Nullstellensatz Size-Degree Trade-offs from Reversible Pebbling

Susanna F. de Rezende
KTH Royal Institute of Technology, Stockholm, Sweden

Jakob Nordström
University of Copenhagen, Denmark
KTH Royal Institute of Technology, Stockholm, Sweden

Or Meir
University of Haifa, Israel
http://cs.haifa.ac.il/~ormeir/
ormeir@cs.haifa.ac.il

Robert Robere
DIMACS, New Brunswick, U.S.A.

Abstract

We establish an exactly tight relation between reversible pebblings of graphs and Nullstellensatz refutations of pebbling formulas, showing that a graph G can be reversibly pebbled in time t and space s if and only if there is a Nullstellensatz refutation of the pebbling formula over G in size $t+1$ and degree s (independently of the field in which the Nullstellensatz refutation is made). We use this correspondence to prove a number of strong size-degree trade-offs for Nullstellensatz, which to the best of our knowledge are the first such results for this proof system.

2012 ACM Subject Classification Theory of computation → Proof complexity

Keywords and phrases proof complexity, Nullstellensatz, pebble games, trade-offs, size, degree

Digital Object Identifier 10.4230/LIPIcs.CCC.2019.18

Funding This work was mostly carried out while 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.
Susanna F. de Rezende: was supported by the *Knut and Alice Wallenberg* grant KAW 2016.0066 *Approximation and Proof Complexity*.
Jakob Nordström: was supported by the *Knut and Alice Wallenberg* grant KAW 2016.0066 *Approximation and Proof Complexity* and by the Swedish Research Council grants 621-2012-5645 and 2016-00782.
Or Meir: was supported by the Israel Science Foundation (grant No. 1445/16).
Robert Robere: was supported by NSERC, and also conducted part of this work at DIMACS with support from the National Science Foundation under grant number CCF-1445755.

Acknowledgements We are grateful for many interesting discussions about matters pebbling-related (and not-so-pebbling-related) with Arkadev Chattopadhyay, Toniann Pitassi, and Marc Vinyals.

1 Introduction

In this work, we obtain strong trade-offs in proof complexity by making a connection to pebble games played on graphs. In this introductory section we start with a brief overview of these two areas and then explain how our results follow from connecting the two.

34th Computational Complexity Conference (CCC 2019).
Editor: Amir Shpilka; Article No. 18; pp. 18:1–18:16

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

1.1 Proof Complexity

Proof complexity is the study of efficiently verifiable certificates for mathematical statements. More concretely, statements of interest claim to provide correct answers to questions like:

- Given a CNF formula, does it have a satisfying assignment or not?
- Given a set of polynomials over some finite field, do they have a common root?

There is a clear asymmetry here in that it seems obvious what an easily verifiable certificate for positive answers to the above questions should be, while it is not so easy to see what a concise certificate for a negative answer could look like. The focus of proof complexity is therefore on the latter scenario.

In this paper we study the algebraic proof system system *Nullstellensatz* introduced by Beame et al. [7]. A *Nullstellensatz refutation* of a set of polynomials $\mathcal{P} = \{p_i \mid i \in [m]\}$ with coefficients in a field $\mathbb{F}$ is an expression

$$\sum_{i=1}^{m} r_i \cdot p_i + \sum_{j=1}^{n} s_j \cdot (x_j^2 - x_j) = 1 \tag{1}$$

(where r_i, s_j are also polynomials), showing that 1 lies in the polynomial ideal in the ring $\mathbb{F}[x_1, \ldots, x_n]$ generated by $\mathcal{P} \cup \left\{x_j^2 - x_j \,\middle|\, j \in [n]\right\}$. By (a slight extension of) Hilbert's Nullstellensatz, such a refutation exists if and only if there is no common $\{0,1\}$-valued root for the set of polynomials $\mathcal{P}$.

Nullstellensatz can also be viewed as a proof system for certifying the unsatisfiability of CNF formulas. If we translate a clause like, e.g., $C = x \vee y \vee \overline{z}$ to the polynomial $p(C) = (1-x)(1-y)z = z - yz - xz + xyz$, then an assignment to the variables in a CNF formula $F = \bigwedge_{i=1}^{m} C_i$ (where we think of 1 as true and 0 as false) is satisfying precisely if all the polynomials $\{p(C_i) \mid i \in [m]\}$ vanish.

The *size* of a Nullstellensatz refutation (1) is the total number of monomials in all the polynomials $r_i \cdot p_i$ and $s_j \cdot (x_j^2 - x_j)$ expanded out as linear combinations of monomials. Another, more well-studied, complexity measure for Nullstellensatz is *degree*, which is defined as $\max\{\deg(r_i \cdot p_i), \deg(s_j \cdot (x_j^2 - x_j))\}$.

In order to prove a lower bound d on the Nullstellensatz degree of refuting a set of polynomials $\mathcal{P}$, one can construct a *d-design*, which is a map D from degree-d polynomials in $\mathbb{F}[x_1, \ldots, x_n]$ to $\mathbb{F}$ such that

1. D is linear, i.e., $D(\alpha p + \beta q) = \alpha D(p) + \beta D(q)$ for $\alpha, \beta \in \mathbb{F}$;
2. $D(1) = 1$;
3. $D(rp) = 0$ for all $p \in \mathcal{P}$ and $r \in \mathbb{F}[x_1, \ldots, x_n]$ such that $\deg(rp) \leq d$;
4. $D(x^2 s) = D(xs)$ for all $s \in \mathbb{F}[x_1, \ldots, x_n]$ such that $\deg(s) \leq d - 2$.

Designs provide a characterization of Nullstellensatz degree in that there is a d-design for $\mathcal{P}$ if and only if there is no Nullstellensatz refutation of $\mathcal{P}$ in degree d [18]. Another possible approach to prove degree lower bounds is by computationally efficient versions of Craig's interpolation theorem. It was shown in [53] that constant-degree Nullstellensatz refutations yield polynomial-size monotone span programs, and that this is also tight: every span program is a unique interpolant for some set of polynomials refutable by Nullstellensatz. This connection has not been used to obtain Nullstellensatz degree lower bounds, however, due to the difficulty of proving span program lower bounds.

Lower bounds on Nullstellensatz degree have been proven for sets of polynomials encoding combinatorial principles such as the pigeonhole principle [6], induction principle [20], house-sitting principle [26, 18], matching [19], and pebbling [17]. It seems fair to say that research in algebraic proof complexity soon moved on to stronger systems such as *polynomial calculus* [26, 1], where the proof that 1 lies in the ideal generated by $\mathcal{P} \cup \left\{x_j^2 - x_j \,\middle|\, j \in [n]\right\}$ can be

constructed dynamically by a step-by-step derivation. However, the Nullstellensatz proof system has been the focus of renewed interest in a recent line of works [54, 50, 51, 29] showing that Nullstellensatz lower bounds can be lifted to stronger lower bounds for more powerful computational models using composition with gadgets. The size complexity measure for Nullstellensatz has also received attention in recent papers such as [14, 5].

