



Shrinkage Under Random Projections, and Cubic Formula Lower Bounds for AC^0

Yuval Filmus 

Technion – Israel Institute of Technology, Haifa, Israel
<https://yuvalfilmus.cs.technion.ac.il>
yuvalfi@cs.technion.ac.il

Or Meir 

Department of Computer Science, University of Haifa, Israel
<https://cs.haifa.ac.il/~ormeir/>
ormeir@cs.haifa.ac.il

Avishay Tal 

Department of Electrical Engineering and Computer Sciences,
University of California at Berkeley, CA, USA
<http://www.avishaytal.org>
atal@berkeley.edu

Abstract

Håstad showed that any De Morgan formula (composed of AND, OR and NOT gates) shrinks by a factor of $O(p^2)$ under a random restriction that leaves each variable alive independently with probability p [SICOMP, 1998]. Using this result, he gave an $\tilde{\Omega}(n^3)$ formula size lower bound for the Andreev function, which, up to lower order improvements, remains the state-of-the-art lower bound for any explicit function.

In this work, we extend the shrinkage result of Håstad to hold under a far wider family of random restrictions and their generalization – random projections. Based on our shrinkage results, we obtain an $\tilde{\Omega}(n^3)$ formula size lower bound for an explicit function computed in AC^0 . This improves upon the best known formula size lower bounds for AC^0 , that were only quadratic prior to our work. In addition, we prove that the KRW conjecture [Karchmer et al., Computational Complexity 5(3/4), 1995] holds for inner functions for which the unweighted quantum adversary bound is tight. In particular, this holds for inner functions with a tight Khrapchenko bound.

Our random projections are tailor-made to the function's structure so that the function maintains structure even under projection – using such projections is necessary, as standard random restrictions simplify AC^0 circuits. In contrast, we show that any De Morgan formula shrinks by a quadratic factor under our random projections, allowing us to prove the cubic lower bound.

Our proof techniques build on the proof of Håstad for the simpler case of balanced formulas. This allows for a significantly simpler proof at the cost of slightly worse parameters. As such, when specialized to the case of p -random restrictions, our proof can be used as an exposition of Håstad's result.

2012 ACM Subject Classification Theory of computation → Circuit complexity

Keywords and phrases De Morgan formulas, KRW Conjecture, shrinkage, random restrictions, random projections, bounded depth circuits, constant depth circuits, formula complexity

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.89

Category Extended Abstract

Related Version A full version of the paper is available at <https://yuvalfilmus.cs.technion.ac.il/Papers/shrinkage.pdf>.

Funding *Yuval Filmus*: Taub Fellow – supported by the Taub Foundations. The research was funded by ISF grant 1337/16.

Or Meir: Partially supported by the Israel Science Foundation (grant No. 1445/16).



© Yuval Filmus, Or Meir, and Avishay Tal;
licensed under Creative Commons License CC-BY
12th Innovations in Theoretical Computer Science Conference (ITCS 2021).
Editor: James R. Lee; Article No. 89; pp. 89:1–89:7



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Acknowledgements A.T. would like to thank Igor Carboni Oliveira for bringing the question of proving formula size lower bounds for \mathbf{AC}^0 to his attention. We are also grateful to Robin Kothari for posing this open question on “Theoretical Computer Science Stack Exchange” [23], and to Kaveh Ghasemloo and Stasys Jukna for their feedback on this question. We would like to thank Anna Gál for very helpful discussions.

1 Introduction

1.1 Background

Is there an efficient computational task that cannot be perfectly parallelized? Equivalently, is $\mathbf{P} \not\subseteq \mathbf{NC}^1$? The answer is still unknown. The question can be rephrased as follows: is there a function in \mathbf{P} that does not have a (De Morgan) formula of polynomial size?

The history of formula lower bounds for functions in \mathbf{P} goes back to the 1960s, with the seminal result of Subbotovskaya [30] that introduced the technique of random restrictions. Subbotovskaya showed that the Parity function on n variables requires formulas of size at least $\Omega(n^{1.5})$. Khrapchenko [22], using a different proof technique, showed that in fact the Parity function on n variables requires formulas of size $\Theta(n^2)$. Later, Andreev [3] came up with a new explicit function (now known as the Andreev function) for which he was able to obtain an $\Omega(n^{2.5})$ size lower bound. This lower bound was subsequently improved by [18, 25, 14, 31] to $n^{3-o(1)}$.

