.415\ldots$. If we choose $\mathcal{D}$ by sampling a pair $(x, 1^n)$ where x is uniform random in $\{0,1\}^n$, then we have effectively reduced IP_n to n-bit parity and we obtain $\alpha \ge 0.5$ [2], which is tight for parity.

To get our improved lower bound on α, we analyse a more general distribution.

(Section 3) We consider a distribution where the $2n$ input bits are *iid*, that is, $\mathcal{D}$ is the binomial distribution with some parameter $p \in (0,1)$. (Note that while $\mathcal{D}$ is not supported on $\mathrm{IP}_n^{-1}(1)$ it does place a constant probability mass on it.) We prove a structure lemma for consistent 2-CNFs and characterise those that have the highest acceptance probability under $\mathcal{D}$. Optimising the choice of p, we will obtain $\alpha \ge \log \frac{9}{5} = 0.847\ldots$.

---

[1] If $\mathcal{D}, p$ is feasible for (Dual LP), then $A(x) \coloneqq \mathcal{D}(x)/p$ is feasible and has the same objective function value in the other program. In the other direction, set $p \coloneqq 1/\sum_y A(y)$ and $\mathcal{D}(x) \coloneqq p \cdot A(x)$.

1.3 Discussion and open problems

The challenge in proving a better lower bound in Theorem 1 is that our techniques rely heavily on the hard distribution having independence between the n coordinates. One way we could try to improve the lower bound is to consider a slightly more general *coordinate-wise iid* distribution. That is, we choose a distribution μ over one coordinate pair $(x_i, y_i) \in \{0,1\}^2$ and then define a product distribution by $\mathcal{D} \coloneqq \mu^n \coloneqq \mu \times \cdots \times \mu$. We carried out this approach (using computer-aided calculations) only to find out that we get no improvement this way: the hardest $\mathcal{D}$ is still the bit-wise *iid* that we consider in Section 3. A candidate for the absolute hardest distribution (not necessarily coordinate-wise *iid*) is merely a *symmetric* distribution that is invariant under permuting the n coordinates. We leave it as an open problem to analyse such non-*iid* distributions.

Another open problem that could be amenable to an LP-based attack is to determine the Σ_3^k-circuit complexity of inner product in case $k = 3$, as was originally asked by Golovnev, Kulikov, and Williams [2]. The best lower bound known is $2^{n/3}$ [14], and one could hope to show an improved lower bound even relative to an *iid* distribution. Here the obvious challenge is that 3-CNFs are notoriously much more difficult (even NP-hard) to analyse than 2-CNFs. Our overall approach in this paper is still applicable even for $k > 2$. Namely, one needs to "merely" prove an analogue of our structure lemma (Lemma 7) for k-CNFs.

2 Upper bound

In this section, we prove the upper bound $\alpha \le 0.965...$ as claimed in Theorem 1. The circuit will be constructed in two parts. To explain this, we denote, for an input $(x, y) \in \{0,1\}^{2n}$ and a 2-bit pattern $s \in \{0,1\}^2$, the fraction of occurrences of this pattern by

$$p_s(x, y) \;\coloneqq\; \tfrac{1}{n}\,|\{i \in [n] \colon (x_i, y_i) = s\}|\,.$$

We use one Σ_3^2-circuit to accept every input $(x, y) \in \mathrm{IP}_n^{-1}(1)$ with $p_{11}(x, y) \le p$ where p is a carefully chosen threshold, and another Σ_3^2-circuit to accept those inputs with $p_{11}(x, y) \ge p$.

The following two lemmas (proved in Sections 2.1 and 2.2) record the two types of circuits we will construct. To state these lemmas, recall that a circuit C is *consistent* with IP_n if $C(x, y) \le \mathrm{IP}_n(x, y)$ for all inputs (x, y). We let $\mathrm{H}(p) \coloneqq -p \log p - (1-p) \log(1-p)$ denote the binary entropy function. Moreover, we let $\mathbb{H}(X)$ denote the usual Shannon entropy of a random variable X. Finally, for $p \in [0, 1]$, we define a random variable $X_p \in \{0,1\}^2$ such that $\Pr[X_p = 11] = p$ and $\Pr[X_p = s] = (1-p)/3$ for $s \in \{00, 01, 10\}$.

▶ **Lemma 3.** *For every $p \in [0, \frac{1}{2}]$ there exists a Σ_3^2-circuit of size $2^{n\mathrm{H}(p)+o(n)}$ that is consistent with IP_n and that accepts all $(x, y) \in \mathrm{IP}^{-1}(1)$ with $p_{11}(x, y) \le p$.*

▶ **Lemma 4.** *For every $p \in [\frac{1}{4}, 1]$ there exists a Σ_3^2-circuit of size $2^{\frac{1}{2}n\mathbb{H}(X_p)+o(n)}$ that is consistent with IP_n and that accepts all $(x, y) \in \mathrm{IP}^{-1}(1)$ with $p_{11}(x, y) \ge p$.*

The final Σ_3^2-circuit for IP_n is the OR of the two Σ_3^2-circuits above. It is easy to see that using any constant $p \in (\frac{1}{4}, \frac{1}{2})$ we get a circuit of size $2^{\beta n}$ with $\beta < 1$. We can further optimise the choice of p by equating the two circuit size expressions, solving for p numerically (using any numerical computation software), which comes to $p \coloneqq 0.3909...$, and then plugging this value of p into the size expressions to get a circuit of size $2^{0.965\ldots \cdot n+o(n)}$, as desired.

