Towards an Understanding of Polynomial Calculus: New Separations and Lower Bounds (Extended Abstract)

Yuval Filmus[1], Massimo Lauria[2], Mladen Mikša[2], Jakob Nordström[2], and Marc Vinyals[2]

[1] University of Toronto, Toronto, Ontario M5S 2E4, Canada
[2] KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden

Abstract. During the last decade, an active line of research in proof complexity has been into the space complexity of proofs and how space is related to other measures. By now these aspects of resolution are fairly well understood, but many open problems remain for the related but stronger polynomial calculus (PC/PCR) proof system. For instance, the space complexity of many standard "benchmark formulas" is still open, as well as the relation of space to size and degree in PC/PCR.

We prove that if a formula requires large resolution width, then making XOR substitution yields a formula requiring large PCR space, providing some circumstantial evidence that degree might be a lower bound for space. More importantly, this immediately yields formulas that are very hard for space but very easy for size, exhibiting a size-space separation similar to what is known for resolution. Using related ideas, we show that if a graph has good expansion and in addition its edge set can be partitioned into short cycles, then the Tseitin formula over this graph requires large PCR space. In particular, Tseitin formulas over random 4-regular graphs almost surely require space at least $\Omega(\sqrt{n})$.

Our proofs use techniques recently introduced in [Bonacina-Galesi '13]. Our final contribution, however, is to show that these techniques provably cannot yield non-constant space lower bounds for the functional pigeonhole principle, delineating the limitations of this framework and suggesting that we are still far from characterizing PC/PCR space.

1 Introduction

Proof complexity studies how hard it is to provide succinct certificates for tautological formulas in propositional logic—i.e., proofs that formulas always evaluate to true under any truth value assignment, where these proofs are verifiable in time polynomial in their size. It is widely believed that there is no proof system where such efficiently verifiable proofs can always be found of size at most polynomial in the size of the formulas they prove. Showing this would establish $\mathsf{NP} \neq \mathsf{co\text{-}NP}$, and hence $\mathsf{P} \neq \mathsf{NP}$, and the study of proof complexity was initiated by Cook and Reckhow [16] as an approach towards this (still very distant) goal.

A second prominent motivation for proof complexity is the connection to applied SAT solving. By a standard transformation, any propositional logic formula F can be transformed to another formula F' in conjunctive normal form (CNF) such that F' has the same size up to constant factors and is unsatisfiable if and only if F is a tautology. Any algorithm for solving SAT defines a proof system in the sense that the execution trace of the algorithm constitutes a polynomial-time verifiable witness of unsatisfiability (such a witness is often referred to as a *refutation* rather than a *proof*, and we will use the two terms interchangeably in this paper). In the other direction, most modern SAT solvers can in fact be seen to search for proofs in systems studied in proof complexity, and upper and lower bounds for these proof systems hence give information about the potential and limitations of such SAT solvers.

In addition to running time, a major concern in SAT solving is memory consumption. In proof complexity, these two resources are modelled by *proof size/length* and *proof space*. It is thus interesting to understand these complexity measures and how they are related to each other, and such a study reveals intriguing connections that are also of intrinsic interest to proof complexity. In this context, it is natural to focus on proof systems at comparatively low levels in the proof complexity hierarchy that are, or could plausibly be, used as a basis for SAT solvers. Such proof systems include resolution and polynomial calculus. This paper takes as its starting point the former system but focuses on the latter.

Previous Work The *resolution* proof system was introduced in [12], and is at the foundation of state-of-the-art SAT solvers based on so-called conflict-driven clause learning (CDCL) [4, 23]. In resolution, one derives new disjunctive clauses from the clauses of the original CNF formula until contradiction is reached. One of the early breakthroughs in proof complexity was the (sub)exponential lower bound on proof length (measured as the number of clauses in a proof) obtained by Haken [19]. Truly exponential lower bounds—i.e., bounds $\exp(\Omega(n))$ in the size n of the formula—were later established in [14, 25] and other papers.

Ben-Sasson and Wigderson [11] identified *width* as a crucial resource, where the width is the size of a largest clause in a resolution proof. They proved that strong lower bounds on width imply strong lower bounds on length, and used this to rederive essentially all known length lower bounds in terms of width.

The study of space in resolution was initiated by Esteban and Torán [17], measuring the space of a proof (informally) as the maximum number of clauses needed to be kept in memory during proof verification. Alekhnovich et al. [1] later extended the concept of space to a more general setting, including other proof systems. The (clause) space measure can be shown to be at most linear in the formula size, and matching lower bounds were proven in [1, 8, 17].

