Lifting with Simple Gadgets and Applications to Circuit and Proof Complexity

Susanna de Rezende*, Or Meir†, Jakob Nordström‡, Toniann Pitassi§, Robert Robere¶, and Marc Vinyals‖
* Institute of Mathematics of the Czech Academy of Sciences
† University of Haifa
‡ University of Copenhagen and Lund University
§ University of Toronto and Institute for Advanced Study
¶ McGill University
‖ Technion

2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS) | 978-1-7281-9621-3/20/$31.00 ©2020 IEEE | DOI: 10.1109/FOCS46700.2020.00011

***Abstract*—We significantly strengthen and generalize the theorem lifting Nullstellensatz degree to monotone span program size by Pitassi and Robere (2018) so that it works for any gadget with high enough rank, in particular, for useful gadgets such as *equality* and *greater-than*. We apply our generalized theorem to solve three open problems:**

- **We present the first result that demonstrates a separation in proof power for cutting planes with unbounded versus polynomially bounded coefficients. Specifically, we exhibit CNF formulas that can be refuted in quadratic length and constant line space in cutting planes with unbounded coefficients, but for which there are no refutations in subexponential length and subpolynomial line space if coefficients are restricted to be of polynomial magnitude.**
- **We give the first explicit separation between monotone Boolean formulas and monotone real formulas. Specifically, we give an explicit family of functions that can be computed with monotone real formulas of nearly linear size but require monotone Boolean formulas of exponential size. Previously only a non-explicit separation was known.**
- **We give the strongest separation to-date between monotone Boolean formulas and monotone Boolean circuits. Namely, we show that the classical GEN problem, which has polynomial-size monotone Boolean circuits, requires monotone Boolean formulas of size $2^{\Omega(n/\operatorname{polylog}(n))}$.**

An important technical ingredient, which may be of independent interest, is that we show that the Nullstellensatz degree of refuting the pebbling formula over a DAG G over any field coincides exactly with the reversible pebbling price of G. In particular, this implies that the standard decision tree complexity and the parity decision tree complexity of the corresponding falsified clause search problem are equal.

This is an extended abstract. The full version of the paper is available at https://arxiv.org/abs/2001.02144.

***Index Terms*—proof complexity; communication complexity; circuit complexity; cutting planes; trade-offs; pebble games**

I. Introduction

Lifting theorems in complexity theory are a method of transferring lower bounds in a weak computational model into lower bounds for a more powerful computational model via function composition. There has been an explosion of lifting theorems in the last ten years, essentially reducing communication lower bounds to query complexity lower bounds.

Early papers that establish lifting theorems include Raz and McKenzie's separation of the monotone NC hierarchy [46] (by lifting decision tree complexity to deterministic communication complexity), and Sherstov's pattern matrix method [52] which lifts (approximate) polynomial degree to (approximate) matrix rank. More recent works have established query-to-communication lifting theorems in a variety of models, leading to the resolution of many longstanding open problems in many areas of computer science. This includes problems in communication complexity [21]–[23], [25], [26], monotone complexity [43], [44], [50], proof complexity [12], [17], [24], [29], extension complexity of linear and semidefinite programs [20], [34], [38], data structures [8] and finite model theory [4].

Lifting theorems have the following form: given functions $f\colon \{0,1\}^n \to \{0,1\}$ (the "outer function") and $g\colon \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ (the "gadget"), a lower bound for f in a weak computational model implies a lower bound on $f \circ g^n$ in a stronger computational model. Lifting theorems should preferably be as general as possible. First, they should hold for *any* outer function, and ideally f should be allowed to be a partial function or a relation (i.e., a search problem). Indeed, nearly all of the applications mentioned above require lifting where the outer function is a relation or a partial function. Second, it is often desirable that the gadget be as small as possible. The most general lifting theorems established so far, for example those for deterministic and randomized communication complexity, require at least logarithmically-sized gadgets; if these theorems could be improved generically to hold for constant-sized gadgets, then many results would be vastly improved. Some notable examples where constant-sized gadgets are possible include Sherstov's degree-to-rank lifting [52], critical block-sensitivity lifting [24], [29], and lifting for monotone span programs [43], [44], [49].

II. A New Lifting Theorem

In this work, we generalize a lifting theorem of Pitassi and Robere [44] to use any gadget that has nontrivial rank. This theorem takes a search problem associated with an unsatisfiable CNF formula, and lifts a lower bound on the Nullstellensatz degree of refuting this formula to a lower bound on a related communication problem.