In this work, we are interested in understanding the relation between size and degree in Nullstellensatz. In this context it is relevant to compare and contrast Nullstellensatz with polynomial calculus as well as with the well-known *resolution* proof system [15], which operates directly on the clauses of a CNF formula and repeatedly derives resolvent clauses $C \vee D$ from clauses of the form $C \vee x$ and $D \vee \overline{x}$ until contradiction, in the form of the empty clause without any literals, is reached. For resolution, size is measured by counting the number of clauses, and *width*, measured as the number of literals in a largest clause in a refutation, plays an analogous role to degree for Nullstellensatz and polynomial calculus.

By way of background, it is not hard to show that for all three proof systems upper bounds on degree/width imply upper bounds on size, in the sense that if a CNF formula over n variables can be refuted in degree/width d, then such a refutation can be carried out in size $n^{O(d)}$. Furthermore, this upper bound has been proven to be tight up to constant factors in the exponent for resolution and polynomial calculus [4], and it follows from [44] that this also holds for Nullstellensatz. In the other direction, it has been shown for resolution and polynomial calculus that strong enough lower bounds on degree/width imply lower bounds on size [36, 11]. This is known to be false for Nullstellensatz, and the pebbling formulas discussed in more detail later in this paper provide a counter-example [17].

The size lower bounds in terms of degree/width in [36, 11] can be established by transforming refutations in small size to refutations in small degree/width. This procedure blows up the size of the refutations exponentially, however. It is natural to ask whether such a blow-up is necessary or whether it is just an artifact of the proof. More generally, given that a formula has proofs in small size and small degree/width, it is an interesting question whether both measures can be optimized simultaneously, or whether there has to be a trade-off between the two.

For resolution this question was finally answered in [59], which established that there are indeed strong trade-offs between size and width making the size blow-up in [11] unavoidable. For polynomial calculus, the analogous question remains open.

In this paper, we show that there are strong trade-offs between size and degree for Nullstellensatz. We do so by establishing a tight relation between Nullstellensatz refutations of pebbling formulas and reversible pebblings of the graphs underlying such formulas. In order to discuss this connection in more detail, we first need to describe what reversible pebblings are. This brings us to our next topic.

1.2 Pebble Games

In the *pebble game* first studied by Paterson and Hewitt [48], one places pebbles on the vertices of a directed acyclic graph (DAG) G according to the following rules:

- If all (immediate) predecessors of an empty vertex v contain pebbles, a pebble may be placed on v.
- A pebble may be removed from any vertex at any time.

The game starts and ends with the graph being empty, and a pebble should be placed on the (unique) sink of G at some point. The complexity measures to minimize are the total number of pebbles on G at any given time (the *pebbling space*) and the number of moves (the *pebbling time*).

The pebble game has been used to study flowcharts and recursive schemata [48], register allocation [56], time and space as Turing-machine resources [27, 35], and algorithmic time-space trade-offs [25, 57, 55, 58, 60]. In the last two decades, pebble games have seen a revival in the context of proof complexity (see, e.g., [46]), and pebbling has also turned out to be useful for applications in cryptography [30, 2]. An excellent overview of pebbling up to ca. 1980 is given in [49] and some more recent developments are covered in the upcoming survey [47].

Bennett [13] introduced the *reversible pebble game* as part of a broader program [12] aimed at eliminating or reducing energy dissipation during computation. Reversible pebbling has also been of interest in the context of quantum computing. For example, it was noted in [45] that reversible pebble games can be used to capture the problem of "uncomputing" intermediate values in quantum algorithms.

The reversible pebble game adds the requirement that the whole pebbling performed in reverse order should also be a correct pebbling, which means that the rules for pebble placement and removal become symmetric as follows:

- If all predecessors of an empty vertex v contain pebbles, a pebble may be placed on v.
- If all predecessors of a pebbled vertex v contain pebbles, the pebble on v may be removed.

Reversible pebblings have been studied in [43, 39, 38] and have been used to prove time-space trade-offs in reversible simulation of irreversible computation in [42, 40, 61, 16]. In a different context, Potechin [52] implicitly used reversible pebbling to obtain lower bounds in monotone space complexity, with the connection made explicit in later works [24, 31]. The paper [23] (to which this overview is indebted) studied the relative power of standard and reversible pebblings with respect to space, and also established PSPACE-hardness results for estimating the minimum space required to pebble graphs (reversibly or not).

1.3 Our Contributions

In this paper, we obtain an exactly tight correspondence between on the one hand reversible pebblings of DAGs and on the other hand Nullstellensatz refutations of pebbling formulas over these DAGs. We show that for any DAG G it holds that G can be reversibly pebbled in time t and space s if and only if there is a Nullstellensatz refutation of the pebbling formula over G in size $t+1$ and degree s. This correspondence holds regardless of the field in which the Nullstellensatz refutation is operating, and so, in particular, it follows that pebbling formulas have exactly the same complexity for Nullstellensatz regardless of the ambient field.

We then revisit the time-space trade-off literature for the standard pebble game, focusing on the papers [21, 22, 41]. The results in these papers do not immediately transfer to the reversible pebble game, and we are not fully able to match the tightness of the results for standard pebbling, but we nevertheless obtain strong time-space trade-off results for the reversible pebble game.

This allows us to derive Nullstellensatz size-degree trade-offs from reversible pebbling time-space trade-offs as follows. Suppose that we have a DAG G such that:

1. G can be reversibly pebbled in time $t_1 \ll t_2$.
2. G can be reversibly pebbled in space $s_1 \ll s_2$.
3. There is no reversible pebbling of G that simultaneously achieves time t_1 and space s_1.

Then for Nullstellensatz refutations of the pebbling formula Peb_G over G (which will be formally defined shortly) we can deduce that:

1. Nullstellensatz can refute Peb_G in size $t_1 + 1 \ll t_2 + 1$.
2. Nullstellensatz can also refute Peb_G in degree $s_1 \ll s_2$.
3. There is no Nullstellensatz refutation of Peb_G that simultaneously achieves size $t_1 + 1$ and degree s_1.

We prove four such trade-off results, which can be found in Section 4. The following theorem is one example of such a result (specifically, it is a simplified version of Theorem 4).

▶ **Theorem 1.** *There is a family of 3-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that:*
1. *There is a Nullstellensatz refutation of F_n in degree $s_1 = \mathrm{O}\big(\sqrt[6]{n}\log n\big)$.*
2. *There is a Nullstellensatz refutation of F_n of near-linear size and degree $s_2 = \mathrm{O}\big(\sqrt[3]{n}\log n\big)$.*
3. *Any Nullstellensatz refutation of F_n in degree at most $\sqrt[3]{n}$ must have exponential size.*

It should be noted that this is not the first time proof complexity trade-off results have been obtained from pebble games. Pebbling formulas were used in [9, 10] to obtain size-space trade-offs for resolution, and later in [8] also for polynomial calculus. However, the current reductions between pebbling and Nullstellensatz are much stronger in that they go in both directions and are exact even up to additive constants.