The line of work initiated by Subbotovskaya and Andreev relies on the *shrinkage* of formulas under p -random restrictions. A p -random restriction is a randomly chosen partial assignment to the inputs of a function. Set a parameter $p \in (0, 1)$. We fix each variable independently with probability $1 - p$ to a uniformly random bit, and we keep the variable alive with probability p . Under such a restriction, formulas shrink (in expectation) by a factor more significant than p . Subbotovskaya showed that De Morgan formulas shrink to at most $p^{1.5}$ times their original size, whereas subsequent works of [25, 18] improved the bound to $p^{1.55}$ and $p^{1.63}$, respectively. Finally, Håstad [14] showed that the shrinkage exponent of De Morgan formulas is 2, or in other words, that De Morgan formulas shrink by a factor of $p^{2-o(1)}$ under p -random restrictions. Tal [31] improved the shrinkage factor to $O(p^2)$ – obtaining a tight result, as exhibited by the Parity function.

In a nutshell, shrinkage results are useful to proving lower bounds as long as the explicit function being analyzed maintains structure under such restrictions and does not trivialize. For example, the Parity function does not become constant as long as at least one variable remains alive. Thus any formula F that computes Parity must be of at least quadratic size, or else the formula F under restriction, keeping each variable alive with probability $100/n$, would likely become a constant function, whereas Parity would not. Andreev’s idea is similar, though he manages to construct a function such that under a random restriction keeping only $\Theta(\log n)$ of the variables, the formula size should be at least $\tilde{\Omega}(n)$ (in expectation). This ultimately gives the nearly cubic lower bound.

The KRW Conjecture

Despite much effort, proving $\mathbf{P} \not\subseteq \mathbf{NC}^1$, and even just breaking the cubic barrier in formula lower bounds, have remained a challenge for more than two decades. An approach to solve the \mathbf{P} versus \mathbf{NC}^1 problem was suggested by Karchmer, Raz and Wigderson [20]. They conjectured that when composing two Boolean functions, f and g , the formula size of the

resulting function, $f \diamond g$, is (roughly) the product of the formula sizes of f and g .¹ We will refer to this conjecture as the “KRW conjecture”. Under the KRW conjecture (and even under weaker variants of it), [20] constructed a function in \mathbf{P} with no polynomial-size formulas. It remains a major open challenge to settle the KRW conjecture.

A few special cases of the KRW conjecture are known to be true. The conjecture holds when either f or g is the AND or the OR function. Håstad’s result [14] and its improvement [31] show that the conjecture holds when the inner function g is the Parity function and the outer function f is any function. This gives an alternative explanation to the $n^{3-o(1)}$ lower bound for the Andreev function. Indeed, the Andreev function is at least as hard as the composition of a maximally-hard function f on $\log n$ bits and $g = \text{Parity}_{n/\log n}$, where the formula size of f is $\tilde{\Omega}(n)$ and the formula size of $\text{Parity}_{n/\log n}$ is $\Theta(n^2/\log^2 n)$. Since the KRW conjecture holds for this special case, the formula size of the Andreev function is at least $\tilde{\Omega}(n^3)$. In other words, the state-of-the-art formula size lower bounds for explicit functions follow from a special case of the KRW conjecture – the case in which g is the Parity function. Moreover, this special case follows from the shrinkage of De Morgan formulas under p -random restrictions.

Bottom-Up versus Top-Down Techniques

Whereas random restrictions are a “bottom-up” proof technique [15], a different line of work suggested a “top-down” approach using the language of communication complexity. The connection between formula size and communication complexity was introduced in the seminal work of Karchmer and Wigderson [21]. They defined for any Boolean function f a two-party communication problem KW_f : Alice gets an input x such that $f(x) = 1$, and Bob gets an input y such that $f(y) = 0$. Their goal is to identify a coordinate i on which $x_i \neq y_i$, while minimizing their communication. It turns out that there is a one-to-one correspondence between any protocol tree solving KW_f and any formula computing the function f . Since protocols naturally traverse the tree from root to leaf, proving lower bounds on their size or depth is done usually in a top-down fashion. This framework has proven to be very useful in proving formula lower bounds in the monotone setting (see, e.g., [21, 10, 28, 20, 27, 11, 26]). Moreover, a recent work by Dinur and Meir [6] was able to reprove Håstad’s cubic lower bound using the framework of Karchmer and Wigderson. As Dinur and Meir’s proof showed that top-down techniques can replicate Håstad’s cubic lower bound, a natural question (which motivated this project) arose:

Are top-down techniques superior to bottom-up techniques?