2575-8454/20/$31.00 ©2020 IEEE
DOI 10.1109/FOCS46700.2020.00011

More specifically, let $\mathcal{C}$ be an unsatisfiable k-CNF formula. The search problem associated with $\mathcal{C}$, which we denote $\mathsf{Search}(\mathcal{C})$, takes as input an assignment to the underlying variables, and outputs a clause that is falsified by the assignments. It was proven in [44] that for any unsatisfiable $\mathcal{C}$, and for a sufficiently rich gadget g, deterministic communication complexity lower bounds for the composed search problem $\mathsf{Search}(\mathcal{C}) \circ g^n$ follow from Nullstellensatz degree lower bounds for $\mathcal{C}$.[1] We significantly improve this lifting theorem so that it holds for *any* gadget of large enough rank.

Theorem 1. *Let $\mathcal{C}$ be a CNF formula over n variables, let $\mathbb{F}$ be any field, and let g be any gadget of rank at least r. Then the deterministic communication complexity of $\mathsf{Search}(\mathcal{C}) \circ g^n$ is at least $\mathsf{NS}_{\mathbb{F}}(\mathcal{C})$, the Nullstellensatz degree of refuting $\mathcal{C}$, as long as $r \geq cn/\mathsf{NS}_{\mathbb{F}}(\mathcal{C})$ for some large enough constant c.*

An important special case of our generalized theorem is when the gadget g is the equality function. In this work, we apply our theorem to resolve two open problems in proof complexity and circuit complexity. Both solutions depend crucially on the ability to use the equality gadget.

We remark that lifting with the equality gadget has recently been the focus of another paper by Loff and Mukhopadhyay [39], who observe that a lifting theorem with equality for *total functions* can be proven using a rank argument. Surprisingly, they also show that it is *not* possible to lift query complexity to communication complexity for arbitrary relations, giving an example of a relation with linear query complexity whose composition with equality has only polylogarithmic communication complexity. However, they prove a lifting theorem for general relations using the equality gadget by replacing standard query complexity with a stronger complexity measure (namely, the 0-query complexity of the relation).

Unfortunately, we cannot use either of the lifting theorems of [39] for our applications. We need to lift a search problem (and therefore cannot use the result for total functions), and this search problem has small 0-query complexity (meaning that we cannot use the lifting theorem for general relations). Indeed, this shows that our lifting theorem is incomparable to the results of [39], even when specialized to the equality gadget. One similarity, though, is that our theorem also bypasses the impossibility result of [39] by using a stronger complexity measure, which in our case is Nullstellensatz degree.

III. A Separation in Proof Complexity

The main application of our lifting theorem is the first separation in proof complexity between *cutting planes* proofs with high-weight versus low-weight coefficients. In the cutting planes proof system, an unsatisfiable CNF formula is refuted by first translating it into a system of 0-1 linear inequalities and then showing that this system has no integral solutions. The latter is achieved by a sequence of steps that derive new integer inequalities from old ones until the plainly contradictory inequality $0 \geq 1$ is reached. The efficiency of such a refutation can be measured by its *length* (i.e., the number of steps) and *space* (i.e., the maximal number of inequalities that have to be stored simultaneously during the derivation).

The standard version of the cutting planes proof system, commonly denoted by CP, allows the inequalities to use coefficients of arbitrary size. However, it is also interesting to consider the variant in which the coefficients are polynomially bounded, sometimes denoted by CP*. It is natural to ask how CP and CP* are related: are they polynomially equivalent or is there a super-polynomial length separation? This question appeared in [7] and remains stubbornly open to date. In this work we finally make progress by exhibiting for the first time a setting in which unbounded coefficients afford an exponential increase in proof power over polynomially bounded coefficients.

Theorem 2. *There is a family of CNF formulas of size N that have cutting planes refutations of length $\widetilde{O}(N^2)$ and space $O(1)$, but for which any refutation in length L and space s with polynomially bounded coefficients must satisfy $s \log L = \widetilde{\Omega}(N)$.*

Although this result is, to the best of our knowledge, the first of its kind in proof complexity, for Boolean functions the relative power of high-weight and low-weight linear threshold functions has been understood for a long time. The greater-than function can be computed by high-weight threshold functions but not by low-weight threshold functions, and weights of bit-length polynomial in n suffice [40] for Boolean functions. For higher-depth threshold formulas, it is known that depth-d threshold formulas of high-weight can efficiently be computed by depth-$(d+1)$ threshold formulas of low-weight [19].