With regard to Nullstellensatz, it was shown in [17] that Nullstellensatz degree is lower-bounded by standard pebbling price. This was strengthened in [29], which used the connection between designs and Nullstellensatz degree discussed above to establish that the degree needed to refute a pebbling formula exactly coincides with the reversible pebbling price of the corresponding DAG (which is always at least the standard pebbling price, but can be much larger). Our reduction significantly improves on [29] by constructing a more direct reduction, inspired by [34], that can simultaneously capture both time and space.

1.4 Outline of This Paper

After having discussed the necessary preliminaries in Section 2, we prove the reductions between Nullstellensatz and reversible pebblings in Section 3. In Section 4, we present the size-degree trade-offs for Nullstellensatz we obtain for different degree regimes. Section 5 contains some concluding remarks with suggestions for future directions of research.

2 Preliminaries

All logarithms in this paper are base 2 unless otherwise specified. For a positive integer n we write $[n]$ to denote the set of integers $\{1, 2, \ldots, n\}$.

A *literal* a over a Boolean variable x is either the variable x itself or its negation $\overline{x}$ (a *positive* or *negative* literal, respectively). A *clause* $C = a_1 \vee \cdots \vee a_k$ is a disjunction of literals. A *k-clause* is a clause that contains at most k literals. A formula F in *conjunctive normal form (CNF)* is a conjunction of clauses $F = C_1 \wedge \cdots \wedge C_m$. A *$k$-CNF formula* is a CNF formula consisting of k-clauses. We think of clauses and CNF formulas as sets, so that the order of elements is irrelevant and there are no repetitions. A truth value assignment ρ to the variables of a CNF formula F is satisfying if every clause in F contains a literal that is true under ρ.

2.1 Nullstellensatz

Let $\mathbb{F}$ be any field and let $\vec{x} = \{x_1, \ldots, x_n\}$ be a set of variables. We identify a set of polynomials $\mathcal{P} = \{p_i(\vec{x}) \mid i \in [m]\}$ in the ring $\mathbb{F}[\vec{x}]$ with the statement that all $p_i(\vec{x})$ have a common $\{0,1\}$-valued root. A *Nullstellensatz refutation* of this claim is a syntactic equality

$$\sum_{i=1}^{m} r_i(\vec{x}) \cdot p_i(\vec{x}) + \sum_{j=1}^{n} s_j(\vec{x}) \cdot (x_j^2 - x_j) = 1 \ , \tag{2}$$

where r_i, s_j are also polynomials in $\mathbb{F}[\vec{x}]$. We sometimes refer to the polynomials $p_i(\vec{x})$ as axioms and $(x_j^2 - x_j)$ as Boolean axioms.

As discussed in the introduction, Nullstellensatz can be used as a proof system for CNF formulas by translating a clause $C = \bigvee_{x \in P} x \vee \bigvee_{y \in N} \overline{y}$ to the polynomial $p(C) = \prod_{x \in P}(1 - x) \cdot \prod_{y \in N} y$ and viewing Nullstellensatz refutations of $\{p(C_i) \mid i \in [m]\}$ as refutations of the CNF formula $F = \bigwedge_{i=1}^{m} C_i$.

The *degree* of a Nullstellensatz refutation (1) is $\max\{\deg(r_i(\vec{x}) \cdot p_i(\vec{x})), \deg(s_j(\vec{x}) \cdot (x_j^2 - x_j))\}$. We define the *size* of a refutation (2) to be the total number of monomials encountered when all products of polynomials are expanded out as linear combinations of monomials. To be more precise, let $mSize(p)$ denote the number of monomials in a polynomial p written as a linear combination of monomials. Then the size of a Nullstellensatz refutation on the form (1) is

$$\sum_{i=1}^{m} mSize\big(r_i(\vec{x})\big) \cdot mSize\big(p_i(\vec{x})\big) + \sum_{j=1}^{n} 2 \cdot mSize\big(s_j(\vec{x})\big) \ . \tag{3}$$

This is consistent with how size is defined for the "dynamic version" of Nullstellensatz known as polynomial calculus [26, 1], and also with the general size definitions for so-called algebraic and semialgebraic proof systems in [4, 14, 5].

We remark that this is not the only possible way of measuring size, however. It can be noted that the definition (3) is quite wasteful in that it forces us to represent the proof in a very inefficient way. Other papers in the semialgebraic proof complexity literature, such as [33, 37, 28], instead define size in terms of the polynomials in isolation, more along the lines of

$$\sum_{i=1}^{m} \big(mSize\big(r_i(\vec{x})\big) + mSize\big(p_i(\vec{x})\big)\big) + \sum_{j=1}^{n} \big(mSize\big(s_j(\vec{x})\big) + 2\big) \ , \tag{4}$$

or as the bit size or "any reasonable size" of the representation of all polynomials $r_i(\vec{x}), p_i(\vec{x})$, and $s_j(\vec{x})$.

In the end, the difference is not too important since the two measures (3) and (4) are at most a square apart, and for size we typically want to distinguish between polynomial and superpolynomial. In addition, and more importantly, in this paper we will only deal with k-CNF formulas with $k = \mathrm{O}(1)$, and in this setting the two definitions are the same up to a constant factor 2^k. Therefore, we will stick with (3), which matches best how size is measured in the closely related proof systems resolution and polynomial calculus, and which gives the cleanest statements of our results.[1]

When proving lower bounds for algebraic proof systems it is often convenient to consider a *multilinear* setting where refutations are presented in the ring $\mathbb{F}[\vec{x}]/\{x_j^2 - x_j \mid j \in [n]\}$. Since the Boolean axioms $x_j^2 - x_j$ are no longer needed, the refutation (2) can be written simply as

$$\sum_{i=1}^{m} r_i(\vec{x}) \cdot p_i(\vec{x}) = 1 \ , \tag{5}$$

where we assume that all results of multiplications are implicitly multilinearized. It is clear that any refutation on the form (2) remains valid after multilinearization, and so the size and degree measures can only decrease in a multilinear setting. In this paper, we prove our lower bound in our reduction in the multilinear setting and the upper bound in the non-multilinear setting, making the tightly matching results even stronger.

[1] We refer the reader to Section 2.4 in [3] for a more detailed discussion of the definition of proof size in algebraic and semialgebraic proof systems.

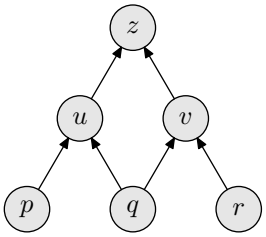

(a) Pyramid of height 2.

x_p	$x_p - 1$
$\wedge\ x_q$	$x_q - 1$
$\wedge\ x_r$	$x_r - 1$
$\wedge\ (\overline{x}_p \vee \overline{x}_q \vee x_u)$	$x_p x_q (1 - x_u)$
$\wedge\ (\overline{x}_q \vee \overline{x}_r \vee x_v)$	$x_q x_r (1 - x_v)$
$\wedge\ (\overline{x}_u \vee \overline{x}_v \vee x_z)$	$x_u x_v (1 - x_z)$
$\wedge\ \overline{x}_z$	x_z

(b) Pebbling formula in CNF. **(c)** Polynomial translation.

Figure 1 Example pebbling contradiction for the pyramid graph of height 2.