It remains to prove Lemmas 3 and 4, which we do in the rest of this section.

2.1 Proof of Lemma 3

In this lemma we focus on finding efficient Σ_3^2-circuits accepting inputs $(x,y) \in \mathrm{IP}^{-1}(1)$ with a small value of $p_{11}(x,y) \le p \le 1/2$. Given a subset $I \subseteq [n]$, define the *brute-force CNF* by

$$\phi_{\mathrm{BF}}^{(I)} \;:=\; \bigwedge_{i\in I}(x_i \wedge y_i) \wedge \bigwedge_{i\in[n]\setminus I}(\neg x_i \vee \neg y_i).$$

Note that $\phi_{\mathrm{BF}}^{(I)}$ accepts an input (x,y) iff I equals the set of all i such that $(x_i,y_i)=(1,1)$. Hence, to accept every input with $p_{11}(x,y)\le p$, our Σ_3^2-circuit will consider all suitable I:

$$C \;:=\; \bigvee_{\substack{I\subseteq[n]\\ |I|\le pn\\ |I|\text{ odd}}} \phi_{\mathrm{BF}}^{(I)} \;. \tag{2}$$

The size of C is at most $\binom{n}{\le pn}\cdot O(n)$ where $\binom{n}{\le pn} := \sum_{i=0}^{pn}\binom{n}{i}$ can be estimated from above via Stirling's approximation by $2^{n\mathrm{H}(p)+o(n)}$ for all $p \le 1/2$. Finally, it is clear from the construction that C is consistent with IP_n. This concludes the proof of Lemma 3. ◀

2.2 Proof of Lemma 4

In this lemma we focus on finding efficient Σ_3^2-circuits accepting inputs $(x,y)\in\mathrm{IP}_n^{-1}(1)$ with a large value of $p_{11}(x,y)\ge p\ge 1/4$. To illustrate our idea, we first construct a circuit for a simpler related function, and then explain how to modify it to get circuits for IP_n.

Simple warm-up circuit. We first describe a circuit that computes the following partial function (which is consistent with $\neg\mathrm{IP}_n$, but we will address this later):

$$f_n(x,y) \;:=\; \begin{cases} 0 & \text{if } n\cdot p_{11}(x,y) \text{ is odd,}\\ 1 & \text{if } n\cdot p_s(x,y) \text{ is even for all } s\in\{0,1\}^2, \text{ and } p_{11}(x,y)\ge p,\\ * & \text{otherwise.}\end{cases}$$

The interesting case here is when n is even, as otherwise $f_n(x,y)\in\{0,*\}$ for all (x,y). Let $M\subseteq\binom{[n]}{2} := \{e\subseteq[n] : |e|=2\}$ be a *perfect matching* of $[n]$ (that is, partition of $[n]$ into pairs). We define the *collision CNF* associated with M by

$$\phi_{\mathrm{Coll}}^{(M)} \;:=\; \bigwedge_{\{i,j\}\in M} (x_i\leftrightarrow x_j)\wedge(y_i\leftrightarrow y_j).$$

This is a 2-CNF since we can write an equivalence as $a\leftrightarrow b \equiv (a\vee\neg b)\wedge(\neg a\vee b)$. Note that a collision CNF accepts iff for every pair $\{i,j\}\in M$ we have $(x_i,y_i)=(x_j,y_j)$. Hence it only accepts inputs where $n\cdot p_s(x,y)$ is even for all $s\in\{0,1\}^2$. Thus $\phi_{\mathrm{Coll}}^{(M)}$ is consistent with f_n.

To construct a Σ_3^2-circuit for f_n, it is enough, as discussed in Section 1.2, to design a feasible solution to the (LP) associated with f_n. (We note that the (LP) formulation works equally well for partial functions.) To this end, we calculate in the following claim (proved in Section 2.3) the probability that a *random* collision CNF accepts a fixed 1-input of f_n.

▷ Claim 5. Let $(x,y)\in f_n^{-1}(1)$. For a uniformly chosen perfect matching $M\subseteq\binom{[n]}{2}$,

$$\Pr_M\left[\phi_{\mathrm{Coll}}^{(M)}(x,y)=1\right] \;\ge\; 2^{-\frac{1}{2}n\mathbb{H}(X_p)-o(n)} \;=:\; L(p).$$

We now construct a feasible solution to (LP) for f_n. Let Φ_{Coll} denote the set of all collision CNFs, one for each perfect matching of $[n]$. Consider the weight assignment corresponding to the uniform distribution over Φ_{Coll}; namely, set $w_\phi \coloneqq 1/|\Phi_{\text{Coll}}|$ for every $\phi \in \Phi_{\text{Coll}}$ and $w_\phi \coloneqq 0$ for all the rest. Note that the objective function value is $\sum_\phi w_\phi = 1$. However, the assignment may not be feasible: for a covering constraint indexed by $(x, y) \in f_n^{-1}(1)$, we are only guaranteed a weak lower bound (much smaller than 1):

$$\textstyle\sum_\phi w_\phi \phi(x,y) \;=\; \Pr_M \left[\phi_{\text{Coll}}^{(M)}(x,y) = 1\right] \;\geq\; L(p).$$

We can, however, transform this weight assignment into a feasible one by scaling all the weights up by a factor of $1/L(p)$ (and truncating any resulting weight > 1 to 1). In the scaled assignment, the objective function value is at most $1/L(p)$. We conclude (using Lemma 2) that f_n has a circuit of size $O(n)/L(p) = 2^{\frac{1}{2}n\mathbb{H}(X_p)+o(n)}$.