It is all the more striking, then, that almost nothing is known about the relative power of high versus low weights in the context of proof complexity. Buss and Clote [7], building on work by Cook, Coullard, and Turán [9], proved an analog of the result of Muroga et al. for cutting planes, showing that it suffices to have weights of bit-length polynomial in the length of the proof (i.e., of exponential magnitude). Very recently Dadush and Tiwari [10] extended this to the more general linear threshold proof system *stabbing planes* [2], or equivalently the tree-like restriction of Krajíček's threshold logic proof system R(CP) [36], where one can additionally branch on linear threshold formulas. However, there is still no nontrivial upper bound on the weights of the unrestricted threshold logic proof system R(CP) nor of its extension LK(CP). Prior to our result, there was no separation between weights of exponential and polynomial magnitude for any linear threshold proof system.

IV. Separations in Circuit Complexity

Next, we describe two applications of our lifting theorem to lower bounds in circuit complexity. Our first application relates to *monotone real circuits*, which were introduced by Pudlák [45]. Monotone real circuits are a generalization of

[1] In fact the result is quite a bit stronger—it applies to Razborov's rank measure [48], which is a strict strengthening of deterministic communication complexity.

monotone Boolean circuits where each gate is allowed to compute any non-decreasing real function of its inputs, but the inputs and output of the circuit are Boolean. A *formula* is a tree-like circuit, that is, every gate has fan-out 1. The first (exponential) lower bound for monotone real circuits was proven already in [45] by extending the lower bound for computing the clique-colouring function with monotone Boolean circuits [1], [47]. This lower bound, together with a generalization of Krajíček's interpolation technique [35], was used by Pudlák to obtain the first exponential lower bounds for CP.

Shortly after monotone real circuits were introduced, there was an interest in understanding the power of monotone real computation in comparison to monotone Boolean computation. By extending techniques in [46], Bonet et al. proved that there are functions with polynomial size monotone Boolean circuits that require monotone real formulas of exponential size [5], [30]. This illustrates the power of DAG-like computation over tree-like computation. A related question is whether monotone real circuits are exponentially stronger than monotone Boolean circuits. Rosenbloom [51] answered this question in the affirmative by presenting an elegant and simple (but non-explicit) proof that monotone real formulas can be exponentially stronger than (even non-monotone) Boolean circuits, since slice functions can be computed by linear-size monotone real formulas, whereas by a counting argument we know that most slice functions require exponential size Boolean circuits.

The question of finding explicit functions demonstrating that monotone real circuits are stronger than general Boolean circuits is much more challenging since it involves proving explicit lower bounds for Boolean circuits—a task that currently seems completely out of reach. A more tractable problem is that of finding explicit functions showing that monotone real circuits or formulas are stronger than *monotone* Boolean circuits or formulas, but prior to this work, no such separation was known either. We provide an explicit separation for *monotone formulas*, that is, we provide a family of explicit functions that can be computed with monotone real formulas of near-linear size but require exponential size monotone Boolean formulas. This is the first explicit example that illustrates the strength of monotone real computation.

Theorem 3. *There is an explicit family of monotone functions f_n on n variables that can be computed by monotone real formulas of size $O(n \operatorname{polylog} n)$ but for which every monotone Boolean formula requires size $\exp\big(\Omega(n/\operatorname{polylog} n)\big)$.*

Our second application of our lifting theorem gives the strongest separation to-date between monotone Boolean formulas[2] and monotone Boolean circuits. Namely, we prove the following.[3]

[2]In fact, our lower bound holds for circuit models stronger than monotone Boolean formulas, such as monotone switching networks, monotone span programs, and monotone comparator circuits.

[3]We thank an anonymous reviewer for pointing out this corollary of our main lifting theorem.

Theorem 4. *There is an explicit family of functions f_n on n variables that can be computed by polynomial size monotone boolean circuits, but for which every monotone Boolean formula requires size $\exp\big(\Omega(n/\operatorname{polylog} n)\big)$.*

The first superpolynomial separation between monotone Boolean formulas and monotone Boolean circuits is due to Karchmer and Wigderson [32], who proved $n^{\Omega(\log n)}$ lower bounds on the size of any monotone Boolean formula computing the *st-connectivity* function. The first *exponential* separation is due to Raz and McKenzie [46], who proved $2^{\Omega(n^{\varepsilon})}$ monotone formula size lower bounds for a new function that they defined called GEN, which is a natural generalization of *st*-connectivity. Raz and McKenzie's GEN lower bound was strengthened by Göös and Pitassi [24], who proved $2^{\Theta(\sqrt{n})}$ lower bounds, which was the strongest result prior to our work. We also note that our $2^{\Omega(n/\operatorname{polylog} n)}$ lower bound also holds for a restriction of the GEN function.