2.2 Reversible Pebbling and Pebbling Formulas

Throughout this paper $G = (V, E)$ denotes a directed acyclic graph (DAG) of constant fan-in with vertices $V(G) = V$ and edges $E(G) = E$. For an edge $(u, v) \in E$ we say that u is a *predecessor* of v and v a *successor* of u. We write $pred_G(v)$ to denote the sets of all predecessors of v, and drop the subscript when the DAG is clear from context. Vertices with no predecessors/successors are called *sources*/*sinks*. Unless stated otherwise we will assume that all DAGs under consideration have a unique sink z.

A *pebble configuration* on a DAG $G = (V, E)$ is a subset of vertices $\mathbb{P} \subseteq V$. A *reversible pebbling strategy* for a DAG G with sink z, or a *reversible pebbling* of G for short, is a sequence of pebble configurations $\mathcal{P} = (\mathbb{P}_0, \mathbb{P}_1, \ldots, \mathbb{P}_t)$ such that $\mathbb{P}_0 = \mathbb{P}_t = \emptyset$, $z \in \bigcup_{0 \leq t \leq t} \mathbb{P}_t$, and such that each configuration can be obtained from the previous one by one of the following rules:

1. $\mathbb{P}_{i+1} = \mathbb{P}_i \cup \{v\}$ for $v \notin \mathbb{P}_i$ such that $pred_G(v) \subseteq \mathbb{P}_i$ (a *pebble placement* on v).
2. $\mathbb{P}_{i+1} = \mathbb{P}_i \setminus \{v\}$ for $v \in \mathbb{P}_i$ such that $pred_G(v) \subseteq \mathbb{P}_i$ (a *pebble removal* from v).

The *time* of a pebbling $\mathcal{P} = (\mathbb{P}_0, \ldots, \mathbb{P}_t)$ is $\mathsf{time}(\mathcal{P}) = t$ and the *space* is $\mathsf{space}(\mathcal{P}) = \max_{0 \leq t \leq t}\{|\mathbb{P}_t|\}$.

We could also say that a reversible pebbling $\mathcal{P} = (\mathbb{P}_0, \ldots, \mathbb{P}_t)$ should be such that $\mathbb{P}_0 = \emptyset$ and $z \in \mathbb{P}_t$, and define the time of such a pebbling to be $2t$. This is so since once we have reached a configuration containing z we can simply run the pebbling backwards (because of reversibility) until we reach the empty configuration again, and without loss of generality all time- and space-optimal reversible pebblings can be turned into such pebblings. For simplicity, we will often take this viewpoint in what follows.

Pebble games can be encoded in CNF by so-called *pebbling formulas* [11], or *pebbling contradictions*. Given a DAG $G = (V, E)$ with a single sink z, we associate a variable x_v with every vertex v and add clauses encoding that

- the source vertices are all true;
- if all immediate predecessors are true, then truth propagates to the successor;
- but the sink is false.

In short, the pebbling formula over G consists of the clauses $x_v \vee \bigvee_{u \in pred(v)} \neg x_u$ for all $v \in V$ (note that if v is a source $pred(v) = \emptyset$), and the clause $\neg x_z$.

We encode this formula by a set of polynomials in the standard way. Given a set $U \subseteq V$, we denote by x_U the monomial $\prod_{u \in U} x_u$ (in particular, $x_\emptyset = 1$). For every vertex $v \in V$, we have the polynomial

$$A_v := (1 - x_v) \cdot x_{\mathrm{pred}(v)} \ , \tag{6}$$

and for the sink z we also have the polynomial

$$A_{\text{sink}} := x_z \ . \tag{7}$$

See Figure 1 for an illustration, including how the CNF formula is translated to a set of polynomials.

3 Reversible Pebblings and Nullstellensatz Refutations

In this section, we prove the correspondence between the reversible pebbling game on a graph G and Nullstellensatz refutation of the pebbling contradiction of G. Specifically, we prove the following result.

▶ **Theorem 2.** *Let G be a directed acyclic graph with a single sink, let ϕ be the corresponding pebbling contradiction, and let $\mathbb{F}$ be a field. Then, there is a reversible pebbling strategy for G with time at most t and space at most s if and only if there is a Nullstellensatz refutation for ϕ over $\mathbb{F}$ of size at most $t+1$ and degree at most s. Moreover, the same holds for multilinear Nullstellensatz refutations.*

We prove each of the directions of Theorem 2 separately in Sections 3.1 and 3.2 below.

3.1 From Pebbling to Refutation

We start by proving the "only if" direction of Theorem 2. Let

$$\mathbb{P} = (\mathbb{P}_0, \ldots, \mathbb{P}_t) \tag{8}$$

be an optimal reversible pebbling strategy for G. Let $\mathbb{P}_{t'}$ be the first configuration in which there is a pebble on the sink z. Without loss of generality, we may assume that $t = 2 \cdot t'$: if the last $t - t'$ steps were more efficient than the first t' steps, we could have obtained a more efficient strategy by replacing the first t' steps with the (reverse of) the last $t - t'$ steps, and vice versa.

We use $\mathbb{P}$ to construct a Nullstellensatz refutation over $\mathbb{F}$ for the pebbling contradiction ϕ. To this end, we will first construct for each step $i \in [t']$ of $\mathbb{P}$ a Nullstellensatz derivation of the polynomial $x_{\mathbb{P}_{i-1}} - x_{\mathbb{P}_i}$. The sum of all these polynomials is a telescoping sum, and is therefore equal to

$$x_{\mathbb{P}_0} - x_{\mathbb{P}_{t'}} = 1 - x_{\mathbb{P}_{t'}} \ . \tag{9}$$

We will then transform this sum into a Nullstellensatz refutation by adding the polynomial

$$x_{\mathbb{P}_{t'}} = A_{\text{sink}} \cdot x_{\mathbb{P}_{t'} - \{z\}} \ . \tag{10}$$

We turn to constructing the aforementioned derivations. To this end, for every $i \in [t']$, let $v_i \in V$ denote the vertex which was pebbled or unpebbled during the i-th step, i.e., during the transition from $\mathbb{P}_{i-1}$ to $\mathbb{P}_i$. Then, we know that in both configurations $\mathbb{P}_{i-1}$ and $\mathbb{P}_i$ the predecessors of v_i have pebbles on them, i.e., $\text{pred}(v) \subseteq \mathbb{P}_{i-1}, \mathbb{P}_i$. Let us denote by $R_i = \mathbb{P}_i - \{v_i\} - \text{pred}(v_i)$ the set of other vertices that have pebbles during the i-th step. Finally, let p_i be a number that equals to 1 if v_i was pebbled during the i-th step, and equals to -1 if v_i was unpebbled. Now, observe that

$$x_{\mathbb{P}_{i-1}} - x_{\mathbb{P}_i} = p_i \cdot x_{\mathbb{P}_{i-1}}(1 - x_{v_i}) = p_i \cdot x_{R_i} A_{v_i} \ , \tag{11}$$

where the last step follows since the predecessors of v_i are necessarily in $\mathbb{P}_{i-1}$. Therefore, our final refutation for ϕ is

$$\begin{aligned}\sum_{i=1}^{t'} A_{v_i} \cdot p_i \cdot x_{R_i} + A_{\text{sink}} \cdot x_{\mathbb{P}_{t'}-\{z\}} &= x_{\mathbb{P}_{t'}} + \sum_{i=1}^{t'} x_{\mathbb{P}_{i-1}} - x_{\mathbb{P}_i} \\ &= x_{\mathbb{P}_{t'}} + (x_{\mathbb{P}_0} - x_{\mathbb{P}_{t'}}) \\ &= x_{\mathbb{P}_{t'}} + (1 - x_{\mathbb{P}_{t'}}) = 1 \ .\end{aligned} \tag{12}$$

Note, in fact, it is a multilinear Nullstellensatz refutation, since it contains only multilinear monomials and does not use the Boolean axioms. It remains to analyze its degree and size.

