comput. complex. (2021) 30:4

https://doi.org/10.1007/s00037-021-00201-y

computational complexity

NULLSTELLENSATZ SIZE-DEGREE TRADE-OFFS FROM REVERSIBLE PEBBLING

Susanna F. de Rezende, Or Meir, Jakob Nordström, and Robert Robere

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.

Keywords. Proof complexity, Nullstellensatz, Trade-offs, Pebbling

Subject classification. 68Q17

1. Introduction

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

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:

Published online: 12 February 2021

- 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 or not?

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.* (1994). 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

$$\sum_{i=1}^{m} p_i \cdot r_i + \sum_{j=1}^{n} (x_j^2 - x_j) \cdot s_j = 1 \tag{1.1}$$

(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 the set of polynomials $\mathcal{P} \cup \{x_j^2 - x_j \,|\, j \in [n]\}$. 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 (1.1) is the total number of monomials in all the polynomials $p_i \cdot r_i$ and $(x_j^2 - x_j) \cdot s_j$ expanded out as linear combinations of monomials. Another, more well-studied, complexity measure for Nullstellensatz is *degree*, which is defined as $\max\{\deg(p_i \cdot r_i), \deg((x_j^2 - x_j) \cdot s_j)\}$.

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 polynomials $p \in \mathcal{P}$ 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 polynomials $s \in \mathbb{F}[x_1, \ldots, x_n]$ 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 (Buss 1998). Another possible approach to prove degree lower bounds is by computationally efficient versions of Craig's interpolation theorem. It was shown in (Pudlák & Sgall 1998) 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 (Beame *et al.* 1998), induction principle (Buss & Pitassi 1998), house-sitting principle (Buss 1998; Clegg *et al.* 1996), matching (Buss *et al.* 1997), and pebbling (Buresh-Oppenheim *et al.* 2002). It seems fair to say that research in algebraic proof complexity soon moved on to stronger proof systems such as *polynomial calculus* (Alekhnovich *et al.* 2002; Clegg *et al.* 1996), 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, Nullstellensatz has been the focus of renewed interest in a recent line of works (de Rezende *et al.* 2020; Pitassi & Robere 2017, 2018; Robere *et al.* 2016) 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 (Atserias & Ochremiak 2019; Berkholz 2018).

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 (Blake 1937), 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 $n^{\mathrm{O}(d)}$. Furthermore, this upper bound has been proven to be tight up to constant factors in the exponent—that is, there are formulas that can be refuted in degree/width d but require refutations of size $n^{\Omega(d)}$ regardless of the degree/width of the refutation—for resolution and polynomial calculus (Atserias *et al.* 2016), and it follows from (Loera *et al.* 2009) 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 (Ben-Sasson & Wigderson 2001; Impagliazzo *et al.* 1999). This is known to be false for Nullstellensatz, and the pebbling formulas discussed in more detail later in this paper provide a counter-example (Buresh-Oppenheim *et al.* 2002).

The size lower bounds in terms of degree/width in (Ben-Sasson & Wigderson 2001; Impagliazzo *et al.* 1999) 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 artefact 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 by Thapen (2016), which established that there are indeed strong trade-offs between size and width making the size blow-up in (Ben-Sasson & Wigderson 2001) unavoidable. For polynomial calculus, an analogous result was obtained in (Lagarde *et al.* 2020) (after the publication of the conference version (de Rezende *et al.* 2019) of the current paper).

In this work, we complete the picture by showing that there are strong trade-offs between size and degree also 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 & Hewitt (1970), 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 *pebbling price* of G is the minimum space required to pebble G without any constraint on time.

The pebble game has been used to study flowcharts and recursive schemata (Paterson & Hewitt 1970), register allocation (Sethi 1975), time and space as Turing-machine resources (Cook 1974; Hopcroft *et al.* 1977), and algorithmic time-space trade-offs (Chandra 1973; Savage & Swamy 1978, 1979; Swamy & Savage 1983; Tompa 1978). In the last two decades, pebble games have seen a revival in the context of proof complexity (see, e.g., Nordström

2013), and pebbling has also turned out to be useful for applications in cryptography (Alwen & Serbinenko 2015; Dwork *et al.* 2005). An excellent overview of the first decade of pebbling research can be found in (Pippenger 1980), and another in-depth treatment of some classic results can be found in (Savage 1998, Chapter 10). Some more recent developments are covered in the upcoming survey (Nordström 2020).

Bennett (1989) introduced the *reversible pebble game* as part of a broader program (Bennett 1973) 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 by Meuli *et al.* (2019) 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.

We refer to the minimum space required for a reversible pebbling of a graph G as the *reversible pebbling price* of G and for a (standard) pebbling of G—without the extra restriction on pebble removals—as the *standard pebbling price*.

The reversible pebble game has been studied in (Komarath *et al.* 2018; Královič 2004; Li & Vitányi 1996) and has been used to prove time-space trade-offs in reversible simulations of irreversible computation in (Buhrman *et al.* 2001; Lange *et al.* 2000; Li *et al.* 1998; Williams 2000). In a different context, Potechin (2010) implicitly used reversible pebbling to obtain lower bounds in monotone space complexity, with the connection made explicit in later works (Chan & Potechin 2014; Filmus *et al.* 2013). In (Torán & Wörz 2020), reversible pebbling was used as a tool to study space complexity in tree-like resolution as compared to general resolution.

Chan *et al.* (2015) (to which paper 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 (Carlson & Savage 1980, 1982; Lengauer & Tarjan 1982). 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 that have the following form. Suppose that we have a DAG G such that:

1. G can be reversibly pebbled in space s_1.

2. G can be reversibly pebbled in time t_1 and space $s_2 \gg s_1$.

3. There is no reversible pebbling of G that simultaneously achieves space s_1 and time t_1. More specifically, any reversible pebbling of G in space slightly less than s_2 must take time $t_2 \gg t_1$.

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 degree s_1.
2. Nullstellensatz can also refute Peb_G in simultaneous size t_1+1 and space $s_2 \gg s_1$.
3. There is no Nullstellensatz refutation of Peb_G that simultaneously achieves degree s_1 and size $t_1 + 1$. More specifically, any Nullstellensatz refutation of Peb_G in degree slightly less than s_2 must have size $t_2 + 1 \gg t_1 + 1$.

We prove four such trade-off results, which can be found in Section 4. The following theorem (which is a simplified version of Theorem 4.1) is one example of such a result.

THEOREM 1.2. *There is a family of 3-CNF formulas* $\{F_n\}_{n=1}^{\infty}$ *of size* $\Theta(n)$ *such that:*

(i) There is a Nullstellensatz refutation of F_n *in degree* $s_1 = \mathrm{O}\big(\sqrt[6]{n}\log n\big)$.

(ii) There is a Nullstellensatz refutation of F_n *of near-linear size and degree* $s_2 = \mathrm{O}\big(\sqrt[3]{n}\log n\big)$.

(iii) 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 (Alwen *et al.* 2017; Ben-Sasson 2009; Ben-Sasson & Nordström 2011) to obtain size-space trade-offs for resolution, and in (Beck *et al.* 2013) 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 the Nullstellensatz proof system, it was shown by Buresh-Oppenheim *et al.* (2002) that Nullstellensatz degree is lower-bounded by standard pebbling price. This was strengthened by de Rezende *et al.* (2020), who 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 de Rezende *et al.* (2020) by constructing a more direct reduction, inspired by Göös *et al.* (2019), that can simultaneously capture both time and space.

1.4. Outline of this paper. After having discussed the necessary preliminaries in Section 2, we present the reductions between Nullstellensatz and reversible pebblings in Section 3. In Section 4, we prove time-space trade-offs for reversible pebblings in order to obtain size-degree trade-offs for Nullstellensatz. 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

$$\sum_{i=1}^{m} p_i(\vec{x}) \cdot r_i(\vec{x}) + \sum_{j=1}^{n} (x_j^2 - x_j) \cdot s_j(\vec{x}) = 1 \,, \tag{2.1}$$

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

As discussed in the introduction, the Nullstellensatz proof system can be used also 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 of the form (2.1) is $\max\{\deg(p_i(\vec{x}) \cdot r_i(\vec{x})), \deg((x_j^2 - x_j) \cdot s_j(\vec{x}))\}$. We define the *size* of a refutation 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 $\mathit{mSize}(p)$ 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 (2.1) is

$$\sum_{i=1}^{m} \mathit{mSize}\big(p_i(\vec{x})\big) \cdot \mathit{mSize}\big(r_i(\vec{x})\big) + \sum_{j=1}^{n} 2 \cdot \mathit{mSize}\big(s_j(\vec{x})\big) \; . \tag{2.2}$$

This is consistent with how size is defined for the "dynamic version" of Nullstellensatz known as polynomial calculus (Alekhnovich *et al.* 2002; Clegg *et al.* 1996), and also with the general size definitions for so-called algebraic and semialgebraic proof systems in (Atserias *et al.* 2016; Atserias & Ochremiak 2019; Berkholz 2018).