It remains to explain how a circuit of this size can also be constructed for IP_n.

Actual circuit for IP. To prove Lemma 4, we would like to use the Σ_3^2-circuit we constructed above for f_n to design a circuit for the partial function

$$\text{IP}_n^{(p)}(x,y) \;\coloneqq\; \begin{cases} 0 & \text{if } n \cdot p_{11}(x,y) \text{ is even,} \\ 1 & \text{if } n \cdot p_{11}(x,y) \text{ is odd, and } p_{11}(x,y) \geq p, \\ * & \text{otherwise.} \end{cases}$$

Consider the following nondeterministic algorithm for $\text{IP}_n^{(p)}$. On input $(x, y) \in \{0, 1\}^{2n}$:

1. Nondeterministically guess a subset $S \subseteq \{0, 1\}^2$ where $11 \in S$. The intention is that patterns in S should appear in (x, y) an odd number of times.
2. For each $s \in S$, guess a coordinate $i(s) \in [n]$.
3. For each $s \in S$, check that $(x_{i(s)}, y_{i(s)}) = s$. If not, reject.
4. Output the same as the function $f_{n-|S|}$ on input $(x_i, y_i)_{i\in[n]\setminus i(S)}$.

It is straightforward to check that this computes $\text{IP}_n^{(p)}$ correctly. (A minor technical detail is that when computing $f_{n-|S|}$, the p_{11} value may slightly drop because we remove one occurrence of the 11-pattern. However, this is not really a problem since the slight drop will not affect the asymptotics of the circuit size.) The question remains: How can it be implemented as a Σ_3^2-circuit? We do it as follows. Consider any guess outcome $O \coloneqq (S, (i(s))_{s\in S})$. We can modify the circuit C for $f_{n-|S|}$ (applied to the input bits $(x_i, y_i)_{i\in[n]\setminus i(S)}$) to perform the check in Item 3 by adding to each 2-CNF in C the singleton terms $(x_{i(s)} = s_1)$ and $(y_{i(s)} = s_2)$ for all $s = (s_1, s_2) \in S$. Call the resulting circuit C_O. Our final Σ_3^2-circuit computes the OR of all circuits C_O. Since there are only $O(n^4)$ many different guess outcomes, the resulting circuit is only a factor $O(n^4)$ larger than our circuit for f_n. This concludes the proof of Lemma 4. ◀

2.3 Proof of Claim 5

Write $n!! \coloneqq \prod_{i=0}^{\lfloor n/2 \rfloor}(n - 2i)$ for the double factorial. The number of perfect matchings on $[n]$ is well-known to be given by $(n-1)!!$ when n is even. Therefore, $(np_s - 1)!!$ gives the number of ways to match the coordinates with pattern s. We have

$$\Pr_M \left[\phi_{\text{Coll}}^{(M)}(x,y) = 1\right] \;=\; \frac{\prod_{s\in\{0,1\}^2}(np_s - 1)!!}{(n-1)!!}. \tag{3}$$

Taking logarithms and using Stirling's approximation ($\log n!! = \frac{1}{2}n\log n - \frac{1}{2}n \pm o(n)$) we get

$$\begin{aligned}\log \Pr_M\left[\phi_{\text{Coll}}^{(M)}(x,y)=1\right] &= \tfrac{1}{2}\textstyle\sum_s np_s\log(np_s) - \tfrac{1}{2}n\log n \pm o(n)\\ &= \tfrac{1}{2}n\cdot \textstyle\sum_s p_s \log p_s \pm o(n)\\ &= -\tfrac{1}{2}n\cdot \mathbb{H}(P)\pm o(n).\end{aligned}$$

Here $P \in \{0,1\}^2$ is the random variable defined by $\Pr[P = s] = p_s$. We ask: which random variable $X \in \{0,1\}^2$ maximises the entropy $\mathbb{H}(X)$ subject to the constraint $\Pr[X = 11] = p^*$? By the concavity of $\mathbb{H}$ and symmetry (we can relabel outcomes without affecting the entropy), it is the random variable X_{p^*} such that

$$\Pr[X_{p^*} = 11] = p^*, \qquad \Pr[X_{p^*} = 00] = \Pr[X_{p^*} = 10] = \Pr[X_{p^*} = 01] = (1-p^*)/3.$$

The univariate map $p^* \mapsto \mathbb{H}(X_{p^*})$ is also concave. It is maximised at $p^* = 1/4$ (when X_{p^*} is uniform), and decreasing for $p^* > 1/4$. This means that, since $1/4 \le p \le p_{11}$, we have that $\mathbb{H}(X_p) \ge \mathbb{H}(X_{p_{11}}) \ge \mathbb{H}(P)$. Hence we obtain the claimed lower bound:

$$\log \Pr_M\left[\phi_{\text{Coll}}^{(M)}(x,y)=1\right] \;\ge\; -\tfrac{1}{2}n\cdot\mathbb{H}(X_p) - o(n). \qquad ◀$$

3 Lower bound

In this section, we prove the lower bound $\alpha \ge \log\frac{9}{5} = 0.847...$ as claimed in Theorem 1. We will follow the Dual LP strategy discussed in Section 1.2. Namely, we will choose a hard distribution over $\text{IP}_n^{-1}(1)$ and then bound the acceptance probability of any 2-CNF consistent with IP_n. In fact, it is convenient to prove a slightly stronger statement and bound the acceptance probability of any 2-CNF consistent with IP_n or $\neg\text{IP}_n$. Indeed, we let Φ_n denote the set of 2-CNFs consistent with IP_n or $\neg\text{IP}_n$.