For the degree, observe that every monomial in the proof is of the form $x_{\mathbb{P}_i}$, and the degree of each such monomial is exactly the number of pebbles used in the corresponding configuration. It follows that the maximal degree is exactly the space of the pebbling strategy $\mathbb{P}$.

Let us turn to considering the size. Observe that for each of the configurations $\mathbb{P}_1, \ldots, \mathbb{P}_{t'}$, the refutation contains exactly two monomials: for all $i \in [t'-1]$, one monomial for $\mathbb{P}_i$ is generated in the i-th step, and another in the $(i+1)$-th step, and for the configuration $\mathbb{P}_{t'}$ the second monomial is generated when we add $A_{\text{sink}} \cdot x_{\mathbb{P}_{t'}-\{z\}}$. In addition, the refutation contains exactly one monomial for the configuration $\mathbb{P}_0$, which is generated in the first step. Hence, the total number of monomials generated in the refutation is exactly $2 \cdot t' + 1 = t + 1$, as required.

3.2 From Refutation to Pebbling

We turn to prove the "if" direction of Theorem 2. We note that it suffices to prove it for multilinear Nullstellensatz refutations, since every standard Nullstellensatz refutation implies the existence of a multilinear one with at most the same size and degree. Let

$$\sum_{v \in V} A_v \cdot Q_v + A_{\text{sink}} \cdot Q_{\text{sink}} = 1 \tag{13}$$

be a multilinear Nullstellensatz refutation of ϕ over $\mathbb{F}$ of degree s. We use this refutation to construct a reversible pebbling strategy $\mathbb{P}$ for G.

To this end, we construct a "configuration graph" $\mathbb{C}$, whose vertices consist of all possible configurations of at most s pebbles on G (i.e., the vertices will be all subsets of V of size at most s). The edges of $\mathbb{C}$ will be determined by the polynomials Q_v of the refutation, and every edge $\{U_1, U_2\}$ in $\mathbb{C}$ will constitute a legal move in the reversible pebbling game (i.e., it will be legal to move from U_1 to U_2 and vice versa). We will show that $\mathbb{C}$ contains a path from the empty configuration $\emptyset$ to a configuration U_z that contains the sink z, and our pebbling strategy will be generated by walking on this path from $\emptyset$ to U_z and back.

The edges of the configuration graph $\mathbb{C}$ are defined as follows: Let $v \in V$ be a vertex of G, and let q be a monomial of Q_v that *does not contain* x_v. Let $W \subseteq V$ be the set of vertices such that $q = x_W$ (such a set W exists since the refutation is multilinear). Then, we put an edge e_q in $\mathbb{C}$ that connects $W \cup \text{pred}(v)$ and $W \cup \text{pred}(v) \cup \{v\}$ (we allow parallel edges). It is easy to see that the edge e_q connects configurations of size at most s, and that it indeed constitutes a legal move in the reversible pebbling game. We note that $\mathbb{C}$ is a bipartite graph: to see it, note that every edge e_q connects a configuration of an odd size to a configuration of an even size.

For the sake of the analysis, we assign the edge e_q a weight in $\mathbb{F}$ that is equal to coefficient of q in Q_v. We define *the weight of a configuration* U to be the sum of the weights of all the edges that touch U (where the addition is done in the field $\mathbb{F}$). We use the following technical claim, which we prove at the end of this section.

▷ Claim 3. Let $U \subseteq V$ be a configuration in $\mathbb{C}$ that does not contain the sink z. If U is empty, then its weight is 1. Otherwise, its weight is 0.

We now show how to construct the required pebbling strategy $\mathbb{P}$ for G. To this end, we first prove that there is a path in $\mathbb{C}$ from the empty configuration to a configuration that contains the sink z. Suppose for the sake of contradiction that this is not the case, and let $\mathbb{H}$ be the connected component of $\mathbb{C}$ that contains the empty configuration. Our assumption says that none of the configurations in $\mathbb{H}$ contains z.

The connected component $\mathbb{H}$ is bipartite since $\mathbb{C}$ is bipartite. Without loss of generality, assume that the empty configuration is in the left-hand side of $\mathbb{H}$. Clearly, the sum of the weights of the configurations on the left-hand side should be equal to the corresponding sum on the right-hand side, since they are both equal to the sum of the weights of the edges in $\mathbb{H}$. However, the sum of the weights of the configurations on the right-hand side is 0 (since all these weights are 0 by Claim 3), while the sum of the weights of the left-hand side is 1 (again, by Claim 3). We reached a contradiction, and therefore $\mathbb{H}$ must contain some configuration U_z that contains the sink z.

Next, let $\emptyset = \mathbb{P}_0, \mathbb{P}_1, \ldots, \mathbb{P}_{t'} = U_z$ be a path from the empty configuration to U_z. Our reversible pebbling strategy for G is

$$\mathbb{P} = (\mathbb{P}_0, \ldots, \mathbb{P}_{t'-1}, \mathbb{P}_{t'}, \mathbb{P}_{t'-1}, \ldots, \mathbb{P}_0) \quad . \tag{14}$$

This is a legal pebbling strategy since, as noted above, every edge of $\mathbb{C}$ constitutes a legal move of the reversible pebbling game. The strategy $\mathbb{P}$ uses space s, since all the configurations in $\mathbb{C}$ contain at most s pebbles by definition. The time of $\mathbb{P}$ is $t = 2 \cdot t'$. It therefore remains to show that the size of the Nullstellensatz refutation from Equation 13 is at least $t + 1$.

To this end, note that every edge e_q in the path corresponds to some monomial q in some polynomial Q_v. When the monomial q is multiplied by the axiom A_v, it generates two monomials in the proof: the monomial $q \cdot x_{\mathrm{pred}(v)}$ and the monomial $q \cdot x_{\mathrm{pred}(v)} \cdot x_v$. Hence, the Nullstellensatz refutation contains at least $2 \cdot t'$ monomials that correspond to edges from the path. In addition, the product $A_{\mathrm{sink}} \cdot Q_{\mathrm{sink}}$ must contains at least one monomial, since the refutation must use the sink axiom A_{sink} (because ϕ without this axiom is not a contradiction). It follows that the refutation contains at least $2 \cdot t' + 1 = t + 1$ monomials, as required. We conclude this section by proving Claim 3.