Observe that our lower bound is close to the strongest possible separation of $O(\operatorname{poly}(n))$ size monotone circuits vs. $2^{\Omega(n)}$ size monotone formulas. While $2^{\Omega(n)}$ lower bounds are known for monotone Boolean formulas computing a monotone function in NP [43], it seems that we cannot obtain such a lower bound for GEN without using different techniques. For example, if one could improve our main lifting theorem to use *constant size gadgets* (i.e. if we could choose $r = O(1)$ in the statement of Theorem 1), then one can easily prove that any monotone Boolean formula for GEN requires size $2^{\Omega(n/\log n)}$; however, this appears to be the fundamental limit of our lower bound technique, due to an upper bound of $O(n/\log n)$ on the Nullstellensatz degree of the underlying search problem.

Finally, another motivation for studying lifting theorems with simple gadgets in circuit complexity (and, in particular, the equality gadget) is the connection with proving *non-monotone* formula size lower bounds. As noted earlier, lifting theorems have been extremely successful in proving monotone circuit lower bounds, and it has also been shown to be useful in some computational settings that are only "partially" monotone; notably monotone span programs [43], [44], [50] and extended formulations [20], [34].

This raises the question of to what extent lifting techniques can help prove *non-monotone* lower bounds. The beautiful work by Karchmer, Raz and Wigderson [31] initiated such an approach for separating P from NC^1—this opened up a line of research popularly known as the *KRW conjecture*. Intriguingly, steps towards resolving the KRW conjecture are closely connected to proving lifting theorems for the equality gadget. The first major progress was made in [14], where lower bounds for the universal relation game were proven, which is an important special case of the KRW conjecture. This result was recently improved in several papers [18], [27], [33], and Dinur and Meir [13] gave a new top-down proof of the state-of-the-art $\Omega(n^3)$ formula-size lower bounds via the KRW approach.

The connection to lifting using the equality gadget can be made by observing that the KRW conjecture involves

communication problems in which Alice and Bob are looking for a bit on which they differ—this is exactly an *equality* problem. A close examination of the results in [14], [27] shows that they are equivalent to proving lower bounds for the search problem associated with the pebbling formula when lifted with a 1-bit equality gadget on a particular graph [42]. In the full version of this work we establish near-optimal lower bounds on the communication complexity of the pebbling formula lifted with equality for *any* graph, but with size of the equality gadget larger than 1. If our main lifting theorem could be improved with one-bit equality gadgets, this would imply the results of [14], [27] as a direct corollary and with significantly better parameters.

V. Overview of Techniques

We now give a brief overview of our techniques, while also trying to convey some of the simplicity of the proofs which we believe is an extra virtue of these results.

Lifting theorem: In order to prove their lifting theorem, Pitassi and Robere [44] defined a notion of "good" gadgets. They then showed that if we compose a polynomial p with a good gadget g, the rank of the resulting matrix $p \circ g^n$ is determined *exactly* by the non-zero coefficients of p and the rank of g. Their lifting theorem follows by using this correspondence to obtain bounds on the ranks of certain matrices, which in turn yield the required communication complexity lower bound.

In this work, we observe that every gadget g can be turned into a good gadget using a simple transformation. This observation allows us to get an approximate bound on the rank of $p \circ g^n$ for any g with nontrivial rank. While the correspondence we get in this way is only an approximation and not an exact correspondence as in [44], it turns out that this approximation is sufficient to prove the required lower bounds. We thus get a lifting theorem that works for every gadget g with sufficiently large rank.

Cutting planes separation: The crux of our separation between CP and CP* is the following observation: CP can encode a conjunction of linear equalities with a single equality by using exponentially large coefficients. This allows CP refutations to obtain a significant saving in space when working with linear equalities, and in fact Filmus et al. [15] exploit it to separate the semantic and syntactic variants of cutting planes. Achieving such savings is not possible with the polynomially bounded coefficients in CP*, and this difference between the proof systems is what allows us to establish the separation.

In order to exploit this observation, one of our main innovations is to concoct the separating formulas. To do this, we must come up with candidate formulas that can only be refuted by reasoning about a large conjunction of linear equalities, show that cutting planes (CP) can efficiently refute them, and then proof that low-weight cutting planes (CP*) cannot do so.