We remark that this is not the only possible way of measuring size, however. It can be noted that the definition (2.2) 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 (Dantchev *et al.* 2009; Grigoriev *et al.* 2002; Kojevnikov & Itsykson 2006), instead define size in terms of the polynomials in isolation, more along the lines of

$$\sum_{i=1}^{m} \Big(\mathit{mSize}\big(p_i(\vec{x})\big) + \mathit{mSize}\big(r_i(\vec{x})\big)\Big) + \sum_{j=1}^{n} \Big(2 + \mathit{mSize}\big(s_j(\vec{x})\big)\Big) \; , \tag{2.3}$$

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 (2.2) and (2.3) 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 (2.2), 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. We refer the reader to Section 2.4 in (Atserias & Hakoniemi 2019) for a more detailed discussion of the definition of proof size in algebraic and semialgebraic proof systems.

When proving lower bounds for algebraic proof systems it is often convenient to consider a *multilinear* setting where refutations are presented in the ring $\mathbb{F}[\vec{x}]/\{x_j^2 - x_j \mid j \in [n]\}$, so that no variable appears raised to a higher power than 1 in any polynomial. Since the Boolean axioms $x_j^2 - x_j$ are no longer needed, the refutation (2.1) can be written simply as

$$\sum_{i=1}^{m} p_i(\vec{x}) \cdot r_i(\vec{x}) = 1\,, \tag{2.4}$$

where we assume that all results of multiplications are implicitly multilinearized. It is clear that any refutation on the form (2.1) 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.

2.2. Reversible pebbling and pebbling formulas. In what follows, $G = (V, E)$ will always denote 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 $\mathit{pred}_G(v)$ to denote the sets of all predecessors of v, and drop the subscript when the DAG G 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 $\mathbb{P}_0 = \mathbb{P}_t = \emptyset$, $z \in \bigcup_{0\le t \le t} \mathbb{P}_t$, 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 $\mathit{time}(\mathcal{P}) = t$ and the *space* is $\mathit{space}(\mathcal{P}) = \max_{0\le t\le t}\{|\mathbb{P}_t|\}$.

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. For technical reasons, it is sometimes important to distinguish between *visiting pebblings*, for which $z \in \mathbb{P}_t$, and *persistent pebblings*, which meet the more stringent requirement that $z \in \mathbb{P}_t = \{z\}$. (It can be noted that for the more relaxed standard pebble game discussed in the introductory section any pebbling can be assumed to be persistent without loss of generality.)

Pebble games can be encoded in CNF by so-called *pebbling formulas* (Ben-Sasson & Wigderson 2001), also referred to as *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;

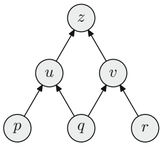

(a) Pyramid.

$$\begin{array}{l} x_p \\ \wedge\ x_q \\ \wedge\ x_r \\ \wedge\ (\overline{x}_p \vee \overline{x}_q \vee x_u) \\ \wedge\ (\overline{x}_q \vee \overline{x}_r \vee x_v) \\ \wedge\ (\overline{x}_u \vee \overline{x}_v \vee x_z) \\ \wedge\ \overline{x}_z \end{array}$$

(b) CNF formula.

$$\begin{array}{l} 1 - x_p \\ 1 - x_q \\ 1 - x_r \\ x_p x_q (1 - x_u) \\ x_q x_r (1 - x_v) \\ x_u x_v (1 - x_z) \\ x_z \end{array}$$

(c) Polynomials.

Figure 2.1: Example pebbling contradiction for the pyramid graph of height 2 in CNF and translated to polynomials

- 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, then $pred(v) = \emptyset$), 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

$$A_v := x_{\mathrm{pred}(v)} \cdot (1 - x_v)\,, \tag{2.5}$$

and for the sink z we also have the polynomial

$$A_{\mathrm{sink}} := x_z \ . \tag{2.6}$$

See Figure 2.1 for an illustration, including how the CNF formula is translated to a set of polynomials.

3. 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, which can be stated formally as follows.

THEOREM 3.1. *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 in* *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 3.1 separately in Sections 3.1 and 3.2 below.

3.1. From pebbling strategies to Nullstellensatz refutations. We start by proving the "only if" direction of Theorem 3.1. Let

$$\mathcal{P} = (\mathbb{P}_0, \ldots, \mathbb{P}_t) \tag{3.2}$$

be a 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 $\mathcal{P}$ 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 $\mathcal{P}$ 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

$$x_{\mathbb{P}_0} - x_{\mathbb{P}_{t'}} = 1 - x_{\mathbb{P}_{t'}} \ . \tag{3.3}$$

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

$$x_{\mathbb{P}_{t'}} = A_{\text{sink}} \cdot x_{\mathbb{P}_{t'} - \{z\}} \ . \tag{3.4}$$

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 ith 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., $\text{pred}(v_i) \subseteq \mathbb{P}_{i-1} \cap \mathbb{P}_i$. 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 ith step. Finally, let b_i

be 1 if v_i was pebbled during the ith step or -1 if v_i was unpebbled. Now, observe that

$$x_{\mathbb{P}_{i-1}} - x_{\mathbb{P}_i} = b_i \cdot x_{\mathbb{P}_{i-1}-\{v_i\}}(1 - x_{v_i}) = b_i \cdot x_{R_i} A_{v_i}\,, \tag{3.5}$$

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

$$\begin{aligned}\sum_{i=1}^{t'} A_{v_i} \cdot b_i \cdot x_{R_i} + x_{\mathbb{P}_{t'}} &= x_{\mathbb{P}_{t'}} + \sum_{i=1}^{t'} \left(x_{\mathbb{P}_{i-1}} - x_{\mathbb{P}_i}\right) \\ &= x_{\mathbb{P}_{t'}} + (x_{\mathbb{P}_0} - x_{\mathbb{P}_{t'}}) = 1\ .\end{aligned} \tag{3.6}$$

Note that it is a multilinear Nullstellensatz refutation, since it contains only multilinear monomials and does not use the Boolean axioms. It remains to analyse the 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. Therefore, the maximal degree is exactly the space of the pebbling strategy $\mathcal{P}$.

As for the size of the refutation, using the definition in (2.2) we obtain

$$\begin{aligned}&\sum_{i=1}^{t'} \mathit{mSize}\big(A_{v_i}\big) \cdot \mathit{mSize}\big(b_i \cdot x_{R_i}\big) + \mathit{mSize}\big(x_{\mathbb{P}_{t'}}\big) = \\ &\qquad\qquad = \sum_{i=1}^{t'} 2 \cdot 1 + 1 = t + 1\,,\end{aligned} \tag{3.7}$$

where for the first equality we recall that $\mathit{mSize}\big(A_{v_i}\big) = 2$ for every vertex v_i.

3.2. From Nullstellensatz refutations to pebbling strategies. We turn to prove the "if" direction of Theorem 3.1. 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

$$\sum_{v \in V} A_v \cdot q_v + A_{\text{sink}} \cdot q_{\text{sink}} = 1 \tag{3.8}$$

be a multilinear Nullstellensatz refutation of ϕ over $\mathbb{F}$ of degree s. We will use this refutation to construct a reversible pebbling strategy $\mathcal{P}$ for G. Let us note right away that without loss of generality we have that no monomial m in any q_v in (3.8) contains the variable x_v, since if so the factor $1 - x_v$ in A_v will make $m \cdot A_v$ cancel due to multilinearity. We will use this observation in what follows.