Proof of Claim 3. We start by introducing some terminology. First, observe that a monomial m may be generated multiple times in the refutation of Equation 13, and we refer to each time it is generated as an *occurrence* of m. We say that an occurrence of m is *generated by a monomial q_v of Q_v* if it is generated by the product $A_v \cdot q_v$. Throughout the proof, we identify a configuration U with the monomial x_U.

We first prove the claim for the non-empty case. Let $U \subseteq V$ be a non-empty configuration. We would like to prove the weight of U is 0. Recall that the weight of U is the sum of the coefficients of the occurrences of U that are generated by monomials q_v *that do not contain the corresponding vertex* v. Observe that Equation 13 implies that the sum of the coefficients of *all* the occurrences of U is 0: the coefficient of U on the right-hand side is 0, and it must be equal to the coefficient of U on the left-hand side, which is the sum of the coefficients of all the occurrences.

To complete the proof, we argue that every monomial q_v that does contain the vertex v contributes 0 to that sum. Let q_v be a monomial of Q_v that contains the vertex v and

generates an occurrence of U. Let α be the coefficient of q. Then, it must hold that

$$\begin{aligned} A_v \cdot q_v &= x_{\mathrm{pred}(v)} \cdot q_v - x_v \cdot x_{\mathrm{pred}(v)} \cdot q_v \\ &= x_{\mathrm{pred}(v)} \cdot q_v - x_{\mathrm{pred}(v)} \cdot q_v \\ &= \alpha \cdot x_U - \alpha \cdot x_U \ , \end{aligned} \tag{15}$$

where the second equality holds since we q_v contains v and we are working with a multilinear refutation, and the third equality holds since we assumed that q_v generates an occurrence of U. It follows that q_v generates two occurrences of U, one with coefficient α and one with coefficient $-\alpha$, and therefore it contributes 0 to the sum of the coefficients of all the occurrences of U.

We have shown that the sum of the coefficients of all the occurrences of U is 0, and that the occurrences generated by monomials q_v that contain v contribute 0 to this sum, and therefore the sum of coefficients of occurrences that are generated by monomials q_v that do not contain v must be 0, as required. In the case that U is the empty configuration, the proof is identical, except that the sum of the coefficients of all occurrences is 1, since the coefficient of $\emptyset$ is 1 on the right hand side of Equation 13. ◁

4 Nullstellensatz Trade-offs from Reversible Pebbling

In this section we present the Nullstellensatz size-degree trade-offs we obtain for different degree regimes. Let us first recall what is known with regards to degree and size. In what follows, a Nullstellensatz refutation of a CNF formula F refers to a Nullstellensatz refutation of the translation of F to polynomials. As mentioned in the introduction, if a CNF formula over n variables can be refuted in degree d then it can be refuted in simultaneous degree d and size $n^{\mathrm{O}(d)}$. However, for Nullstellensatz it is not the case that strong enough degree lower bounds imply size lower bounds.

A natural question is whether for any given function $d_1(n)$ there is a family of CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that

1. F_n has a Nullstellensatz refutation $d_1(n)$;
2. F_n has a Nullstellensatz refutation of (close to) linear size and degree $d_2(n) \gg d_1(n)$;
3. Any Nullstellensatz refutation of F_n in degree only slightly below $d_2(n)$ must have size nearly $n^{d_1(n)}$.

We present explicit constructions of formulas providing such trade-offs in several different parameter regimes. We first show that there are formulas that require exponential size in Nullstellensatz if the degree is bounded by some polynomial function, but if we allow slightly larger degree there is a nearly linear size proof.

▶ **Theorem 4.** *There is a family of explicitly constructible unsatisfiable 3-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that:*

1. *There is a Nullstellensatz refutation of F_n in degree $d_1 = \mathrm{O}\big(\sqrt[6]{n} \log n\big)$.*
2. *For any constant $\epsilon > 0$, there is a Nullstellensatz refutation of F_n of size $\mathrm{O}(n^{1+\epsilon})$ and degree $d_2 = \mathrm{O}\big(d_1 \cdot \sqrt[6]{n}\big) = \mathrm{O}\big(\sqrt[3]{n} \log n\big)$.*
3. *There exists a constant $K > 0$ such that any Nullstellensatz refutation of F_n in degree at most $d = Kd_2/\log n = \mathrm{O}\big(\sqrt[3]{n}\big)$ must have size $\big(\sqrt[6]{n}\big)!$.*

We also analyse a family of formulas that can be refuted in close to logarithmic degree and show that even if we allow up to a certain polynomial degree, the Nullstellensatz size required is superpolynomial.

▶ **Theorem 5.** *Let $\delta > 0$ be an arbitrarily small positive constant and let $g(n)$ be any arbitrarily slowly growing monotone function $\omega(1) = g(n) \leq n^{1/4}$. Then there is a family of explicitly constructible unsatisfiable 3-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that:*

1. *There is a Nullstellensatz refutation of F_n in degree $d_1 = g(n)\log(n)$.*
2. *For any constant $\epsilon > 0$, there is a Nullstellensatz refutation of F_n of size $\mathrm{O}(n^{1+\epsilon})$ and degree*

$$d_2 = \mathrm{O}\big(d_1 \cdot n^{1/2}/g(n)^2\big) = \mathrm{O}\big(n^{1/2}\log n/g(n)\big).$$

3. *Any Nullstellensatz refutation of F_n in degree at most*

$$d = \mathrm{O}\big(d_2/n^{\delta}\log n\big) = \mathrm{O}\big(n^{1/2-\delta}/g(n)\big)$$

must have size superpolynomial in n.

Still in the small-degree regime, we present a very robust trade-off in the sense that superpolynomial size lower bound holds for degree from $\log^2(n)$ to $n/\log(n)$.

▶ **Theorem 6.** *There is a family of explicitly constructible unsatisfiable 3-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that:*

1. *There is a Nullstellensatz refutation of F_n in degree $d_1 = \mathrm{O}(\log^2 n)$.*
2. *For any constant $\delta > 0$, there is a Nullstellensatz refutation of F_n of size $\mathrm{O}(n)$ and degree*

$$d_2 = \mathrm{O}(d_1 \cdot n/\log^{3-\delta} n) = \mathrm{O}(n/\log^{1-\delta} n).$$

3. *There exists a constant $K > 0$ such that any Nullstellensatz refutation of F_n in degree at most $d = Kd_2/\log^{\delta} n = O(n/\log n)$ must have size $n^{\Omega(\log\log n)}$.*

Finally, we study a family of formulas that have Nullstellensatz refutation of quadratic size and that present a smooth size-degree trade-off.

▶ **Theorem 7.** *There is a family of explicitly constructible unsatisfiable 3-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that any Nullstellensatz refutation of F_n that optimizes size given degree constraint $d = n^{\Theta(1)}$ (and less than n) has size $\Theta\big(n^2/d\big)$.*

We prove these results by obtaining the analogous time-space trade-offs for reversible pebbling and then applying the tight correspondence between size and degree in Nullstellensatz and time and space in reversible pebbling. We differ the reader to the upcoming full version for the details.