To find such candidates we resort to the *pebbling formulas* Peb_G which have played a major role in many proof complexity trade-off results. Let G be any directed acyclic graph with a unique sink node t. Formally, the pebbling formula Peb_G is the following CNF formula. For each vertex $u \in V$ there is a variable z_u (intuitively, z_u should take the value "true" if and only if it is possible to place a "pebble" on u). The variables are constrained by the following clauses.

- a clause z_s for each source vertex s (i.e., we can always place a pebble on any source),
- a clause $\bigvee_{u \in \mathrm{pred}(v)} \neg z_u \vee z_v$ for each non-source vertex v with predecessors $\mathrm{pred}(v)$ (i.e., if we can place a pebble on the predecessors of v, then we can place a pebble on v), and
- a clause $\neg z_t$ for the sink t (i.e., it is impossible to place a pebble on t).

Interestingly, pebbling formulas have short refutations in the *resolution* proof system that reason in terms of large conjunctions of literals. We show that when pebbling formulas are "lifted" with an equality gadget—by replacing each variable z_u with an equality $\mathrm{EQ}(x, y)$ on "fresh" variables x, y—then the efficient resolution refutations of Peb_G can be simulated in cutting planes for $\mathrm{Peb}_G \circ \mathrm{EQ}^n$ using the large coefficients to encode the conjunction of many lifted literals with a single equality. In this way, we can construct cutting planes refutations of any pebbling formula in quadratic length and constant space.

For cutting planes with bounded coefficients, however, we establish a time-space lower bound showing that any CP* refutation requires large length or large space for the right type of pebbling formulas. To prove this lower bound, the first step is to instantiate the connection in [29] linking time-space lower bounds for many proof systems to communication complexity lower bounds for lifted search problems. This connection means that we can obtain the desired CP*-lower bounds for our lifted pebbling formulas $\mathrm{Peb}_G \circ \mathrm{EQ}^n$ by proving communication complexity lower bounds for the corresponding lifted search problem $\mathsf{Search}(\mathrm{Peb}_G) \circ \mathrm{EQ}^n$. In order to show the latter bounds, we first prove lower bounds on the Nullstellensatz degree of refuting pebbling formulas Peb_G, and then invoke our new lifting theorem to translate such bounds into communication complexity lower bounds for $\mathsf{Search}(\mathrm{Peb}_G) \circ \mathrm{EQ}^n$. Our Nullstellensatz degree lower bounds, in turn, follow from the next lemma, which establishes an equivalence between Nullstellensatz degree and *reversible pebbling* price, a result that we find to be interesting in its own right.

Lemma 5. *For any field $\mathbb{F}$ and any directed acyclic graph G, the Nullstellensatz degree of refuting the pebbling formula Peb_G is equal to the reversible pebbling price of G.*

Connections between Nullstellensatz degree and pebbling were previously shown in [6], but were not tight. We remark that building on our work, the connection between Nullstellensatz and reversible pebbling in Lemma 5 was further strengthened in [11]. We also want to point out that, thanks to previously known results in query and proof complexity, this lemma immediately implies that Nullstellensatz degree coincides with (deterministic) decision tree and parity decision

tree complexity. We record this corollary here as another result that can potentially be of independent interest.

Corollary 6. *For any field $\mathbb{F}$ and any directed acyclic graph G, the Nullstellensatz degree over $\mathbb{F}$ of refuting Peb_G, the decision tree depth of $\mathsf{Search}(\mathrm{Peb}_G)$, and the parity decision tree depth of $\mathsf{Search}(\mathrm{Peb}_G)$ all coincide and are equal to the reversible pebbling price of G.*

Returning from this brief detour to the cutting planes separation, by considering the family of graphs with maximal pebbling price in [41] and appealing to Lemma 5, we obtain the time-space lower bound for CP* stated in Theorem 2.

We wish to highlight the specific combination of lifting theorem and gadget that we need in order to achieve our separation of CP from CP*. On the one hand, the gadget should be strong enough, so that the lifting result applies for deterministic communication, which can simulate small-size small-space CP* refutations efficiently. On the other hand, the gadget also has to be weak enough, so that the lifted problem does not also become hard for stronger communication models such as randomized or real communication, which would immediately imply lower bounds for cutting planes with unbounded coefficients. The reason that we are focusing on the equality gadget is that it hits this sweet spot—it requires large deterministic communication complexity, yet admits efficient randomized and real protocols. Furthermore, when conjunctions of literals are lifted with the equality gadget, the lifted conjunction can be efficiently represented with a single linear equality with exponentially large coefficients.