To extract a pebbling strategy $\mathcal{P}$ for G from the Nullstellensatz refutation (3.8), we construct a "configuration graph" $\mathcal{C}$, 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 $\mathcal{C}$ will be determined by the polynomials q_v of the refutation, and every edge $\{U_1, U_2\}$ in $\mathcal{C}$ will constitute a legal move in the reversible pebbling game (i.e., it will be legal to transition from U_1 to U_2 and vice versa). We will show that $\mathcal{C}$ 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 $\mathcal{C}$ are defined as follows: Let $v \in V$ be a vertex of G, and let m be a monomial of q_v in (3.8) (where, as observed above, we can assume that m does not contain x_v). Let $W \subseteq V$ be the set of vertices such that $m = x_W$ (such a set W exists since the refutation is multilinear). We put an edge e_m in $\mathcal{C}$ that connects $U_1 = W \cup \mathrm{pred}(v)$ and $U_2 = W \cup \mathrm{pred}(v) \cup \{v\}$ (we allow parallel edges). It is easy to see that the edge e_m connects configurations of size at most s, and that it indeed constitutes a legal move in the reversible pebbling game. We note that $\mathcal{C}$ is a bipartite graph: to see this, note that every edge connects a configuration of odd size to a configuration of even size.

For the sake of analysis, we assign weights to edges in $\mathcal{C}$ in the following way. Let $e_m = \{U_1, U_2\}$ be an edge as defined above and let c be the coefficient of m in q_v. Note that e_m represents an occurrence of the monomial x_{U_1} with coefficient c and of x_{U_2} with coefficient $-c$ in the polynomial $A_v \cdot q_v$. We assign the edge e_m a weight in $\mathbb{F}$ that is equal to $(-1)^{|U_1|} \cdot c = (-1)^{|U_2|} \cdot (-c)$. Observe that both sides of the equation are indeed the same since every edge connects an odd-sized to an even-sized configuration.

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 3.9. *Let* $U \subseteq V$ *be a configuration in* $\mathcal{C}$ *that does not contain the sink* z*. If* U *is the unique empty configuration, then its weight is* 1*. Otherwise, its weight is* 0*.*

We now show how to construct the required pebbling strategy $\mathcal{P}$ for G. To this end, we first prove that there is a path in $\mathcal{C}$ 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 $\mathcal{H}$ be the connected component of $\mathcal{C}$ that contains the empty configuration. Note that the empty configuration cannot be an isolated vertex, since it has weight 1 according to our claim. What our assumption says is that none of the configurations in $\mathcal{H}$ contains z.

The connected component $\mathcal{H}$ is bipartite since $\mathcal{C}$ is bipartite. Without loss of generality, assume that the empty configuration is in the left-hand side of $\mathcal{H}$. 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 $\mathcal{H}$. However, the sum of the weights of the configurations on the right-hand side is 0 (since all these weights are 0 by Claim 3.9), while the sum of the weights of the left-hand side is 1 (again, by Claim 3.9). We reached a contradiction, and therefore, $\mathcal{H}$ 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

$$\mathcal{P} = (\mathbb{P}_0, \ldots, \mathbb{P}_{t'-1}, \mathbb{P}_{t'}, \mathbb{P}_{t'-1}, \ldots, \mathbb{P}_0) \ . \tag{3.10}$$

This is a valid pebbling strategy since, as noted above, every edge of $\mathcal{C}$ constitutes a legal move in the reversible pebbling game. The strategy $\mathcal{P}$ uses space s, since all the configurations in $\mathcal{C}$ contain at most s pebbles by definition. The time of $\mathcal{P}$ is $t = 2 \cdot t'$. It therefore remains to show that the size of the Nullstellensatz refutation in (3.8) is at least $t + 1$.

To this end, note that every edge e_m in the path corresponds to some monomial m in some polynomial q_v. When the monomial m is multiplied by the axiom A_v, it generates two monomials in the proof: the monomial $m \cdot x_{\mathrm{pred}(v)}$ and the monomial $m \cdot x_{\mathrm{pred}(v)} \cdot x_v$. 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 contain at least one monomial, since the refutation must use the sink axiom A_{sink} (because without this axiom the rest of ϕ is satisfiable, and so cannot have any Nullstellensatz refutation). It follows that the refutation contains at least $2 \cdot t' + 1 = t + 1$ monomials, as required. We conclude the proof of the "if" direction of Theorem 3.1 by establishing Claim 3.9.

Proof (Proof of Claim 3.9.). A monomial m may be generated multiple times in the refutation (3.8). We refer to each time it is generated as an *occurrence* of m and say that such an occurrence is *generated by a monomial m_v of q_v* in (3.8) if m appears in the product $A_v \cdot m_v$.

We first prove the claim for the non-empty case. Let $U \subseteq V$ be a non-empty configuration such that $z \notin U$. We would like to prove that the weight of U is 0. Note that by definition the weight of U is equal to the sum of the weights of all the edges that touch U, i.e., $(-1)^{|U|}$ times the sum of the coefficients of the occurrences of x_U generated by monomials m_v of q_v in (3.8). Since $z \notin U$, these are all occurrences of x_U in (3.8)—i.e., x_U can only be generated by products $A_v \cdot q_v$ and can never appear in $A_{\mathrm{sink}} \cdot q_{\mathrm{sink}} = x_z \cdot q_{\mathrm{sink}}$—and so the (multi-)set of edge weights for edges incident to U in our configuration graph $\mathcal{C}$ is precisely the (multi-)set of coefficients (multiplied by $(-1)^{|U|}$) of all occurrences of x_U in (3.8). But from (3.8), we can also see that the sum of these coefficients must be 0 in $\mathbb{F}$, since the coefficient of x_U on the right-hand side is 0. Hence, the weight of U is 0.

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 (3.8). $\square$

3.3. An alternative perspective. Another way to view this proof is by considering a Nullstellensatz refutation as a solution to a matrix equation in the following manner. Let $\mathcal{P} = \{p_i : i \in [m]\}$ be a set of polynomials with no common root. Given $d \in \mathbb{N}$, let M_d be a matrix such that

- the rows of M_d are indexed by all multilinear monomials x_U of degree at most d, with the first row indexed by the monomial 1;
- the columns of M_d are indexed by all the (multilinearized) products of the form $p_i \cdot x_V$ where x_V is a multilinear monomial such that $\deg(p_i \cdot x_V) \leq d$; and
- the entry of M_d at $(x_U, p_i \cdot x_V)$ contains the coefficient of the monomial x_U in the polynomial $p_i \cdot x_V$.

It is not hard to see that $\mathcal{P}$ has a multilinear Nullstellensatz refutation of degree at most d if and only if the equation $M_d \cdot x = (1, 0, 0, \ldots, 0)^T$ has a solution. Moreover, the size of the refutation is the sum of the number of nonzero entries in all columns i such that the ith entry of x is nonzero.

In order to prove the "only if" direction of Theorem 3.1, that is, that from a pebbling strategy in time t and space s we can extract a Nullstellensatz refutation of size at most $t + 1$ and degree at most s, it is enough to show how to translate a pebbling strategy into a solution to $M_d \cdot x = (1, 0, 0, \ldots, 0)^T$. This can be argued along the lines of what was done in Section 3.1.

The other direction—that from a Nullstellensatz refutation we can extract a pebbling strategy—is where this perspective proves more elucidating. The crucial observation here is that in the specific case of pebbling contradictions the matrix M_d is *totally unimodular*, that is, the determinant of every square submatrix of M_d is in $\{0, \pm 1\}$. Indeed, it is easy to see that this matrix satisfies the following sufficient condition for total unimodularity.

FACT 3.11 (Heller & Tompkins 1957). *Let A be matrix over $\mathbb{F}$ with entries in $\{0, \pm 1\}$.*

- *If the characteristic of $\mathbb{F}$ is not 2, and every column of A contains at most one 1 and at most one -1, then A is totally unimodular.*

◦ *If the characteristic of* $\mathbb{F}$ *is* 2*, and every column of* A *contains at most two* 1*'s, then* A *is totally unimodular.*

Now consider a Nullstellensatz refutation of size $t+1$ and degree d. Let $M = M_d$ be the matrix defined by the pebbling contradiction and let x^* be the solution to the equation $Mx = (1,0,0,\ldots,0)^T$ corresponding to this Nullstellensatz refutation. The proof proceeds in two steps. First we show that there exists a solution y^* that satisfies the following two conditions: (1) y^* has entries in $\{0,\pm 1\}$, and (2) the support of y^* is contained in the support of x^*. This latter condition implies that y^* corresponds to a Nullstellensatz refutation of size at most $t+1$ (the fact that the degree of this refutation is at most d follows by definition of M). In the second step, we show that from y^*, we can extract a pebbling strategy in time at most t and space at most d.