Towards that, we focused on a candidate problem: prove a cubic lower bound for an explicit function in \mathbf{AC}^0 .² Based on the work of Dinur and Meir [6], we suspected that such a lower bound could be achieved using top-down techniques. We were also *certain* that the problem cannot be solved using the random restriction technique. Indeed, in order to prove a lower bound on a function f using random restrictions, one should argue that f remains hard under a random restriction, however, it is well-known that functions in \mathbf{AC}^0 trivialize under p -random restrictions [7, 1, 32, 12]. Based on this intuition, surely random restrictions cannot show that a function in \mathbf{AC}^0 requires cubic size. Our intuition turned out to be false.

¹ More precisely, the original KRW conjecture [20] concerns depth complexity rather than formula complexity. The variant of the conjecture for formula complexity, which is discussed above, was posed in [9].

² Recall that \mathbf{AC}^0 is the class of functions computed by constant depth polynomial size circuits composed of AND and OR gates of unbounded fan-in, with variables or their negation at the leaves.

1.2 Our results

In this work, we construct an explicit function in \mathbf{AC}^0 which requires De Morgan formulas of size $n^{3-o(1)}$. Surprisingly, our proof is conducted via the bottom-up technique of random projections, which is a generalization of random restrictions (more details below).

► **Theorem 1.** *There exists a family of Boolean functions $h_n : \{0, 1\}^n \rightarrow \{0, 1\}$ for $n \in \mathbb{N}$ such that*

1. h_n can be computed by uniform depth-4 unbounded fan-in formulas of size $O(n^3)$.
2. The formula size of h_n is at least $n^{3-o(1)}$.

Prior to our work, the best formula size lower bounds on an explicit function in \mathbf{AC}^0 were only quadratic [24, 5, 19, 4].

Our hard function is a variant of the Andreev function. More specifically, recall that the Andreev function is based on the composition $f \diamond g$, where f is a maximally-hard function and g is the Parity function. Since Parity is not in \mathbf{AC}^0 , we cannot take g to be the Parity function in our construction. Instead, our hard function is obtained by replacing the Parity function with the Surjectivity function of [4].

As in the case of the Andreev function, we establish the hardness of our function by proving an appropriate special case of the KRW conjecture. To this end, we introduce a generalization of the complexity measure of Khrapchenko [22], called the *min-entropy Khrapchenko bound*. We prove the KRW conjecture for the special case in which the outer function f is any function, and g is a function whose formula complexity is bounded tightly by the min-entropy Khrapchenko bound. We then obtain Theorem 1 by applying this version of the KRW conjecture to the case where g is the Surjectivity function. We note that our KRW result also implies the known lower bounds in the cases where g is the Parity function [14] and the Majority function [8].

Our KRW result in fact applies more generally, to functions g whose formula complexity is bounded tightly by the “soft-adversary method”, denoted $\text{Adv}_s(g)$, which is a generalization of Ambainis’ unweighted adversary method [2].

Our proof of the special case of the KRW conjecture follows the methodology of Håstad [13], who proved the special case in which g is Parity on m variables. Håstad proved that De Morgan formulas shrink by a factor of (roughly) p^2 under p -random restrictions. Choosing $p = 1/m$ shrinks a formula for $f \diamond g$ by a factor of roughly m^2 , which coincides with the formula complexity of g . On the other hand, on average each copy of g simplifies to a single input variable, and so $f \diamond g$ simplifies to f . This shows that $L(f \diamond g) \gtrsim L(f) \cdot L(g)$.

Our main technical contribution is a new shrinkage theorem that works in a far wider range of scenarios than just p -random restrictions. Given a function g with soft-adversary bound $\text{Adv}_s(g)$, we construct a random projection³ which, on the one hand, shrinks De Morgan formulas by a factor of $\text{Adv}_s(g)$, and on the other hand, simplifies $f \diamond g$ to f . We thus show that $L(f \diamond g) \gtrsim L(f) \cdot \text{Adv}_s(g)$, and in particular, if $\text{Adv}_s(g) \approx L(g)$, then $L(f \diamond g) \gtrsim L(f) \cdot L(g)$, just as in Håstad’s proof. Using these random projections, that are tailored specifically to the structure of the function $f \diamond g$ so that $f \diamond g$ simplifies to f under projection, enables us to overcome the aforementioned difficulty. In contrast, p -random restrictions that do not respect the structure of $f \diamond g$ would likely result in a restricted function that is much simpler than f and in fact would be a constant function with high probability.