5 Concluding Remarks

In this paper we prove that size and degree of Nullstellensatz refutations in any field of pebbling formulas are exactly captured by time and space of the reversible pebble game on the underlying graph. This allows us to prove a number of strong size-degree trade-offs for Nullstellensatz. To the best of our understanding no such results have been known previously.

The most obvious, and also most interesting, open question is whether there are also size-degree trade-offs for the stronger polynomial calculus proof system. Such trade-offs cannot be exhibited by the pebbling formulas considered in this work, since such formulas have small-size low-degree polynomial calculus refutations, but the formulas exhibiting size-width trade-offs for resolution [59] appear to be natural candidates.

Another interesting question is whether the tight relation between Nullstellensatz and reversible pebbling could make it possible to prove even sharper trade-offs for size versus degree in Nullstellensatz, where just a small constant drop in the degree would lead to an

exponential blow-up in size. Such results for pebbling time versus space are known for the standard pebble game, e.g., in [32]. It is conceivable that a similar idea could be applied to the reversible pebbling reductions in [23], but it is not obvious whether just adding a small amount of space makes it possible to carry out the reversible pebbling time-efficiently enough.

Finally, it can be noted that our results crucially depend on that we are in a setting with variables only for positive literals. For polynomial calculus it is quite common to consider the stronger setting with "twin variables" for negated literals (as in the generalization of polynomial calculus in [26] to *polynomial calculus resolution* in [1]). It would be nice to generalize our size-degree trade-offs for Nullstellensatz to this setting, but it is not obvious whether the reductions in the current work could be made to work or not.

References

1 Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space Complexity in Propositional Calculus. *SIAM Journal on Computing*, 31(4):1184–1211, 2002. Preliminary version in *STOC '00*.

2 Joël Alwen and Vladimir Serbinenko. High Parallel Complexity Graphs and Memory-Hard Functions. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC '15)*, pages 595–603, June 2015.

3 Albert Atserias and Tuomas Hakoniemi. Size-Degree Trade-Offs for Sums-of-Squares and Positivstellensatz Proofs. Technical report, arXiv.org, November 2018. `arXiv:1811.01351`.

4 Albert Atserias, Massimo Lauria, and Jakob Nordström. Narrow Proofs May Be Maximally Long. *ACM Transactions on Computational Logic*, 17(3):19:1–19:30, May 2016. Preliminary version in *CCC '14*.

5 Albert Atserias and Joanna Ochremiak. Proof Complexity Meets Algebra. *ACM Transactions on Computational Logic*, 20:1:1–1:46, February 2019. Preliminary version in *ICALP '17*.

6 Paul Beame, Stephen A. Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The Relative Complexity of NP Search Problems. *Journal of Computer and System Sciences*, 57(1):3–19, August 1998. Preliminary version in *STOC '95*.

7 Paul Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, and Pavel Pudlák. Lower Bounds on Hilbert's Nullstellensatz and Propositional Proofs. In *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science (FOCS '94)*, pages 794–806, November 1994.

8 Chris Beck, Jakob Nordström, and Bangsheng Tang. Some Trade-off Results for Polynomial Calculus. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC '13)*, pages 813–822, May 2013.

9 Eli Ben-Sasson. Size-Space Tradeoffs for Resolution. *SIAM Journal on Computing*, 38(6):2511–2525, May 2009. Preliminary version in *STOC '02*.

10 Eli Ben-Sasson and Jakob Nordström. Understanding Space in Proof Complexity: Separations and Trade-offs via Substitutions. In *Proceedings of the 2nd Symposium on Innovations in Computer Science (ICS '11)*, pages 401–416, January 2011.

11 Eli Ben-Sasson and Avi Wigderson. Short Proofs are Narrow—Resolution Made Simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version in *STOC '99*.

12 Charles H. Bennett. Logical Reversibility of Computation. *IBM Journal of Research and Development*, 17(6):525–532, November 1973.

13 Charles H. Bennett. Time/Space Trade-offs for Reversible Computation. *SIAM Journal on Computing*, 18(4):766–776, August 1989.

14 Christoph Berkholz. The Relation between Polynomial Calculus, Sherali-Adams, and Sum-of-Squares Proofs. In *Proceedings of the 35th Symposium on Theoretical Aspects of Computer Science (STACS '18)*, volume 96 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 11:1–11:14, February 2018.

15 Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.

16 Harry Buhrman, John Tromp, and Paul Vitányi. Time and Space Bounds for Reversible Simulation. *Journal of physics A: Mathematical and general*, 34:6821–6830, 2001. Preliminary version in *ICALP '01*.

17 Joshua Buresh-Oppenheim, Matthew Clegg, Russell Impagliazzo, and Toniann Pitassi. Homogenization and the Polynomial Calculus. *Computational Complexity*, 11(3-4):91–108, 2002. Preliminary version in *ICALP '00*.

18 Samuel R. Buss. Lower Bounds on Nullstellensatz Proofs via Designs. In *Proof Complexity and Feasible Arithmetics*, volume 39 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 59–71. American Mathematical Society, 1998. Available at `http://www.math.ucsd.edu/~sbuss/ResearchWeb/designs/`.

19 Samuel R. Buss, Russell Impagliazzo, Jan Krajíček, Pavel Pudlák, Alexander A. Razborov, and Jiri Sgall. Proof Complexity in Algebraic Systems and Bounded Depth Frege Systems with Modular Counting. *Computational Complexity*, 6(3):256–298, 1997.

20 Samuel R. Buss and Toniann Pitassi. Good Degree Bounds on Nullstellensatz Refutations of the Induction Principle. *Journal of Computer and System Sciences*, 2(57):162–171, October 1998. Preliminary version in *CCC '96*.

21 David A. Carlson and John E. Savage. Graph Pebbling with Many Free Pebbles Can Be Difficult. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC '80)*, pages 326–332, 1980.

22 David A. Carlson and John E. Savage. Extreme Time-Space Tradeoffs for Graphs with Small Space Requirements. *Information Processing Letters*, 14(5):223–227, 1982.

23 Siu Man Chan, Massimo Lauria, Jakob Nordström, and Marc Vinyals. Hardness of Approximation in PSPACE and Separation Results for Pebble Games (Extended Abstract). In *Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS '15)*, pages 466–485, October 2015.

24 Siu Man Chan and Aaron Potechin. Tight Bounds for Monotone Switching Networks via Fourier Analysis. *Theory of Computing*, 10:389–419, October 2014. Preliminary version in *STOC '12*.

25 Ashok K. Chandra. Efficient Compilation of Linear Recursive Programs. In *Proceedings of the 14th Annual Symposium on Switching and Automata Theory (SWAT '73)*, pages 16–25, 1973.

26 Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner Basis Algorithm to Find Proofs of Unsatisfiability. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 174–183, May 1996.

27 Stephen A. Cook. An Observation on Time-Storage Trade Off. *Journal of Computer and System Sciences*, 9(3):308–316, 1974. Preliminary version in *STOC '73*.

28 Stefan S. Dantchev, Barnaby Martin, and Martin Rhodes. Tight Rank Lower Bounds for the Sherali–Adams Proof System. *Theoretical Computer Science*, 410(21–23):2054–2063, May 2009.