To show that such a solution y^* exists, we use the following known property of totally unimodular matrices, which can be proven by Cramer's rule.

Proposition 3.12. *Let* A *be a totally unimodular matrix. If the equation* $Ax = (1,0,0,\ldots,0)^T$ *has a solution then it has a solution with entries in* $\{0,\pm 1\}$.

Now let I be the support of x^*, and let M' be the restriction of M to the columns in I. Clearly, the matrix M' is totally unimodular and the equation $M' \cdot y = (1,0,0,\ldots,0)^T$ has a solution. Thus, by Proposition 3.12, it has a solution y' with entries in $\{0,\pm 1\}$. Let y^* be the vector of same dimension as x^* that is equal to y' in all coordinates in I and is equal to 0 in all other coordinates. It is easy to see that y^* is a solution to $Mx = (1,0,0,\ldots,0)^T$ that satisfies the two required conditions.

We now show that from y^* we can extract a pebbling strategy in time at most t and space at most d. As in Section 3.2, this can be done by first defining the configuration graph $\mathcal{C}$ and proving there is a path of length at most $t/2$ from the empty configuration to a configuration that contains the sink z, and then showing how to extract a pebbling from such a path. We sketch the first part below—which is simpler since the entries of y^* are in $\{0,\pm 1\}$—but omit the second part since it is exactly the same as in Section 3.2.

We can view M as an incidence matrix of a graph $\mathbb{M}$: the rows of M determine the vertices of $\mathbb{M}$ and the columns of M with two nonzero entries (i.e., a column that come from some axiom A_v) determine the edges. Note that the nonzero entries of y^* define a subgraph $\mathcal{C}$ of $\mathbb{M}$ with at most $t/2$ edges. Moreover, the vertex corresponding to row 1 has odd degree in $\mathcal{C}$ (since the first entry in My^* is 1) and vertices corresponding to monomials that do not contain x_z have even degree in $\mathcal{C}$. Therefore, there must be a path in $\mathcal{C}$ of length at most $t/2$ from row 1 to a row corresponding to a monomial that contains x_z.

We conclude by remarking that, in light of this perspective, a key ingredient for the equivalence between reversible pebbling and Nullstellensatz refutations is that the matrix corresponding to a pebbling refutation is totally unimodular. Moreover, this also gives an explanation as to why degree and size of Nullstellensatz refutations of pebbling contradictions are independent of the field.

4. Nullstellensatz trade-offs from pebbling

In this section, we prove Nullstellensatz refutation size-degree trade-offs for different degree regimes. In what follows, by a Nullstellensatz refutation of a CNF formula F we mean a Nullstellensatz refutation of the translation of F to a set of polynomials as described in Section 2.

In order to obtain our trade-offs, we are looking for non-decreasing and suitably well-behaved functions $d_1(n)$ and families of CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that

1. The formula F_n has a Nullstellensatz refutation of (small) degree $d_1(n)$.

2. The formula F_n has a Nullstellensatz refutation of (close to) linear size, but in (much larger) 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 $n^{d_1(n)}$.

Below, we present explicit constructions of formulas providing such trade-offs in several different parameter regimes. We start by giving an overview of the kind of results we are able to achieve, and then spend the rest of the section on proving the reversible pebbling trade-offs that together with Theorem 3.1 yield these Nullstellensatz size-degree trade-offs.

We remark that a simple trick to achieve some of these results would be to glue together two different formulas (over disjoint set of variables) that have very different properties with respect to proof size and degree, similarly to what was done for other pairs of complexity measures for the resolution proof system in (Nordström 2009). However, the fact that two disjoint formulas can yield a "trade-off result" in this sense when glued together does not seem to be too interesting. Intuitively, we want to find *one single* formula that exhibits this trade-off behaviour. One way of formalizing this is to require that the formulas in question be minimally unsatisfiable (i.e., that removing any axiom of the formula would make it satisfiable). It is straightforward to verify that the pebbling formulas we study in this paper have this minimal unsatisfiability property.

Our first trade-off result says that there are formulas that require exponential size in Nullstellensatz if the degree is bounded by some (sublinear) polynomial function, but that for slightly larger degree admit nearly linear-size proofs.

THEOREM 4.1. *There exists a constant $K > 0$ and a family of explicitly constructible unsatisfiable 3-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that for any constant $\epsilon > 0$:*

(i) There is a Nullstellensatz refutation of F_n in degree $d_1 = \mathrm{O}\big(\sqrt[6]{n}\log n\big)$.

(ii) There is a Nullstellensatz refutation of F_n of size $\mathrm{O}(n^{1+\epsilon})$ and degree

$$d_2 = \mathrm{O}\big(d_1 \cdot \sqrt[6]{n}\big) = \mathrm{O}\big(\sqrt[3]{n}\log n\big) \ .$$

(iii) 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 4.2. *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 for any constant $\epsilon > 0$:*

(i) There is a Nullstellensatz refutation of F_n in degree $d_1 = g(n)\log(n)$.

(ii) There is a Nullstellensatz refutation of F_n of size $\mathrm{O}(n^{1+\epsilon})$ and degree

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

(iii) 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 $\log^2(n)$ all the way up to $n/\log(n)$.

THEOREM 4.3. *There exists a constant $K > 0$ and a family of explicitly constructible unsatisfiable 3-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that for any constant $\delta > 0$:*

(i) There is a Nullstellensatz refutation of F_n in degree $d_1 = \mathrm{O}(\log^2 n)$.

(ii) There is a Nullstellensatz refutation of F_n of size $\mathrm{O}(n)$ and degree

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

(iii) 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 4.4. *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 $d = n^{\Theta(1)} < n$ has size $\Theta\big(n^2/d\big)$.*

As already mentioned, we prove these results by obtaining analogous time-space trade-offs for reversible pebblings and then applying the tight correspondence between size and degree in Nullstellensatz and time and space in reversible pebbling in Theorem 3.1. We proceed to establishing such reversible pebbling trade-offs. Recall that, as mentioned in Section 2.2, we assume that all DAGs under consideration have a single sink, denoted z, and that every other vertex has some path to this sink. Some of the graph constructed below have multiple sinks, but we will explain how to turn them into single-sink DAGs.

4.1. Upper bounds for reversible pebbling time-space trade-offs. Our strategy for proving reversible pebbling trade-offs will be to analyse standard pebbling trade-offs. Clearly, lower bounds from standard pebbling transfer to reversible pebbling; the next theorem shows how, in a limited sense, we can also transfer *upper bounds*. It is based on a reversible simulation of irreversible computation proposed by Bennett (1989) and analysed precisely by Levin & Sherman (1990).

Theorem 4.5 (Bennett 1989; Levin & Sherman 1990). *Let G be an arbitrary DAG and suppose G has a standard pebbling in space s and time $t \geq 2s$. Then for any $\epsilon > 0$, G can be reversibly pebbled in time $t^{1+\epsilon}/s^{\epsilon}$ using $\epsilon(2^{1/\epsilon} - 1)\, s \log(t/s)$ pebbles.*

We also use the following general proposition, which allows upper-bounding the reversible pebbling price of a DAG by its depth and maximum indegree. Here, the *depth* of a DAG is the number of edges in a longest directed path in it, and we remind the reader that *persistent* pebblings were defined in Section 2.2.

Proposition 4.6. *Any DAG with maximum indegree ℓ and depth d has a persistent reversible pebbling strategy in space at most $d\ell + 1$.*

PROOF. The proof is by induction on the depth. For $d = 0$, we can clearly persistently reversibly pebble the graph with 1 pebble.

For $d \geq 1$, we first pebble all the (at most ℓ) predecessors of the sink persistently, so that all these vertices are covered by pebbles but there are no other pebbles in the graph. By the induction hypothesis, this can be done in space at most $(\ell-1)+(d-1)\ell+1 = d\ell$. Now we place a pebble on the sink, and then run the previous pebbling in reverse. Clearly, this adds 1 to the space, so that the total space is at most $d\ell + 1$ as claimed. □

Some of the family of DAGs considered in this section are more naturally described as DAGs with multiple sinks and have been studied as such in the pebbling literature. For the purpose of the analysis, we adopt the commonly used definition of (reversible) pebbling of a multi-sink graph: a (reversible) pebbling that places pebbles on each sink at some point (the pebbles do not need to be present in the last configuration). Let G be a DAG with m sinks and let T be a directed binary tree (arbitrary but fixed) of depth $\lceil \log m \rceil$, with m leaves all being sources and the root being the only sink. We define the single-sink DAG $\widehat{G}$ to be the graph obtained by identifying the sinks of G with the sources of T. We refer to T as the *top binary tree* of $\widehat{G}$. Note that $|V(\widehat{G})| = |V(G)| + m - 1$. Moreover, it is not hard to see that G and $\widehat{G}$ have similar pebbling bounds. We state formally below the relations between these two graphs that we use.