Separations in monotone circuit complexity: First, as is the case for the separation between CP and CP*, in order to establish our separations between monotone Boolean formulas and monotone real formulas/monotone Boolean circuits we must find a function that has just the right level of hardness.

In both cases, to obtain a size lower bound for monotone Boolean formulas we invoke the characterization of monotone formula depth in terms of the communication complexity of the monotone Karchmer–Wigderson game [32]. By a standard reduction [16], [48], for *any* gadget g one can reduce the search problem $\mathsf{Search}(\mathrm{Peb}_G) \circ g^n$ to the monotone Karchmer-Wigderson game separating a family of minterms and maxterms of the GEN function—more precisely, we will *always* obtain minterms and maxterms of GEN, but the exact family of minterms and maxterms produced by the reduction will depend on the choice of G and g. Our main lifting theorem then implies a size lower bound for monotone Boolean formulas from the communication lower bound for this search problem, since monotone Boolean formulas can be balanced. To be precise, our lower bounds will hold for any function that separates the family of minterms and maxterms defined by $\mathsf{Search}(\mathrm{Peb}_G) \circ g^n$, for a suitable choice of G and for any gadget g for which our lifting theorem applies.

In the other direction, it is well-known [46] that the GEN function has polynomial-size monotone Boolean circuits, so the upper bound for Theorem 4 is immediate. However, the upper bound for Theorem 3 is more subtle, since it is known that GEN requires exponential-size monotone real formulas [5], [12]. To handle this issue, we make a special choice of gadget g—namely, the equality gadget—and show that the family of minterms and maxterms produced by the reduction can then be separated by a small monotone real formula. We construct this small monotone real formula in a novel way: namely, we first construct a small cutting planes refutation and then extract a small monotone real formula from it. This is essentially equivalent to *feasible interpolation*, a technique that has been used previously to prove proof complexity lower bounds, but this is perhaps the first time it is used to prove circuit upper bounds.

In more detail: analogous to the Karchmer–Wigderson relation, it was shown in [28] that there is a correspondence between real DAG-like communication protocols (as defined in [37]) and monotone real circuits. Using this relation, a small monotone real *circuit* can be extracted from a short CP refutation of a lifted pebbling formula. However, we would like to establish a monotone real *formula* upper bound. One way to achieve this is by finding small tree-like CP refutations of pebbling formulas lifted with the equality gadget. The problem is that for many gadgets lifted pebbling formulas require exponentially long tree-like proofs. Nevertheless, for pebbling formulas lifted with the equality gadget we are able to exhibit a short *semantic* tree-like CP refutation, which via real communication yields small monotone real formulas.

VI. Concluding Remarks

In this paper, we show that the cutting planes proof system (CP) is stronger than its variant with polynomially bounded coefficients (CP*) with respect to simultaneous length and space. This is the first result in proof complexity demonstrating any situation where high-weight coefficients are more powerful than low-weight coefficients. We also prove an explicit separation between monotone Boolean formulas and monotone real formulas. Previously the result was only known to hold non-constructively. To obtain these results we strengthen a lifting theorem of [44] to allow the lifting to work with *any* gadget with sufficiently large rank, in particular with the equality gadget—a crucial ingredient for obtaining the separations discussed above.

This work raises a number of questions. Prior to our result, no explicit function was known separating monotone real circuits or formulas from monotone Boolean circuits or formula. Although we prove an explicit formula separation, it remains open to obtain an explicit function that separates monotone real circuits from monotone Boolean circuits.

The most glaring open problem related to our cutting planes contribution is to strengthen our result to a true length separation, without any assumption on the space complexity. It is natural to ask whether techniques inspired by [17], [53] can be of use. Another thing to note about our trade-off result for CP* is that it is not a "true trade-off": we know that length and space cannot be optimised simultaneously, but we do not know if there in fact exist small space refutations. An interesting problem is, therefore, to exhibit formulas that present "true

trade-offs" for CP* but are easy with regard to space and length in CP.

It follows from our results that standard decision tree complexity, parity decision tree complexity, and Nullstellensatz degree are equal for the falsified clause search problem of lifted pebbling formulas. In view of this we can ask ourselves what complexity measure we are actually lifting. We know that for general search problem decision tree complexity is not enough for a lifting result. How about parity decision tree complexity? Or can we leverage the fact that we have "well-behaved" rectangle covers and small certificate complexity to lift weaker complexity models? It would be valuable to have a better understanding of the relation between gadgets, outer functions/relations and complexity measures.