Hard distribution. As the hard distribution, we consider the binomial distribution $\mathcal{D}_p$ with parameter $p \in (0,1)$, whose choice we will optimise later. That is, $(X,Y) \sim \mathcal{D}_p$ is such that all bits are *iid*: they are independent and have identical distribution, $\Pr[X_i = 1] = \Pr[Y_i = 1] = p$. Note that $\mathcal{D}_p$ is not in fact supported on $\text{IP}_n^{-1}(1)$, but it still places $\Omega(1)$ probability mass on this set. Consequently, any Σ_3^2-circuit will have to cover $\Omega(1)$ fraction of $\mathcal{D}_p$ with its CNFs, so we can still use $\mathcal{D}_p$ for proving a lower bound.

Max-probability formulas. Our goal will be to argue that any $\phi \in \Phi_n$ has an acceptance probability dominated by one of two "maximum probability formulas" (*max-formulas*, for short). Namely, our first max-formula is the collision CNF (used in our upper bound in Section 2.2 and specialised here for one matching) and our second formula has a NAND constraint for each coordinate.

$$\textit{1st max-formula:}\quad \phi_{\text{Coll}}^{(n)} \;\coloneqq\; \bigwedge_{i\in[n/2]} (x_{2i-1}\leftrightarrow x_{2i})\wedge(y_{2i-1}\leftrightarrow y_{2i}) \quad \text{where } n \text{ is even,}$$

$$\textit{2nd max-formula:}\quad \phi_{\text{Nand}}^{(n)} \;\coloneqq\; \bigwedge_{i\in[n]} (\neg x_i \vee \neg y_i).$$

Writing $\Pr_{\mathcal{D}}[\phi] \coloneqq \Pr_{(X,Y)\sim\mathcal{D}}[\phi(X,Y)=1]$ for short, it is straightforward to see that

$$\Pr_{\mathcal{D}_p}[\phi_{\text{Coll}}^{(n)}] \;=\; (p^2+(1-p)^2)^n \qquad\text{and}\qquad \Pr_{\mathcal{D}_p}[\phi_{\text{Nand}}^{(n)}] \;=\; (1-p^2)^n. \tag{4}$$

Equating these probabilities and solving for p yields our optimal choice $p = p^* := 2/3$. The following lemma states that these formulas have, for $p = p^*$, higher acceptance probabilities than any 2-CNF consistent with IP_n (or $\neg\mathrm{IP}_n$).

▶ **Lemma 6.** $\Pr_{\mathcal{D}_{p^*}}[\phi] \le M^{(n)}_{p^*} := \max\left\{ \Pr_{\mathcal{D}_{p^*}}[\phi^{(n)}_{\mathrm{Coll}}],\, \Pr_{\mathcal{D}_{p^*}}[\phi^{(n)}_{\mathrm{Nand}}] \right\}$ *for any* $\phi \in \Phi_n$.

Using Lemma 6 it is easy to complete our proof. We get for any $\phi \in \Phi_n$,

$$\Pr_{\mathcal{D}_{p^*}}[\phi] \;\le\; M^{(n)}_{2/3} \;=\; (1-(2/3)^2)^n \;=\; 2^{-\log(9/5)\cdot n} \;=\; 2^{-0.847\ldots\cdot n}.$$

As per Dual LP, the reciprocal of this probability yields the claimed circuit lower bound.

It remains to prove Lemma 6, which we do in the rest of this section.

3.1 Proof of Lemma 6

To help us analyse acceptance probabilities, we first prove a *structure lemma* for any consistent 2-CNF formula ϕ. This lemma will find some "structured" formula ϕ' that is (semantically) *implied* by ϕ, denoted $\phi \models \phi'$ (that is, $\phi^{-1}(1) \subseteq \phi'^{-1}(1)$). The formula ϕ' comes from a set of structured formulas $\mathcal{S}_n$, which we will carefully define in Section 3.2. For now, it suffices for us to know that each structured formula $\phi^{(k)} \in \mathcal{S}_n$ only mentions variables among $(x_i, y_i)_{i\in I}$ for some subset $I \subseteq [n]$ of size $|I| = k$ (possibly $k \ll n$).

▶ **Lemma 7** (Structure lemma)**.** *Let* $\phi \in \Phi_n$ *be a 2-CNF consistent with* IP_n *or* $\neg\mathrm{IP}_n$*. Then there is some structured 2-CNF formula* $\phi^{(k)} \in \mathcal{S}_n$ *such that* $\phi \models \phi^{(k)}$*.*

We can now formulate a "localised" version of Lemma 6 for structured formulas. It allows us to locally compare the acceptance probability of $\phi^{(k)}$ with our max-formulas $\phi^{(k)}_{\mathrm{Coll}}$ and $\phi^{(k)}_{\mathrm{Nand}}$, now defined naturally over k many coordinates. Our original definition of $\phi^{(n)}_{\mathrm{Coll}}$ was actually assuming n is even. For technical convenience, for odd n, we define $\phi^{(n)}_{\mathrm{Coll}} := \phi^{(n-1)}_{\mathrm{Coll}} \wedge (x_n \leftrightarrow y_n)$. The bounds in (4) continue to hold for this extended definition.