LEMMA 4.7. *Let G be a DAG with m sinks. We have the following properties of the single-sink DAG $\widehat{G}$.*

(i) If G has reversible pebbling price s then $\widehat{G}$ has reversible pebbling price at most $s + 2\lceil \log m \rceil + 1$.

(ii) If G has a standard pebbling in simultaneous time t and space s then $\widehat{G}$ has a standard pebbling in simultaneous time at most $t + 2(m-1)$ and space at most $s + m$.

PROOF. By Proposition 4.6, we can reversibly pebble a depth-$\lceil \log m \rceil$ binary tree in space $2\lceil \log m \rceil + 1$. To prove item (i), we simulate this pebbling on the top binary tree of $\widehat{G}$ and every time we have to pebble (or unpebble) a leaf of the tree, which coincides with some sink of G, say z_i, we simulate the space s reversible pebbling of G until the moment when we would pebble z_i (except that, in order not to interfere with the pebbling of the top binary tree, we skip steps that would place or remove pebbles from other sinks of G). Let $\mathcal{P}$ be this (partial) simulation of the reversible pebbling of G. We then pebble (or unpebble) z_i, and reverse $\mathcal{P}$ in order to remove any pebbles not on the top binary tree. Note that this adds at most an extra s pebbles on top of the space required for pebbling a depth-$\lceil \log m \rceil$ binary tree.

Let $\mathcal{P}'$ be a standard pebbling of G in time t and space s. To prove item (ii), we first simulate $\mathcal{P}'$ on $\widehat{G}$ except that we do not remove pebbles from the leaves of the top binary tree (i.e., from the sinks of G). Note that this takes time at most $t - m$ and space at most $s + m$. At this point, we only have pebbles on the m leaves of the top binary tree. We can now finish pebbling the binary tree in space $m + 1$ and time $m + 2(m - 1)$. The claimed upper bounds on space and time follow. □

We remark that Lemma 4.7 is very similar to Observation 3.8 in (Nordström 2020), and that it would not be hard to strengthen item (ii) in the lemma to get a pebbling in the same time that uses only space $s + \mathrm{O}(\log m)$ by making slightly stronger assumptions, but since we only care about the asymptotics here we opted for a slightly simpler proof instead.

4.2. Carlson-Savage graphs. The first family of graphs for which we present reversible pebbling trade-offs are the so-called *Carlson-Savage graphs* described next. Carlson & Savage (1980, 1982) defined these graphs with the goal of proving robust trade-offs for the standard pebble game. We refer the reader to Figure 4.1 for an illustration (noting that for this and other graph descriptions below we are relying heavily on Nordström 2020).

DEFINITION 4.8. [Carlson-Savage graphs (Carlson & Savage 1980, 1982; Nordström 2012)] *The graph family $\Gamma(c, r)$, for $c, r \in \mathbb{N}^+$, is defined by induction over the parameter r. The base graph $\Gamma(c, 1)$ is a DAG consisting of two sources s_1, s_2 and c sinks $\gamma_1, \ldots, \gamma_c$*

with directed edges (s_i, γ_j) *for* $i = 1, 2$ *and* $j = 1, \ldots, c$ *from both sources to all sinks. The graph* $\Gamma(c, r+1)$ *has* c *sinks and is built from the following components:*

- c *disjoint copies* $\Pi_r^{(1)}, \ldots, \Pi_r^{(c)}$ *of a so-called pyramid graph of height/depth* r*.*
- *one copy of* $\Gamma(c, r)$*.*
- c *disjoint and identical path graphs, which we call* spines*, where each spine is composed of* r sections *and every section contains* $2c$ *vertices.*

The above components are connected as follows: In every section of every spine, each of the first c *vertices has an incoming edge from the sink of one of the first* c *pyramids, where the* i*th section vertex is connected to the* i*th pyramid, and each of the last* c *vertices has an incoming edge from one of the sinks of* $\Gamma(c, r)$*, with the* i*th vertex in the second half of the section connected to the* i*th sink.*

Note that $\Gamma(c, r)$ has c sinks and maximum indegree 2. We focus for now on these graphs and only later consider their single-sink version as per Lemma 4.7. Carlson and Savage showed that the graphs $\Gamma(c, r)$ are of size $\Theta\big(cr^3 + c^2r^2\big)$ and satisfy the following property.

THEOREM 4.9 (Carlson & Savage 1982). *If* $\mathcal{P}$ *is a standard pebbling of* $\Gamma(c, r)$ *in space less than* $(r+2) + s$*, for* $0 < s \leq c - 3$*, then*

$$\mathit{time}(\mathcal{P}) \geq \left(\frac{c-s}{s+1}\right)^r \cdot r! \,.$$

This lower bound holds for space up to $c + r - 1$. By allowing only a constant factor more pebbles it is possible to pebble the graph in linear time in the standard pebble game.

LEMMA 4.10 (Nordström 2012). *The graphs* $\Gamma(c, r)$ *have standard pebbling strategies in simultaneous space* $\mathrm{O}(c + r)$ *and time linear in the size of the graphs.*

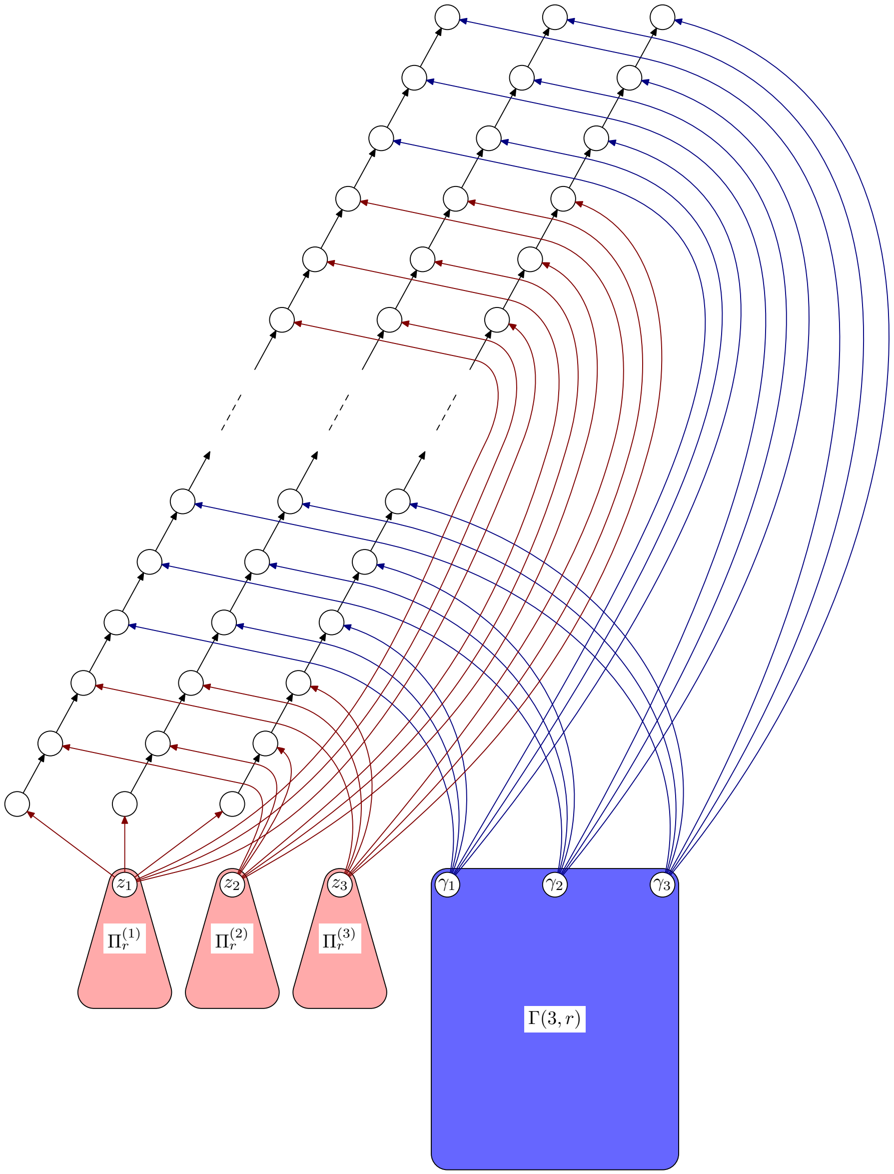

Figure 4.1: Inductive definition of Carlson-Savage graph $\Gamma(3, r + 1)$ with 3 spines and sinks

Carlson and Savage also proved that the standard pebbling price of the graph $\Gamma(c, r)$ is $r + 2$. This upper bound does not carry over to reversible pebbling, because the path graph requires more pebbles in reversible pebbling than in standard pebbling. However, we can adapt the standard pebbling strategy to reversible pebbling using the following fact.