29 Susanna F. de Rezende, Or Meir, Jakob Nordström, Toniann Pitassi, Robert Robere, and Marc Vinyals. Lifting with Simple Gadgets and Applications to Circuit and Proof Complexity. Manuscript in preparation, 2019.

30 Cynthia Dwork, Moni Naor, and Hoeteck Wee. Pebbling and Proofs of Work. In *Proceedings of the 25th Annual International Cryptology Conference (CRYPTO '05)*, volume 3621 of *Lecture Notes in Computer Science*, pages 37–54. Springer, August 2005.

31 Yuval Filmus, Toniann Pitassi, Robert Robere, and Stephen A Cook. Average Case Lower Bounds for Monotone Switching Networks. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS '13)*, pages 598–607, November 2013.

32 John R. Gilbert, Thomas Lengauer, and Robert Endre Tarjan. The Pebbling Problem is Complete in Polynomial Space. *SIAM Journal on Computing*, 9(3):513–524, August 1980. Preliminary version in *STOC '79*.

33 Dima Grigoriev, Edward A. Hirsch, and Dmitrii V. Pasechnik. Exponential Lower Bound for Static Semi-algebraic Proofs. In *Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP '02)*, volume 2380 of *Lecture Notes in Computer Science*, pages 257–268. Springer, July 2002.

34 Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in Monotone Complexity and TFNP. Technical Report TR18-163, Electronic Colloquium on Computational Complexity (ECCC), September 2018.

35 John Hopcroft, Wolfgang Paul, and Leslie Valiant. On Time Versus Space. *Journal of the ACM*, 24(2):332–337, April 1977. Preliminary version in *FOCS '75*.

36 Russell Impagliazzo, Pavel Pudlák, and Jiří Sgall. Lower Bounds for the Polynomial Calculus and the Gröbner Basis Algorithm. *Computational Complexity*, 8(2):127–144, 1999.

37 Arist Kojevnikov and Dmitry Itsykson. Lower Bounds of Static Lovász–Schrijver Calculus Proofs for Tseitin Tautologies. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP '06)*, volume 4051 of *Lecture Notes in Computer Science*, pages 323–334. Springer, July 2006.

38 Balagopal Komarath, Jayalal Sarma, and Saurabh Sawlani. Pebbling meets coloring: Reversible pebble game on trees. *Journal of Computer and System Sciences*, 91:33–41, 2018. `doi:10.1016/j.jcss.2017.07.009`.

39 Richard Královič. Time and Space Complexity of Reversible Pebbling. *RAIRO – Theoretical Informatics and Applications*, 38(02):137–161, April 2004.

40 Klaus-Jörn Lange, Pierre McKenzie, and Alain Tapp. Reversible Space Equals Deterministic Space. *Journal of Computer and System Sciences*, 60(2):354–367, April 2000.

41 Thomas Lengauer and Robert Endre Tarjan. Asymptotically Tight Bounds on Time-Space Trade-offs in a Pebble Game. *Journal of the ACM*, 29(4):1087–1130, October 1982. Preliminary version in *STOC '79*.

42 Ming Li, John Tromp, and Paul Vitányi. Reversible Simulation of Irreversible Computation. *Physica D: Nonlinear Phenomena*, 120(1–2):168–176, September 1998.

43 Ming Li and Paul Vitányi. Reversibility and Adiabatic Computation: Trading Time and Space for Energy. *Proceedings of the Royal Society of London, Series A*, 452(1947):769–789, April 1996.

44 Jesús A. De Loera, Jon Lee, Susan Margulies, and Shmuel Onn. Expressing Combinatorial Problems by Systems of Polynomial Equations and Hilbert's Nullstellensatz. *Combinatorics, Probability and Computing*, 18(4):551–582, July 2009.

45 Giulia Meuli, Mathias Soeken, Martin Roetteler, Nikolaj Bjørner, and Giovanni De Micheli. Reversible Pebbling Game for Quantum Memory Management. *CoRR*, abs/1904.02121, 2019. `arXiv:1904.02121`.

46 Jakob Nordström. Pebble Games, Proof Complexity and Time-Space Trade-offs. *Logical Methods in Computer Science*, 9(3):15:1–15:63, September 2013.

47 Jakob Nordström. New Wine into Old Wineskins: A Survey of Some Pebbling Classics with Supplemental Results. Manuscript in preparation. To appear in *Foundations and Trends in Theoretical Computer Science*. Current draft version available at `http://www.csc.kth.se/~jakobn/research/`, 2019.

48 Michael S. Paterson and Carl E. Hewitt. Comparative Schematology. In *Record of the Project MAC Conference on Concurrent Systems and Parallel Computation*, pages 119–127, 1970.

49 Nicholas Pippenger. Pebbling. Technical Report RC8258, IBM Watson Research Center, 1980. in Proceedings of the 5th IBM Symposium on Mathematical Foundations of Computer Science, Japan.

50 Toniann Pitassi and Robert Robere. Strongly Exponential Lower Bounds for Monotone Computation. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC '17)*, pages 1246–1255, June 2017.

51 Toniann Pitassi and Robert Robere. Lifting Nullstellensatz to Monotone Span Programs over Any Field. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC '18)*, pages 1207–1219, June 2018.

52 Aaron Potechin. Bounds on Monotone Switching Networks for Directed Connectivity. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS '10)*, pages 553–562, October 2010.

53 Pavel Pudlák and Jirí Sgall. Algebraic Models of Computation and Interpolation for Algebraic Proof Systems. In *Proof Complexity and Feasible Arithmetics*, volume 39 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 279–296. American Mathematical Society, 1998. Available at `http://users.math.cas.cz/~pudlak/span.pdf`.

54 Robert Robere, Toniann Pitassi, Benjamin Rossman, and Stephen A. Cook. Exponential Lower Bounds for Monotone Span Programs. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS '16)*, pages 406–415, October 2016.

55 John E. Savage and Sowmitri Swamy. Space-Time Tradeoffs for Oblivious Interger Multiplications. In *Proceedings of the 6th International Colloquium on Automata, Languages and Programming (ICALP '79)*, pages 498–504, 1979.

56 Ravi Sethi. Complete Register Allocation Problems. *SIAM Journal on Computing*, 4(3):226–248, September 1975.

57 Sowmitri Swamy and John E. Savage. Space-Time Trade-offs on the FFT-algorithm. Technical Report CS-31, Brown University, 1977.

58 Sowmitri Swamy and John E. Savage. Space-Time Tradeoffs for Linear Recursion. *Mathematical Systems Theory*, 16(1):9–27, 1983.

59 Neil Thapen. A Trade-off Between Length and Width in Resolution. *Theory of Computing*, 12(5):1–14, August 2016.

60 Martin Tompa. Time-Space Tradeoffs for Computing Functions, Using Connectivity Properties of Their Circuits. In *Proceedings of the 10th annual ACM symposium on Theory of computing (STOC '78)*, pages 196–204, 1978.

61 Ryan Williams. Space-Efficient Reversible Simulations. Technical report, Cornell University, 2000. Available at `http://web.stanford.edu/~rrwill/spacesim9_22.pdf`.