▶ **Lemma 8.** $\Pr_{\mathcal{D}_{p^*}}[\phi^{(k)}] \le M^{(k)}_{p^*} := \max\left\{ \Pr_{\mathcal{D}_{p^*}}[\phi^{(k)}_{\mathrm{Coll}}],\, \Pr_{\mathcal{D}_{p^*}}[\phi^{(k)}_{\mathrm{Nand}}] \right\}$ *for any* $\phi^{(k)} \in \mathcal{S}_n$.

Using Lemmas 7 and 8 (proved below) it is now easy to prove Lemma 6:

Proof of Lemma 6. We prove this by induction on n. The base case $n = 0$ is vacuously true under the convention that $\Pr[\phi_\perp] = M^{(0)}_{p^*} = 1$ for the empty formula $\phi_\perp$. For the inductive case $n \ge 1$, let $\phi \in \Phi_n$ be arbitrary. Apply the structure lemma (Lemma 7) to find some $\phi^{(k)} \in \mathcal{S}_n$ such that $\phi \models \phi^{(k)}$. Suppose for notational convenience $\phi^{(k)}$ involves the first $k \le n$ coordinates. Let $\mathcal{D}^{(k)}_{p^*}$ denote our binomial distribution over $\{0,1\}^{2k}$. Then

$$\Pr_{\mathcal{D}^{(n)}_{p^*}}[\phi] \;\le \sum_{\substack{a,b\in\{0,1\}^k \\ \phi^{(k)}(a,b)=1}} \Pr_{\mathcal{D}^{(k)}_{p^*}}[(a,b)] \cdot \Pr_{\mathcal{D}^{(n-k)}_{p^*}}[\phi|_{a,b}],$$

where $\phi|_{a,b}$ is obtained from ϕ by restricting the first k coordinates to values (a,b). We note that restricting values in a formula consistent with IP_n might yield a formula consistent with $\neg\mathrm{IP}_{n-k}$ (and vice versa). We now apply Lemma 6 inductively for $\phi|_{a,b}$ to conclude

$$\Pr_{\mathcal{D}^{(n)}_{p^*}}[\phi] \;\le\; M^{(n-k)}_{p^*} \cdot \sum_{a,b} \Pr_{\mathcal{D}^{(k)}_{p^*}}[(a,b)] \;=\; M^{(n-k)}_{p^*} \cdot \Pr_{\mathcal{D}^{(k)}_{p^*}}[\phi^{(k)}] \;\le\; M^{(n-k)}_{p^*} M^{(k)}_{p^*} \;=\; M^{(n)}_{p^*},$$

where the last inequality is Lemma 8 and the final equality follows from (4). ◀

The rest of this section is organised as follows. We first define our family of structured formulas $\mathcal{S}_n$ in Section 3.2. Then we will prove Lemmas 7 and 8 in Sections 3.3 and 3.4.

3.2 Structured formulas in $\mathcal{S}_n$

We now proceed to define our family of structured formulas $\mathcal{S}_n$. The family will be closed under *symmetries of* IP_n, as we now explain. The value of inner product IP_n remains unchanged if we permute its n coordinates (e.g., swap (x_i, y_i) with (x_j, y_j)) or transpose two variables inside a single coordinate (i.e., swap (x_i, y_i) with (y_i, x_i)). These permutations generate the group of symmetries of IP_n. We say that two CNFs ϕ and ϕ' are *isomorphic* if there is some symmetry π of IP_n that, when applied to ϕ to yield ϕ^π, makes the two formulas equivalent, $\phi^\pi \equiv \phi'$, that is, to accept the same set of inputs.

Structured family $\mathcal{S}_n$. To define $\mathcal{S}_n$, we list below its various members. Each formula is defined over some $k \le n$ pairs of literals $L_k := \{\tilde{x}_1, \tilde{y}_1, \dots, \tilde{x}_k, \tilde{y}_k\}$ where $\tilde{x}_i \in \{x_i, \neg x_i\}$ and $\tilde{y}_i \in \{y_i, \neg y_i\}$. Each item defines a type of 2-CNF with the understanding that each of its isomorphic copies is included in $\mathcal{S}_n$. See Figure 1 for illustrations. We start with two cases corresponding to our max-formulas.

1. ***Nand*** is $\phi^{(1)}_{\mathrm{Nand}} = (\neg x_1 \vee \neg y_1)$. This is case $n = 1$ of our second max-formula.
2. ***Matching*** is defined relative to a perfect matching $M \subseteq \binom{L_k}{2}$ by

   $$\phi^{(k)}_{\mathrm{Match}} \;=\; \textstyle\bigwedge_{\{\ell,\ell'\}\in M} (\ell \leftrightarrow \ell').$$

   Note that this is a generalisation of our first max-formula (where the literals are positive and the perfect matching is more structured).