Proposition 4.11 (Li & Vitányi 1996). *The visiting reversible pebbling price of the path graph on n vertices is $\lceil \log(n+1) \rceil$, and the persistent reversible pebbling price is $\lfloor \log(n-1) \rfloor + 2$.*

Using this result, we get the following upper bound (which is slightly stronger then what we would get by applying Theorem 4.5).

Lemma 4.12. *The graphs $\Gamma(c, r)$ have reversible pebbling price at most $r(\log(cr) + 3)$.*

Proof. The proof is by induction on r. Clearly, $\Gamma(c, 1)$ can be reversibly pebbled with 3 pebbles.

In order to pebble any sink of $\Gamma(c, r)$ for $r \geq 2$, we can reversibly pebble the corresponding spine with the space-efficient strategy for reversibly pebbling a path graph (as per Proposition 4.11). In order to pebble and unpebble a vertex on the spine, we will also need to have a pebble on the sink of the subgraph $\Pi^{(i)}_{r-1}$ or $\Gamma(c, r-1)$ connected to the spine vertex, and we will achieve this by reversibly pebbling the appropriate subgraph. By Proposition 4.6, pyramids of depth $r - 1$ can be reversibly pebbled with $2(r - 1) + 1$ pebbles, and by the induction hypothesis sinks of $\Gamma(c, r - 1)$ can be pebbled with $(r-1)(\log(c(r-1))+3) \geq 2(r-1)+1$ pebbles. Therefore, by induction on r we get that the reversible pebbling price of $\Gamma(c, r)$ is at most $(r-1)(\log(c(r-1))+3)+\log(cr)+3 \leq r(\log(cr)+3)$. $\square$

We can now choose different values for the parameters c and r and obtain graphs with trade-offs in different space regimes. The first family of graphs we consider are those that exhibit exponential time lower bounds.

Theorem 4.13. *There exists a constant $K > 0$ and an explicitly constructible family of DAGs $\{G_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ and maximum indegree 2 such that for any constant $\epsilon > 0$:*

(i) The graph G_n has reversible pebbling price $s_1 = \mathrm{O}\big(\sqrt[6]{n} \log n\big)$.

(ii) There is a reversible pebbling of G_n in time $\mathrm{O}(n^{1+\epsilon})$ and space

$$s_2 = \mathrm{O}\big(s_1 \cdot \sqrt[6]{n}\big) = \mathrm{O}\big(\sqrt[3]{n} \log n\big) \ .$$

(iii) Any standard pebbling of G_n in space at most

$$s = \frac{K s_2}{\log n} = \mathrm{O}\big(\sqrt[3]{n}\big)$$

must take time at least $\big(\sqrt[6]{n}\big)!$.

PROOF. Let $G_n = \widehat{\Gamma}(c(n), r(n))$ be the single-sink graph obtained from $\Gamma(c(n), r(n))$ as per Lemma 4.7 for parameters $c(n) = \sqrt[3]{n}$ and $r(n) = \sqrt[6]{n}$. Since $\Gamma(c(n), r(n))$ has size $\Theta\big(c(n)(r(n))^3 + (c(n))^2(r(n))^2\big) = \Theta(n)$, so does G_n. By Lemma 4.7, item (i) follows from Lemma 4.12, and item (ii) follows from applying Theorem 4.5 to Lemma 4.10. Finally, item (iii) follows from Theorem 4.9. ☐

It is also interesting to consider families of graphs that can be reversibly pebbled in very small space, close to the logarithmic lower bound on the number of pebbles required to reversibly pebble a single-sink DAG. In this small-space regime, we cannot expect exponential time lower bounds, but we can still obtain superpolynomial ones.

THEOREM 4.14. *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 DAGs $\{G_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ and maximum indegree 2 such that for any constant $\epsilon > 0$:*

(i) The graph G_n has reversible pebbling price $s_1 \leq g(n)\log(n)$.

(ii) There is a reversible pebbling of G_n in time $\mathrm{O}(n^{1+\epsilon})$ and space

$$s_2 = \mathrm{O}\big(s_1 \cdot n^{1/2}/g(n)^2\big) = \mathrm{O}\big(n^{1/2} \log n/g(n)\big) \ .$$

(iii) Any standard pebbling of G_n in space at most

$$s = \mathrm{O}\big(s_2/n^{\delta} \log n\big) = \mathrm{O}\big(n^{1/2-\delta}/g(n)\big)$$

requires time superpolynomial in n.

Proof. The proof is analogous to that of Theorem 4.14 with parameters $r(n) = g(n)$ and $c(n) = n^{1/2}/g(n)$. □

We note that in the second items of both the foregoing theorems, we could have reduced the time of the reversible pebbling to $\mathrm{O}\big(n^{1+\mathrm{o}(1)}\big)$ by applying Theorem 4.5 with $\epsilon = \mathrm{O}(1/\log\log n)$. This would have come at a cost of an extra logarithmic factor in the corresponding space bounds.

Given Theorem 3.1, which proves the tight correspondence between reversible pebbling and Nullstellensatz refutations, Theorem 4.1 follows from Theorem 4.13, and Theorem 4.2 from Theorem 4.14.

4.3. Stacks of superconcentrators. Lengauer & Tarjan (1982) also studied robust superpolynomial trade-offs for standard pebbling and showed that there are graphs that have standard pebbling price $\mathrm{O}(\log^2 n)$, but for which any standard pebbling in space up to $Kn/\log n$, for some constant K, requires superpolynomial time. For reversible pebbling, we get almost the same result for the same family of graphs.

Theorem 4.15. *There exists a constant $K > 0$ and an explicitly constructible family of DAGs $\{G_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ and maximum indegree 2 such that for any constant $\delta > 0$:*

(i) The graph G_n has reversible pebbling price $s_1 = \mathrm{O}(\log^2 n)$.

(ii) There is a reversible pebbling of G_n in time $\mathrm{O}(n)$ and space

$$s_2 = \mathrm{O}(s_1 \cdot n/\log^{3-\delta} n) = \mathrm{O}(n/\log^{1-\delta} n) \ .$$

(iii) Any standard pebbling $\mathcal{P}_n$ of G_n using at most pebbles $s = \frac{Ks_2}{\log^{\delta} n} = O(n/\log n)$ requires time $n^{\Omega(\log\log n)}$.

Note that together with Theorem 3.1 this implies Theorem 4.3. In order to describe the graphs in Theorem 4.15, we need to introduce the notion of superconcentrators.

A directed acyclic graph G is an *m-superconcentrator* if it has m sources $S = \{s_1, \ldots, s_m\}$, m sinks $Z = \{z_1, \ldots, z_m\}$, and for any subsets S' and Z' of sources and sinks of size $\lvert S'\rvert = \lvert Z'\rvert = \ell$

it holds that there are ℓ vertex-disjoint paths between S' and Z' in G.

Pippenger (1977) proved that there are superconcentrators of linear size, constant indegree and logarithmic depth, and Gabber & Galil (1981) gave the first explicit construction. It is easy to see that we can modify these superconcentrators so that the maximum indegree is 2 by substituting each vertex with indegree $\delta > 2$ by a binary tree with δ leaves. Note that this only increases the size and the depth by constant factors. Let us write this down as a formal statement.

THEOREM 4.16 (Gabber & Galil 1981). *There are explicitly constructible m-superconcentrators with $\mathrm{O}(m)$ vertices, maximum indegree 2 and depth $\mathrm{O}(\log m)$.*

Given an m-superconcentrator G_m, we define a stack of r superconcentrators G_m to be r disjoint copies of G_m where each sink of the ith copy is connected to a different source of the $(i+1)$st copy for $i \in [r-1]$. Since these graphs have m sinks, we will later apply Lemma 4.7 to obtain single-sink DAGs. Lengauer & Tarjan (1982) proved the following theorem for stacks of superconcentrators.