Acknowledgements

Different subsets of the authors would like to acknowledge fruitful and enlightening conversations with different subsets of Arkadev Chattopadhyay, Pavel Hrubeš, Christian Ikenmeyer, Bruno Loff, Sagnik Mukhopadhyay, Igor Carboni Oliveira, Pavel Pudlák, and Dmitry Sokolov. We are also grateful for discussions regarding literature references with Albert Atserias, Paul Beame, and Massimo Lauria. We are thankful to the anonymous referees for their comments; in particular, indicating a simplified proof of a lemma and the fact that Theorem 4 is another corollary of our lifting theorem.

Part of this work was carried out while several of the authors were visiting the Simons Institute for the Theory of Computing in association with the DIMACS/Simons Collaboration on Lower Bounds in Computational Complexity, which is conducted with support from the National Science Foundation.

Or Meir was supported by the Israel Science Foundation (grant No. 1445/16). Toniann Pitassi was supported by NSERC. Susanna F. de Rezende and Jakob Nordström were supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement no. 279611, as well as by the Knut and Alice Wallenberg grant KAW 2016.0066. Susanna F. de Rezende also received funding fromg Knut and Alice Wallenberg Foundation grant KAW 2018.0371 and Jakob Nordström from the Swedish Research Council grants 621-2012-5645 and 2016-00782 and from the Independent Research Fund Denmark grant 9040-00389B. Part of this work was completed while Robert Robere was a postdoctoral researcher at DIMACS and the Institute for Advanced Study. This material is based upon work directly supported by the Charles Simonyi Endowment and indirectly supported by the National Science Foundation Grant No. CCF-1900460. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. Marc Vinyals was supported by the Prof. R Narasimhan post-doctoral award.

References

[1] N. Alon and R. B. Boppana, "The monotone circuit complexity of Boolean functions," *Combinatorica*, vol. 7, no. 1, pp. 1–22, Mar. 1987.

[2] P. Beame, N. Fleming, R. Impagliazzo, A. Kolokolova, D. Pankratov, T. Pitassi, and R. Robere, "Stabbing planes," in *Proceedings of the 9th Innovations in Theoretical Computer Science Conference (ITCS '18)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 94, Jan. 2018, pp. 10:1–10:20.

[3] E. Ben-Sasson and A. Wigderson, "Short proofs are narrow—resolution made simple," *Journal of the ACM*, vol. 48, no. 2, pp. 149–169, Mar. 2001, preliminary version in *STOC '99*.

[4] C. Berkholz and J. Nordström, "Near-optimal lower bounds on quantifier depth and Weisfeiler-Leman refinement steps," in *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '16)*, Jul. 2016, pp. 267–276.

[5] M. L. Bonet, J. L. Esteban, N. Galesi, and J. Johannsen, "On the relative complexity of resolution refinements and cutting planes proof systems," *SIAM Journal on Computing*, vol. 30, no. 5, pp. 1462–1484, 2000, preliminary version in *FOCS '98*.

[6] J. Buresh-Oppenheim, M. Clegg, R. Impagliazzo, and T. Pitassi, "Homogenization and the polynomial calculus," *Computational Complexity*, vol. 11, no. 3-4, pp. 91–108, 2002, preliminary version in *ICALP '00*.

[7] S. R. Buss and P. Clote, "Cutting planes, connectivity and threshold logic," *Archive for Mathematical Logic*, vol. 35, pp. 33–63, 1996.

[8] A. Chattopadhyay, M. Koucky, B. Loff, and S. Mukhopadhyay, "Simulation beats richness: New data-structure lower bounds," in *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC '18)*, Jun. 2018, pp. 1013–1020.

[9] W. Cook, C. R. Coullard, and G. Turán, "On the complexity of cutting-plane proofs," *Discrete Applied Mathematics*, vol. 18, no. 1, pp. 25–38, Nov. 1987.

[10] D. Dadush and S. Tiwari, "On the complexity of branching proofs," in *Proceedings of the 35th Annual Computational Complexity Conference (CCC '20)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 169, Jul. 2020, pp. 34:1–34:35.

[11] S. F. de Rezende, J. Nordström, O. Meir, and R. Robere, "Nullstellensatz size-degree trade-offs from reversible pebbling," in *Proceedings of the 34th Annual Computational Complexity Conference (CCC '19)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 137, Jul. 2019, pp. 18:1–18:16.