³ A projection is a mapping from the set of the variables $\{x_1, \dots, x_n\}$ to the set $\{y_1, \dots, y_m, \bar{y}_1, \dots, \bar{y}_m, 0, 1\}$, where y_1, \dots, y_m are formal variables.

Our shrinkage theorem applies more generally to two types of random projections, which we call *fixing projections* and *hiding projections*. Fixing projections are random projections in which fixing the value of a variable results in a projection which is much more probable. Hiding projections are random projections in which fixing the value of a variable hides which coordinates it appeared on. We note that our shrinkage theorem for fixing projections captures Håstad’s result for p -random restrictions as a special case.

The proof of our shrinkage theorem is based on Håstad’s proof [14], but also simplifies it. In particular, we take the simpler argument that Håstad uses for the special case of completely balanced trees, and adapt it to the general case. As such, our proof avoids a complicated case analysis, at the cost of slightly worse bounds. Using our bounds, it is nevertheless easy to obtain the $n^{3-o(1)}$ lower bound for the Andreev function. Therefore, one can see the specialization of our shrinkage result to p -random restrictions as an exposition of Håstad’s cubic lower bound.

An example: our techniques when specialized to $f \diamond \text{Majority}_m$

To illustrate our choice of random projections, we present its instantiation to the special case of $f \diamond g$ where $f : \{0, 1\}^k \rightarrow \{0, 1\}$ is non-constant and $g = \text{Majority}_m$ for some odd integer m . In this case, the input variables to $f \diamond g$ are composed of k disjoint blocks, B_1, \dots, B_k , each containing m variables. We use the random projection that for each block $B_i = \{x_{m(i-1)+1}, \dots, x_{mi}\}$, picks one variable in the block B_i uniformly at random, projects this variable to the new variable y_i , and fixes the rest of the variables in the block in a balanced way so that the number of zeros and ones in the block is equal (i.e., we have exactly $(m-1)/2$ zeros and $(m-1)/2$ ones). It is not hard to see that under this choice, $f \diamond g$ simplifies to f . On the other hand, we show that this choice of random projections shrinks the formula complexity by a factor of $\approx 1/m^2$. Combining the two together, we get that $L(f \diamond \text{Majority}_m) \gtrsim L(f) \cdot m^2$. Note that in this distribution of random projections, the different coordinates are not independent of one another, and this feature allows us to maintain structure.

1.3 Related work

Our technique of using tailor-made random projections was inspired by the celebrated result of Rossman, Servedio, and Tan [29, 16] that proved an average-case depth hierarchy. In fact, the idea to use tailor-made random restrictions goes back to Håstad’s thesis [17, Chapter 6.2]. Similar to our case, in [17, 29, 16], p -random restrictions are too crude to separate depth d from depth $d+1$ circuits. Given a circuit C of depth $d+1$, the main challenge is to construct a distribution of random restrictions or projections (tailored to the circuit C) that on the one hand maintains structure for C , but on the other hand simplify any depth d circuit C' .

Full Version

Due to space constraints, we have only included in this extended abstract the introduction of our paper. We defer the reader to the full version of the paper for more details and complete proofs.