THEOREM 4.17 (Lengauer & Tarjan 1982). *Let $\Phi(m, r)$ denote a stack of r (explicitly constructible) linear-size m-superconcentrator with maximum indegree 2 and depth $\log m$, as per Theorem 4.16. Then the following holds:*

(i) The standard pebbling price of $\Phi(m, r)$ is $\mathrm{O}(r \log m)$.

(ii) There is a linear-time standard pebbling strategy $\mathcal{P}$ for $\Phi(m, r)$ with $\mathsf{space}(\mathcal{P}) = \mathrm{O}(m)$.

(iii) If $\mathcal{P}$ is a standard pebbling strategy for $\Phi(m, r)$ in space $s \leq m/20$, then $\mathsf{time}(\mathcal{P}) \geq m \cdot \left(\frac{rm}{64s}\right)^r$.

With this result in hand we can now proceed to prove Theorem 4.15.

PROOF (Proof of Theorem 4.15). Let $G_n = \widehat{\Phi}(n/\log n, \log n)$ be the single-sink DAG obtained from $\Phi(n/\log n, \log n)$ as per

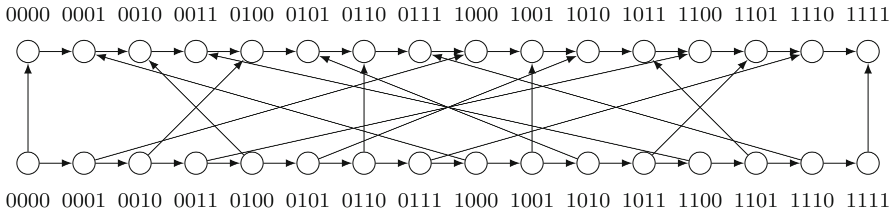

Figure 4.2: A bit-reversal permutation graph

Lemma 4.7. Note that G_n has $\Theta(n)$ vertices, indegree 2 and depth $\mathrm{O}(\log^2 n)$. By Proposition 4.6, we have that G_n can be reversibly pebbled with $\mathrm{O}(\log^2 n)$ pebbles, proving item (i).

By using Lemma 4.7 together with Theorem 4.5 with $\epsilon = 1/(\delta \log\log n)$ applied to item (ii) in Theorem 4.17, we conclude that G_n can be reversibly pebbled in simultaneous time $\mathrm{O}\big(n2^{1/\delta}\big)$ and space $\mathrm{O}\big(n/(\delta \log^{1-\delta} n)\big)$, from which item (ii) follows. Finally, item (iii) in the theorem follows from item (iii) in Theorem 4.17. □

4.4. Permutation graphs. Another family of graphs that has been studied in the context of standard pebbling trade-offs is that of *permutation graphs* as defined next.

Definition 4.18. *Given a permutation $\sigma \in \mathfrak{S}([n])$, the* permutation graph $G(\sigma)$ *consists of two paths $(x_1, \ldots, x_n)$ (the* bottom path*) and $(y_1, \ldots, y_n)$ (*top path*) which are connected as follows: for every $1 \leq i \leq n$, there is an edge from x_i to $y_{\sigma(i)}$.*

Lengauer & Tarjan (1982) proved that permutation graphs present the following smooth trade-off when instantiated with the permutation that reverses the binary representation of the index i (see Fig. 4.2 for an illustration).

Theorem 4.19 (Lengauer & Tarjan 1982). *Let G_n be a bit-reversal permutation graph on $2n$ vertices (for n a power of 2). For any $3 \leq s \leq n$, there is a standard pebbling of G_n in space s and time $\mathrm{O}\big(n^2/s\big)$. Moreover, any standard pebbling $\mathcal{P}_n$ in space s satisfies $\mathsf{time}(\mathcal{P}_n) = \Omega\big(n^2/s\big)$.*

We show that these graphs also present a smooth reversible pebbling trade-off and, in particular, for $s = n^{\Theta(1)}$ and $s \leq n$, any reversible pebbling $\mathcal{P}_n$ in space s satisfies $\mathsf{time}(\mathcal{P}_n) = \Omega\big(n^2/s\big)$ and there are almost matching upper bounds. To this end, we use the following proposition.

PROPOSITION 4.20. *For every natural number k, the path graph over n vertices can be reversibly pebbled using $2k \cdot n^{1/k}$ pebbles in time $2^k \cdot n$.*

PROOF. Observe that there is a standard pebbling of the path graph over n vertices using 2 pebbles and in time $2n$. The proposition follows now by applying Theorem 4.5 with $\epsilon = k/\log n$. □

Using Proposition 4.20, we obtain the following result.

THEOREM 4.21. *Let G_n be a bit-reversal permutation graph on $2n$ vertices (for n a power of 2). Then, G_n satisfies the following properties:*

(i) The reversible pebbling price of G_n is at most $2\log n + 2$.

(ii) If s satisfies $4\log n \leq s \leq 2n$ and k is such that $s = 4kn^{1/k}$, then there is a reversible strategy in simultaneous space s and time $\mathrm{O}\big(k2^{2k} \cdot n^2/s\big)$. In particular, if $s = n^\epsilon$ for some constant ϵ, the time of the strategy becomes $\mathrm{O}\big(n^2/s\big)$, where the big-oh notation hides a factor that depends on ϵ.

(iii) Any standard pebbling $\mathcal{P}_n$ of G_n must satisfy $\mathsf{time}(\mathcal{P}_n) = \Omega\big(n^2/\mathsf{space}(\mathcal{P}_n)\big)$.

PROOF. The upper bounds in items (i) and (ii) hold for any permutation graph.

For item (i), we can simulate a reversible pebbling of the top path that uses space at most $\log n + 1$ (as per Proposition 4.11), and every time we need a pebble on a vertex v of the bottom path in order to place or remove a pebble on the top path, we reversibly pebble the bottom path until v is pebbled (which can be done with $\log n + 1$ pebbles), make the move on the top path, and then unpebble the bottom path.

To obtain item (ii), we consider a two-phase strategy. In the first phase, we place $n^{1/k}$ pebbles spaced equally apart on the bottom path. We refer to these pebbles as *fixed* pebbles, since they will remain on the graph until the sink is pebbled. In the second phase, we simulate a reversible pebbling on the top path with $2kn^{1/k}$ pebbles, and every time we need a pebble on a vertex v on the bottom path to make a move on the top path we reversibly pebble v (with $2(k-1)n^{1/k}$ pebbles) from the nearest fixed pebble, make the move on the top path, and then unpebble the segment on the bottom path.

All that is left to show is that this can be done within the space budget of $4kn^{1/k}$ in time $\mathrm{O}(2^{2k} \cdot n^2/s)$. For the first phase, we reversibly pebble $n^{1/k}$ segments of length $m = n^{1-1/k}$. By Proposition 4.20, each of the segments can be reversibly pebbled using $2(k-1)n^{1/k} = 2(k-1)m^{k-1}$ pebbles in time $2^{k-1}n^{1-1/k}$. Since every segment must be pebbled and then unpebbled, the total time for the first phase is $2 \cdot 2^{k-1}n^{1-1/k} \cdot n^{1/k} = 2^k n$, and the total number of pebbles used is less than $2kn^{1/k}$, where the number of fixed pebbles is $n^{1/k}$ and $2(k-1)n^{1/k}$ pebbles are needed for pebbling each segment.

We turn to analysing the second phase. By Proposition 4.20, the top path can be reversibly pebbled in simultaneous space $2kn^{1/k}$ and time $2^k n$. For each move in the top path, we need to pebble and unpebble a segment of length at most $n^{1-1/k}$. As argued before, this can be done in simultaneous space $2(k-1)n^{1/k}$ and time $2 \cdot 2^{k-1}n^{1-1/k}$. Therefore, at any point in the pebbling strategy there are at most $2kn^{1/k}$ pebbles on the bottom path and at most $2kn^{1/k}$ pebbles on the top path, and the total time of the pebbling is at most $2^k n + 2^{2k}n^{2-1/k} \leq 4k2^{2k}n^2/s$.