The final type of formula will be an extension of the following "ladder" formula

$$\psi^{(k)} \;=\; \textstyle\bigwedge_{i=1}^{k-1} (\tilde{y}_i \leftrightarrow \tilde{x}_{i+1}) \qquad \text{where } k \ge 2.$$

We also define two types of "terminal" constraints (where $\ell, \ell' \in L_k$),

$$\begin{aligned} &\textit{Back-edge:} && \psi^{\mathrm{left}}_{\mathrm{B}} = (\tilde{x}_1 \leftrightarrow \ell), && \psi^{\mathrm{right}}_{\mathrm{B}} = (\tilde{y}_k \leftrightarrow \ell') && \text{where } \ell \ne \tilde{x}_1 \text{ and } \ell' \ne \tilde{y}_k,\\ &\textit{Positive:} && \psi^{\mathrm{left}}_{\mathrm{P}} = (y_1 \to x_1), && \psi^{\mathrm{right}}_{\mathrm{P}} = (x_k \to y_k). \end{aligned}$$

3. ***Ladder*** is given by choosing terminal types $(L, R) \in \{\mathrm{B}, \mathrm{P}\}^2$ and defining

   $$\phi^{(k)}_{LR} \;=\; \psi^{(k)} \wedge \psi^{\mathrm{left}}_{L} \wedge \psi^{\mathrm{right}}_{R}.$$

▶ Remark 9. It can be shown that this list is *irredundant* in that, for each type, there is a formula $\phi^{(k)} \in \mathcal{S}_n$ of that type and $\phi \in \Phi_n$ such that $\phi \models \phi^{(k)}$ but $\phi \not\models \phi'$ for every $\phi' \in \mathcal{S}_n$ of type different than $\phi^{(k)}$. This means that we need all three types for our structure lemma.

3.3 Proof of Structure lemma (Lemma 7)

In the proof of Lemma 7, we use the standard notion of an implication graph of a 2-CNF.

Implication graphs. Given a 2-CNF ϕ over k variables $\{x_1, y_1, \dots, x_k, y_k\}$, its *implication graph* $G_\phi = (V, E)$ is the directed graph given by

$$\begin{aligned} V &:= \{x_1, \neg x_1, y_1, \neg y_1, \dots, x_k, \neg x_k, y_k, \neg y_k\},\\ E &:= \{(u, v) \in V^2 : u \ne v \text{ and } \phi \models (u \to v)\}. \end{aligned}$$

We note that implication graphs are sometimes defined more syntactically: For each clause $(u \vee v)$ of ϕ, include the edges $(\neg u, v)$ and $(\neg v, u)$ in G_ϕ, and moreover, for each

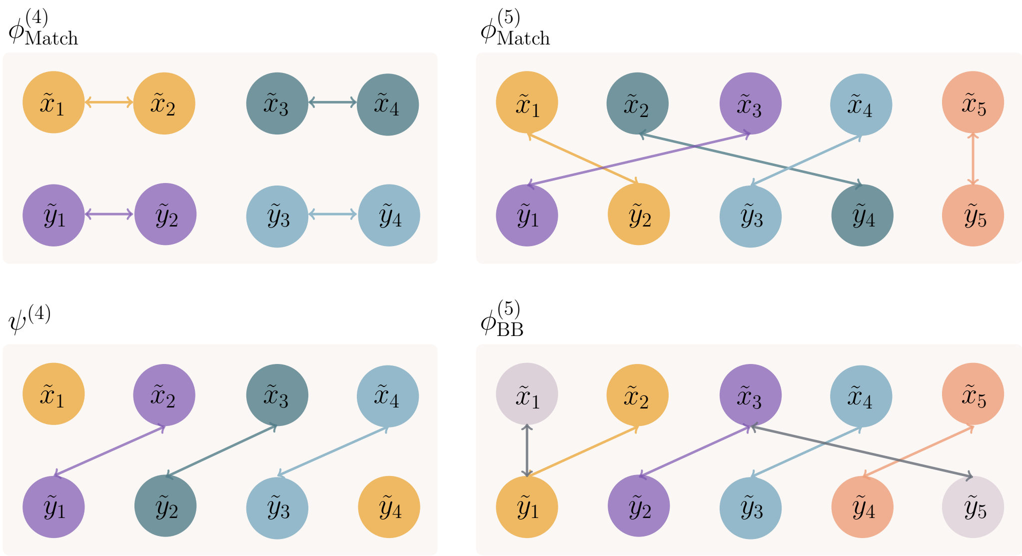

Figure 1 Examples of Matching and Ladder CNFs.

singleton clause (u) of ϕ, include the edges (v, u) in G_ϕ for all v. Taking the transitive closure (add edge (u, v) if there is a directed path from u to v) of this graph yields the graph in our (semantic) definition above.

We call a strongly connected component of G_ϕ a *strong-component* for short. We say that a variable x_i is fixed by ϕ if there is some $b \in \{0, 1\}$ such that for every $(x, y) \in \phi^{-1}(1)$ we have $x_i = b$. The following lemma will be used several times.

▶ **Lemma 10.** *Let $\phi \in \Phi_n$ and suppose y_1 lies in a strong-component of size 1 in G_ϕ. Then we have $\phi \models x_1 \to \tilde{y}_1$ for some $\tilde{y}_1 \in \{y_1, \neg y_1\}$.*

Proof. We may assume that y_1 is not fixed by ϕ, as otherwise the lemma is trivially true. We assume that $\phi \not\models x_1 \to \neg y_1$ and hope to show $\phi \models x_1 \to y_1$. Thus, there is some satisfying assignment $(x', y') \in \phi^{-1}(1)$ such that $(x'_1, y'_1) = (1, 1)$. Denote by $N_{\text{in}} \subseteq V$ the in-neighbours of y_1, that is, all the literals from which there exists an edge (equivalently, directed path, as G_ϕ is transitively closed) to y_1. Note that $\{\ell, \neg\ell\} \not\subseteq N_{\text{in}}$ for every literal ℓ, as otherwise one of ℓ or $\neg\ell$ would always be set to 1, forcing y_1 to always be 1, contradicting that y_1 is not fixed. Modify (x', y') by setting literals in N_{in} to 0. By the properties listed above, it follows that the new assignment, call it (x'', y''), still satisfies ϕ. Moreover, (x'', y'') has the property that we may flip the value of all the literals in the strong-component of y_1—which is just y_1 itself—and still remain a satisfying assignment. Since we can flip y_1 in isolation, we must have that $x''_1 = 0$ (otherwise we would change the parity of the 11 pattern, which contradicts $\phi \in \Phi_n$). But since $x'_1 = 1$ we must have that $x_1 \in N_{\text{in}}$, meaning that (x_1, y_1) is an edge, and hence $\phi \models x_1 \to y_1$, as desired. ◀