[12] S. F. de Rezende, J. Nordström, and M. Vinyals, "How limited interaction hinders real communication (and what it means for proof and circuit complexity)," in *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS '16)*, Oct. 2016, pp. 295–304.

[13] I. Dinur and O. Meir, "Toward the KRW composition conjecture: Cubic formula lower bounds via communication complexity," *Computational Complexity*, vol. 27, no. 3, pp. 375–462, 2018. [Online]. Available: https://doi.org/10.1007/s00037-017-0159-x

[14] J. Edmonds, R. Impagliazzo, S. Rudich, and J. Sgall, "Communication complexity towards lower bounds on circuit depth," *Computational Complexity*, vol. 10, no. 3, pp. 210–246, 2001. [Online]. Available: https://doi.org/10.1007/s00037-001-8195-x

[15] Y. Filmus, P. Hrubeš, and M. Lauria, "Semantic versus syntactic cutting planes," in *Proceedings of the 33rd International Symposium on Theoretical Aspects of Computer Science (STACS '16)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 47, Feb. 2016, pp. 35:1–35:13.

[16] A. Gál, "A characterization of span program size and improved lower bounds for monotone span programs," *Computational Complexity*, vol. 10, no. 4, pp. 277–296, Dec. 2001, preliminary version in *STOC '98*.

[17] A. Garg, M. Göös, P. Kamath, and D. Sokolov, "Monotone circuit lower bounds from resolution," in *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC '18)*, Jun. 2018, pp. 902–911.

[18] D. Gavinsky, O. Meir, O. Weinstein, and A. Wigderson, "Toward better formula lower bounds: The composition of a function and a universal relation," *SIAM Journal on Computing*, vol. 46, no. 1, pp. 114–131, Feb. 2017.

[19] M. Goldmann, J. Håstad, and A. A. Razborov, "Majority gates VS. general weighted threshold gates," *Computational Complexity*, vol. 2, pp. 277–300, 1992, preliminary version in *CCC '92*. [Online]. Available: https://doi.org/10.1007/BF01200426

[20] M. Göös, R. Jain, and T. Watson, "Extension complexity of independent set polytopes," *SIAM Journal on Computing*, vol. 47, no. 1, pp. 241–269, Feb. 2018.

[21] M. Göös, T. S. Jayram, T. Pitassi, and T. Watson, "Randomized communication vs. partition number," in *Proceedings of the 44th International Colloquium on Automata, Languages and Programming (ICALP '17)*,

[42] T. Pitassi, Manuscript, 2016.
[43] T. Pitassi and R. Robere, "Strongly exponential lower bounds for monotone computation," in *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC '17)*, Jun. 2017, pp. 1246–1255.
[44] ——, "Lifting Nullstellensatz to monotone span programs over any field," in *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC '18)*, Jun. 2018, pp. 1207–1219.
[45] P. Pudlák, "Lower bounds for resolution and cutting plane proofs and monotone computations," *Journal of Symbolic Logic*, vol. 62, no. 3, pp. 981–998, Sep. 1997.
[46] R. Raz and P. McKenzie, "Separation of the monotone NC hierarchy," *Combinatorica*, vol. 19, no. 3, pp. 403–435, Mar. 1999, preliminary version in *FOCS '97*.
[47] A. A. Razborov, "Lower bounds for the monotone complexity of some Boolean functions," *Soviet Mathematics Doklady*, vol. 31, no. 2, pp. 354–357, 1985, English translation of a paper in *Doklady Akademii Nauk SSSR*.
[48] ——, "Applications of matrix methods to the theory of lower bounds in computational complexity," *Combinatorica*, vol. 10, no. 1, pp. 81–93, Mar. 1990.
[49] R. Robere, "Unified lower bounds for monotone computation," Ph.D. dissertation, University of Toronto, 2018.
[50] R. Robere, T. Pitassi, B. Rossman, and S. A. Cook, "Exponential lower bounds for monotone span programs," in *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS '16)*, Oct. 2016, pp. 406–415.
[51] A. Rosenbloom, "Monotone real circuits are more powerful than monotone Boolean circuits," *Information Processing Letters*, vol. 61, no. 3, pp. 161–164, Feb. 1997.
[52] A. A. Sherstov, "The pattern matrix method," *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1969–2000, Dec. 2011, preliminary version in *STOC '08*.
[53] D. Sokolov, "Dag-like communication and its applications," in *Proceedings of the 12th International Computer Science Symposium in Russia (CSR '17)*, ser. Lecture Notes in Computer Science, vol. 10304. Springer, Jun. 2017, pp. 294–307.