Finally, item (iii) follows from the standard pebbling lower bound in Theorem 4.19. $\square$

From Theorem 4.21, we obtain the following corollary that, together with Theorem 3.1, implies Theorem 4.4.

Corollary 4.22. *Any reversible pebbling strategy $\mathcal{P}_n$ for the bit-reversal permutation graph G_n on $2n$ vertices that optimizes time* *given the space constraint $n^{\Theta(1)} < n$ exhibits a trade-off of the form $\mathsf{time}(\mathcal{P}_n) = \Theta\big(n^2/\mathsf{space}(\mathcal{P}_n)\big)$.*

5. Concluding remarks

In this paper, we prove that size and degree of Nullstellensatz refutations of pebbling formulas are exactly captured by time and space of reversible pebblings of the underlying graphs, regardless of the ambient field. 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.

An 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 have been shown for the standard pebble game, e.g., in (Gilbert *et al.* 1980). It is conceivable that a similar idea could be applied to the reversible pebbling reductions in (Chan *et al.* 2015), 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. We remark that the techniques in (Ben-Sasson & Nordström 2008, 2011) cannot establish such sharp trade-offs, since the reductions there between so-called black-white pebbling and resolution size/space are only tight up to constant factors, and for polynomial calculus the reductions in (Beck *et al.* 2013) are even more lossy.

Also, it should 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 as defined in Clegg *et al.* 1996 to *polynomial calculus resolution* in Alekhnovich *et al.* 2002). It would be nice to extend our size-degree trade-offs for Nullstellensatz to this setting, but it seems that some additional ideas would be needed to make this work.

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. We would also like to thank the anonymous reviewers for suggestions that improved the presentation of this work, and in particular for suggesting the perspective of totally unimodular matrices.

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. 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. Susanna F. de Rezende and Jakob Nordström were supported by the *Knut and Alice Wallenberg* grant KAW 2016.0066 *Approximation and Proof Complexity.* Jakob Nordström was also supported by the Swedish Research Council grant 2016-00782 and by the Independent Research Fund Denmark grant 9040-00389B. A preliminary version (de Rezende *et al.* 2019) of this work appeared in *CCC 2019.*

Publisher's Note
Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

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

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

Joël Alwen, Susanna F. de Rezende, Jakob Nordström & Marc Vinyals (2017). Cumulative Space in Black-White Pebbling and Resolution. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS '17)*, volume 67 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 38:1–38:21.

Albert Atserias & Tuomas Hakoniemi (2019). Size-Degree Trade-Offs for Sums-of-Squares and Positivstellensatz Proofs. In *Proceedings of the 34th Annual Computational Complexity Conference (CCC '19)*, volume 137 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 24:1–24:20.

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

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

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

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

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

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

Eli Ben-Sasson & Jakob Nordström (2008). Short Proofs May Be Spacious: An Optimal Separation of Space and Length in Resolution.

In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, 709–718.

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

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

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

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

Christoph Berkholz (2018). 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)*, 11:1–11:14.

Archie Blake (1937). *Canonical Expressions in Boolean Algebra*. Ph.D. thesis, University of Chicago.

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

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

Samuel R. Buss (1998). 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*, 59–71. American Mathematical Society. Available at http://www.math.ucsd.edu/~sbuss/ResearchWeb/designs/.

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

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

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

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

Siu Man Chan, Massimo Lauria, Jakob Nordström & Marc Vinyals (2015). 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)*, 466–485.

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

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

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

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

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

Susanna F. de Rezende, Or Meir, Jakob Nordström, Toniann Pitassi, Robert Robere & Marc Vinyals (2020). Lifting with Simple Gadgets and Applications to Circuit and Proof Complexity. In *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS '20)*. To appear.

Susanna F. de Rezende, Jakob Nordström, Or Meir & Robert Robere (2019). Nullstellensatz Size-Degree Trade-offs from Reversible Pebbling. In *Proceedings of the 34th Annual Computational Complexity Conference (CCC '19)*, volume 137 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 18:1–18:16.

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

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

Ofer Gabber & Zvi Galil (1981). Explicit Constructions of Linear-Sized Superconcentrators. *Journal of Computer and System Sciences* **22**(3), 407–420.

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

Dima Grigoriev, Edward A. Hirsch & Dmitrii V. Pasechnik (2002). 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*, 257–268. Springer.

Mika Göös, Pritish Kamath, Robert Robere & Dmitry Sokolov (2019). Adventures in Monotone Complexity and TFNP. In

Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS '19), volume 124 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 38:1–38:19.

Isidore Heller & Charles B. Tompkins (1957). An Extension of a Theorem of Dantzig's. In *Linear Inequalities and Related Systems. (AM-38)*, Annals of Mathematics Studies, 247–254. Princeton University Press.

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

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

Arist Kojevnikov & Dmitry Itsykson (2006). 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*, 323–334. Springer.

Balagopal Komarath, Jayalal Sarma & Saurabh Sawlani (2018). Pebbling meets coloring: Reversible pebble game on trees. *Journal of Computer and System Sciences* **91**, 33–41.

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

Guillaume Lagarde, Jakob Nordström, Dmitry Sokolov & Joseph Swernofsky (2020). Trade-offs Between Size and Degree in Polynomial Calculus. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS '20)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 72:1–72:16.

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

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

ROBERT Y. LEVIN & ALAN T. SHERMAN (1990). A Note on Bennett's Time-Space Tradeoff for Reversible Computation. *SIAM Journal on Computing* **19**(4), 673–677.

MING LI, JOHN TROMP & PAUL VITÁNYI (1998). Reversible Simulation of Irreversible Computation. *Physica D: Nonlinear Phenomena* **120**(1–2), 168–176.

MING LI & PAUL VITÁNYI (1996). Reversibility and Adiabatic Computation: Trading Time and Space for Energy. *Proceedings of the Royal Society of London, Series A* **452**(1947), 769–789.

JESÚS A. DE LOERA, JON LEE, SUSAN MARGULIES & SHMUEL ONN (2009). Expressing Combinatorial Problems by Systems of Polynomial Equations and Hilbert's Nullstellensatz. *Combinatorics, Probability and Computing* **18**(4), 551–582.

GIULIA MEULI, MATHIAS SOEKEN, MARTIN ROETTELER, NIKOLAJ BJØRNER & GIOVANNI DE MICHELI (2019). Reversible Pebbling Game for Quantum Memory Management. In *Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE '19)*, 288–291.

JAKOB NORDSTRÖM (2009). A Simplified Way of Proving Trade-off Results for Resolution. *Information Processing Letters* **109**(18), 1030–1035.

JAKOB NORDSTRÖM (2012). On the Relative Strength of Pebbling and Resolution. *ACM Transactions on Computational Logic* **13**(2), 16:1–16:43. Preliminary version in *CCC '10*.

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

JAKOB NORDSTRÖM (2020). 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/.

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

Nicholas Pippenger (1977). Superconcentrators. *SIAM Journal on Computing* **6**(2), 298–304.

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

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

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

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

Pavel Pudlák & Jiří Sgall (1998). 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*, 279–296. American Mathematical Society. Available at http://users.math.cas.cz/~pudlak/span.pdf.

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

John E. Savage (1998). *Models of Computation: Exploring the Power of Computing*. Addison-Wesley. Available at http://www.modelsofcomputation.org.

John E. Savage & Sowmitri Swamy (1979). Space-Time Tradeoffs for Oblivious Integer Multiplications. In *Proceedings of the*

6th International Colloquium on Automata, Languages and Programming (ICALP '79), 498–504.

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

John E. Savage & Sowmitri Swamy (1978). Space-Time Tradeoffs on the FFT-algorithm. *IEEE Transactions on Information Theory* **24**(5), 563–568

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

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

Martin Tompa (1978). 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)*, 196–204.

Jacobo Torán & Florian Wörz (2020). Reversible Pebble Games and the Relation Between Tree-Like and General Resolution Space. In *Proceedings of the 37th International Symposium on Theoretical Aspects of Computer Science (STACS '20)*, volume 154 of *Leibniz International Proceedings in Informatics (LIPIcs)*, 60:1–60:18.

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

Manuscript received 31 December 2019

Susanna F. De Rezende
Institute of Mathematics of the Czech Academy of Sciences
Prague, Czechia

Or Meir
University of Haifa
Haifa, Israel
ormeir2@gmail.com

Jakob Nordström
University of Copenhagen,
Copenhagen, Denmark and
Lund University
Lund, Sweden

Robert Robere
McGill University
Montreal, Canada