References

- 1 Miklós Ajtai. Σ_1^1 -formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1–48, 1983.
- 2 Andris Ambainis. Quantum lower bounds by quantum arguments. *J. Comput. System Sci.*, 64(4):750–767, 2002. Special issue on STOC 2000 (Portland, OR). doi:10.1006/jcss.2002.1826.
- 3 Alexander E. Andreev. On a method for obtaining more than quadratic effective lower bounds for the complexity of π -schemes. *Moscow University Mathematics Bulletin*, 42(1):24–29, 1987.
- 4 Paul Beame and Widad Machmouchi. The quantum query complexity of AC0. *Quantum Info. Comput.*, 12(7–8):670–676, July 2012.
- 5 Andrew M. Childs, Shelby Kimmel, and Robin Kothari. The quantum query complexity of read-many formulas. In Leah Epstein and Paolo Ferragina, editors, *Algorithms – ESA 2012*, pages 337–348, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- 6 Irit Dinur and Or Meir. Toward the KRW composition conjecture: Cubic formula lower bounds via communication complexity. *Computational Complexity*, 27(3):375–462, 2018.
- 7 Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.
- 8 Anna Gál, Avishay Tal, and Adrian Trejo Nuñez. Cubic formula size lower bounds based on compositions with majority. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10–12, 2019, San Diego, California, USA*, volume 124 of *LIPICs*, pages 35:1–35:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.ITCS.2019.35.
- 9 Dmitry Gavinsky, Or Meir, Omri Weinstein, and Avi Wigderson. Toward better formula lower bounds: The composition of a function and a universal relation. *SIAM J. Comput.*, 46(1):114–131, 2017.
- 10 Mikael Goldmann and Johan Håstad. A simple lower bound for monotone clique using a communication game. *Inf. Process. Lett.*, 41(4):221–226, 1992.
- 11 Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. *SIAM J. Comput.*, 47(5):1778–1806, 2018.
- 12 Johan Håstad. Almost optimal lower bounds for small depth circuits. In *STOC*, pages 6–20, 1986.
- 13 Johan Håstad. The shrinkage exponent is 2. In *Proceedings of the 1993 IEEE 34th Annual Foundations of Computer Science, SFCS '93*, pages 114–123, USA, 1993. IEEE Computer Society. doi:10.1109/SFCS.1993.366876.
- 14 Johan Håstad. The shrinkage exponent of de Morgan formulas is 2. *SIAM J. Comput.*, 27(1):48–64, 1998.
- 15 Johan Håstad, Stasys Jukna, and Pavel Pudlák. Top-down lower bounds for depth-three circuits. *Computational Complexity*, 5(2):99–112, 1995.
- 16 Johan Håstad, Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. An average-case depth hierarchy theorem for boolean circuits. *J. ACM*, 64(5):35:1–35:27, 2017. doi:10.1145/3095799.
- 17 Johan Torkel Håstad. *Computational limitations for small-depth circuits*. MIT press, 1987.
- 18 Russell Impagliazzo and Noam Nisan. The effect of random restrictions on formula size. *Random Struct. Algorithms*, 4(2):121–134, 1993.
- 19 Stasys Jukna. *Boolean Function Complexity - Advances and Frontiers*, volume 27 of *Algorithms and combinatorics*. Springer, 2012.
- 20 Mauricio Karchmer, Ran Raz, and Avi Wigderson. Super-logarithmic depth lower bounds via the direct sum in communication complexity. *Computational Complexity*, 5(3/4):191–204, 1995.
- 21 Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. *SIAM J. Discrete Math.*, 3(2):255–265, 1990.
- 22 V. M. Khrapchenko. A method of obtaining lower bounds for the complexity of π -schemes. *Mathematical Notes Academy of Sciences USSR*, 10:474–479, 1972.

- 23 Robin Kothari. Formula size lower bounds for AC0 functions, 2011. Question on Theoretical Computer Science Stack Exchange. URL: <https://cstheory.stackexchange.com/questions/7156/formula-size-lower-bounds-for-ac0-functions>.
- 24 E. I. Neciporuk. On a Boolean function. *Soviet Mathematics Doklady*, 7(4):999–1000, 1966.
- 25 Mike Paterson and Uri Zwick. Shrinkage of de Morgan formulae under restriction. *Random Struct. Algorithms*, 4(2):135–150, 1993.
- 26 Toniann Pitassi and Robert Robere. Strongly exponential lower bounds for monotone computation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 1246–1255, 2017.
- 27 Ran Raz and Pierre McKenzie. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3):403–435, 1999.
- 28 Ran Raz and Avi Wigderson. Monotone circuits for matching require linear depth. *J. ACM*, 39(3):736–744, 1992.
- 29 Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. An average-case depth hierarchy theorem for Boolean circuits. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 1030–1048, 2015.
- 30 Bella Abramovna Subbotovskaya. Realizations of linear functions by formulas using +, ·, -. *Soviet Mathematics Doklady*, 2:110–112, 1961.
- 31 Avishay Tal. Shrinkage of de Morgan formulae by spectral techniques. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 551–560, 2014.
- 32 Andrew Chi-Chih Yao. Separating the polynomial-time hierarchy by oracles (preliminary version). In *FOCS*, pages 1–10. IEEE Computer Society, 1985.