We now proceed to prove Lemma 7 in two cases by considering G_ϕ for $\phi \in \Phi_n$.

Case 1: *Every strong-component is of size 1.* Applying Lemma 10 twice, the second time with roles of x_1 and y_1 swapped, we learn that $\phi \models x_1 \to \tilde{y}_1$ and $\phi \models y_1 \to \tilde{x}_1$. If $\phi \models x_1 \to \neg y_1$ or $\phi \models y_1 \to \neg x_1$ holds then we have $\phi \models \phi^{(1)}_{\text{Nand}}$, as desired. In the remaining case, both $\phi \models x_1 \to y_1$ and $\phi \models y_1 \to x_1$ hold, which implies $\phi \models \phi^{(1)}_{\text{Match}}$.

Case 2: *There exists a strong-component of size at least 2.* Suppose by symmetry that y_1 lies in a strong-component of size at least 2. If y_1 is bidirectionally connected to $\tilde{x}_1$, that is, $\phi \models (y_1 \leftrightarrow \tilde{x}_1)$, then this means that $\phi \models \phi^{(1)}_{\text{Match}}$ and we are done.

Assume henceforth that y_1 is bidirectionally connected to some literal other than $\tilde{x}_1$, say by symmetry $y_1 \leftrightarrow \tilde{x}_2$. Consider y_2: is it bidirectionally connected to a literal in coordinate greater than 2? If yes, say by symmetry $y_2 \leftrightarrow \tilde{x}_3$. Consider y_3, etc. By this "unravelling" process, we are exposing the bidirectional edges of a ladder formula $\psi^{(k)}$. This process must eventually end at step $k \leq n$ in one of the following two cases.

- **Subcase 2-1:** *y_k is bidirectionally connected to some literal ℓ' in coordinate $\leq k$.* Here we have $\phi \models (y_k \leftrightarrow \ell') = \psi^{\text{right}}_{\text{B}}$.
- **Subcase 2-2:** *y_k lies in a singleton strong-component.* In this case, we apply Lemma 10 to learn that $\phi \models x_k \to \tilde{y}_k$. If $\models x_k \to \neg y_k$, then we would have found a copy of $\phi^{(1)}_{\text{Nand}}$ in coordinate k and we are done. Otherwise $\phi \models x_k \to y_k$, which means $\phi \models \psi^{\text{right}}_{\text{P}}$.

That is, in both cases (if we did not outright prove the lemma) we found either $\phi \models \psi^{\text{right}}_{\text{B}}$ or $\phi \models \psi^{\text{right}}_{\text{P}}$. By a similar argument, we can start unravelling edges starting at x_1 to find either $\phi \models \psi^{\text{left}}_{\text{B}}$ or $\phi \models \psi^{\text{left}}_{\text{P}}$. This will allow us to terminate the left side of the ladder, which completes the proof that $\phi \models \phi^{(k)}_{LR}$.

3.4 Proof of Lemma 8

We show the inequalities for every $\phi \in \mathcal{S}_n$.

- $\phi^{(1)}_{\text{Nand}}$: This is true by definition of $M^{(1)}_p$.
- $\phi^{(k)}_{\text{Match}}$: First note that the structure of the perfect matching for $\phi^{(k)}_{\text{Match}}$ will not change the acceptance probability because all input bits are *iid*. Moreover, when both ℓ and ℓ' are positive, $\Pr[\ell \leftrightarrow \ell'] = p^2 + (1-p)^2$; otherwise, $\Pr[\ell \leftrightarrow \ell'] = \max\{2p(1-p), p^2 + (1-p)^2\} \leq p^2 + (1-p)^2$ for all $p \in [0,1]$. Therefore, we have that $\Pr_{\mathcal{D}_p}[\phi^{(k)}_{\text{Match}}] \leq \Pr_{\mathcal{D}_p}[\phi^{(k)}_{\text{Coll}}]$.
- $\phi^{(k)}_{\text{BB}}$: We show in the above that $\Pr[\ell \leftrightarrow \ell'] \leq p^2 + (1-p)^2$ for any literals ℓ and ℓ'; we can similarly show that, for any literals ℓ, ℓ' and ℓ'', $\Pr[\ell \leftrightarrow \ell', \ell \leftrightarrow \ell''] \leq p^3 + (1-p)^3$. Replacing all literals in $\phi^{(k)}_{\text{BB}}$ by their positive analogues to get a new CNF ϕ, we have that $\Pr_{\mathcal{D}_p}[\phi^{(k)}_{\text{BB}}] \leq \Pr_{\mathcal{D}_p}[\phi]$. Let M be the perfect matching associated with ϕ. Define $M' := M \cup \{(x_1, y_k)\}$. Observe that M' is a perfect matching for all $2k$ literals. Let ϕ' be the matching CNF constructed from M'. Let P be the acceptance probability of ϕ. We know that $\Pr_{\mathcal{D}_p}[\phi'] = P \cdot \frac{[(1-p)^2+p^2]^3}{[(1-p)^3+p^3]^2} \geq P$ since $\frac{[(1-p)^2+p^2]^3}{[(1-p)^3+p^3]^2} \geq 1$ for $p \in [0,1]$.
- $\phi^{(k)}_{PP}$: Similarly, we can replace all literals in $\phi^{(k)}_{PP}$ with their positive analogues and get ϕ. Let M be the perfect matching associated with ϕ. Define $M' := M \cup \{(x_1, y_k)\}$. Observe that M' is a perfect matching for all $2k$ literals. Let ϕ' be the matching CNF constructed from M'. Let P be the acceptance probability of ϕ. If $k = 2$ then we have that $P = (1-p)^2 + p^4 = [(1-p)^2 + p^2]^2 = \Pr_{\mathcal{D}_p}[\phi']$ for $p = \frac{2}{3}$. If $k > 2$, we know that $\Pr_{\mathcal{D}_p}[\phi'] = P \cdot \frac{((1-p)^2+p^2)^3}{((1-p)^2+p^3)^2} > P$ since $\frac{((1-p)^2+p^2)^3}{((1-p)^2+p^3)^2} > 1$ for $p = \frac{2}{3}$.
- $\phi^{(k)}_{BP}$: As we have seen before, we can replace all literals in $\phi^{(k)}_{BP}$ with their positive analogues and get ϕ. Let M be the perfect matching associated with ϕ. Define $M' := M \cup \{(x_1, y_k)\}$. Observe that M' is a perfect matching for all $2k$ literals. Let ϕ' be the matching CNF constructed from M'. Let P be the acceptance probability of ϕ. If $k = 2$ then we have that $P = (1-p)^3 + p^4 < [(1-p)^2 + p^2]^2 = \Pr_{\mathcal{D}_p}[\phi']$ for $p = \frac{2}{3}$. If $k > 2$, we know that $\Pr_{\mathcal{D}_p}[\phi'] = P \cdot \frac{((1-p)^2+p^2)^3}{((1-p)^2+p^3)[(1-p)^3+p^3]} > P$ since $\frac{((1-p)^2+p^2)^3}{((1-p)^2+p^3)[(1-p)^3+p^3]} > 1$ for $p = \frac{2}{3}$.

Acknowledgements

We thank the anonymous reviewers for a careful reading of the paper and comments that helped us improve the presentation. This work was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number MB22.00026.

References

1 Peter Frankl, Svyatoslav Gryaznov, and Navid Talebanfard. A Variant of the VC-Dimension with Applications to Depth-3 Circuits. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215, pages 72:1–72:19, Dagstuhl, 2022. Schloss Dagstuhl. doi:10.4230/LIPIcs.ITCS.2022.72.

2 Alexander Golovnev, Alexander S. Kulikov, and R. Ryan Williams. Circuit Depth Reductions. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, volume 185, pages 24:1–24:20, Dagstuhl, 2021. Schloss Dagstuhl. doi:10.4230/LIPIcs.ITCS.2021.24.

3 J. Håstad, S. Jukna, and P. Pudlák. Top-down lower bounds for depth-three circuits. *Computational Complexity*, 5(2):99–112, jun 1995. doi:10.1007/bf01268140.

4 Suichi Hirahara. A duality between depth-three formulas and approximation by depth-two. Technical report, arXiv, 2017. arXiv:1705.03588.

5 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, dec 2001. doi:10.1006/jcss.2001.1774.

6 Rahul Jain and Hartmut Klauck. The partition bound for classical communication complexity and query complexity. In *Proceedings of the 25th Conference on Computational Complexity (CCC)*, pages 247–258. IEEE, 2010. doi:10.1109/CCC.2010.31.

7 Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.

8 Mauricio Karchmer, Eyal Kushilevitz, and Noam Nisan. Fractional covers and communication complexity. *SIAM Journal on Discrete Mathematics*, 8(1):76–92, feb 1995. doi:10.1137/s0895480192238482.

9 Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.

10 Troy Lee and Adi Shraibman. *Lower Bounds in Communication Complexity*, volume 3. Now Publishers, 2007. doi:10.1561/0400000040.

11 L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975. doi:10.1016/0012-365X(75)90058-8.

12 R. Paturi, M. E. Saks, and F. Zane. Exponential lower bounds for depth three boolean circuits. *computational complexity*, 9(1):1–15, 2000. doi:10.1007/PL00001598.

13 Ramamohan Paturi, Pavel Pudlák, Michael E. Saks, and Francis Zane. An improved exponential-time algorithm for k-SAT. *Journal of the ACM*, 52(3):337–364, may 2005. doi:10.1145/1066100.1066101.

14 Ramamohan Paturi, Pavel Pudlak, and Francis Zane. Satisfiability coding lemma. *Chicago Journal of Theoretical Computer Science*, 5(1):1–19, 1999. doi:10.4086/cjtcs.1999.011.

15 Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In Jozef Gruska, editor, *Mathematical Foundations of Computer Science 1977*, pages 162–176, Berlin, Heidelberg, 1977. Springer Berlin Heidelberg.

16 Emanuele Viola. On the power of small-depth computation. *Foundations and Trends in Theoretical Computer Science*, 5(1):1–72, 2009. doi:10.1561/0400000033.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il