Atserias and Dalmau [3] proved that space is in fact lower-bounded by width, which allowed to rederive all hitherto known space lower bounds as corollaries of width lower bounds. A strong separation of the two measures was obtained in [9], exhibiting formulas with constant width complexity but almost linear space complexity. Also, dramatic space-width trade-offs have been shown in [7], with formulas refutable in constant width and constant space where optimizing one of the measures causes essentially worst-case behaviour of the other.

Regarding the connections between length and space, it follows from [3] that formulas of low space complexity also have short proofs. For the subsystem of *tree-like resolution*, where each line in the proof can only be used once, [17] showed that length upper bounds also imply space upper bounds, but for general resolution [9] established that this is false in the strongest possible sense. Strong trade-offs between length and space were proven in [5, 10].

This paper focuses on the more powerful *polynomial calculus (PC)*[3] proof system [15], which is not at all as well understood. In a PC proof, clauses are interpreted as multilinear polynomials (expanded out to sums of monomials), and one derives contradiction by showing that these polynomials have no common root. Intriguingly, while proof complexity-theoretic results seem to hold out the promise that SAT solvers based on PC could be orders of magnitude faster than CDCL, such algebraic solvers have so far failed to be truly competitive.

Proof size[4] in PC is measured as the total number of monomials and the analogue of resolution space is the number of monomials needed in memory during verification of a proof. Resolution width translates into polynomial degree in PC. While length, space and width in resolution are fairly well understood, our understanding of the corresponding measures in PC is much more limited.

Impagliazzo et al. [21] showed that strong degree lower bounds imply strong size lower bounds. This is a parallel to the length-width relation in [11], and in fact this latter paper can be seen as a translation of [21] from PC to resolution. This size-degree relation has been used to prove exponential lower bounds on size in a number of papers, with [2] perhaps providing the most general setting.

The first lower bounds on space were reported in [1], but only sublinear bounds and only for formulas of unbounded width. The first space lower bounds for k-CNF formulas were presented in [18], and asymptotically optimal (linear) lower bounds were finally proven by Bonacina and Galesi [13]. However, there are several formula families with high resolution space complexity for which the PC space complexity has remained unknown, e.g., Tseitin formulas (encoding that the sum of all vertex degrees in an undirected graph must be even), ordering principle formulas, and functional pigeonhole principle (FPHP) formulas.

Regarding the relation between space and degree, it is open whether degree is a lower bound for space (the analogue of what holds in resolution) and also it has been unknown whether the two measures can be separated in the sense that there are formulas of low degree complexity requiring high space. However, [6] recently proved a space-degree trade-off analogous to the resolution space-width trade-off in [7] (in fact for the very same formulas). This could be interpreted as indicating

[3] Strictly speaking, to get a stronger proof system than resolution we need to look at the generalization *PCR* as defined in [1], but for simplicity we will be somewhat sloppy in this introduction in distinguishing between PC and PCR.

[4] The *length* of a proof is the number of lines, whereas *size* also considers the size of lines. In resolution the two measures are essentially equivalent. In PC size and length can be very different, however, and so size is the right measure to study.

that there should be a space-degree separation analogous to the space-width separation in resolution, and the authors of [13] suggest that their techniques might be a step towards understanding degree and proving that degree lower-bounds space, similar to how this was done for resolution width in [3].

As to size versus space in PC, essentially nothing has been known. It is open whether small space complexity implies small size complexity and/or the other way around. Some size-space trade-offs were recently reported in [6, 20], but these trade-offs are weaker than the corresponding results for resolution.

Our Results We study the relation of size, space, and degree in PC (and the stronger system PCR) and present a number of new results as described below.

1. We prove that if the resolution width of refuting a CNF formula F is w, then by substituting each variable by an exclusive or of two new variables and expanding out we get a new CNF formula $F[\oplus]$ requiring PCR space $\Omega(w)$. In one sense, this is stronger than claiming that degree is a lower bound for space, since high width complexity is a necessary but not sufficient condition for high degree complexity. In another sense, however, this is (much) weaker in that XOR substitution can amplify the hardness of formulas substantially. Nevertheless, to the best of our knowledge this is the first result making any connection between width/degree and space for polynomial calculus.
2. More importantly, this result yields essentially optimal separations between length and degree on the one hand and space on the other. Namely, taking expander graphs and making double copies of all edges, we show that Tseitin formulas over such graphs have proofs in size $\mathrm{O}(n \log n)$ and degree $\mathrm{O}(1)$ in PC but require space $\Theta(n)$ in PCR. (Furthermore, since these small-size proofs are tree-like, this shows that there is no tight correlation between size and space in tree-like PC/PCR in contrast to resolution.)
3. Using related ideas, we also prove strong PCR space lower bounds for Tseitin formulas over (simple or multi-)graphs where the edge set can be partitioned into small cycles. (The two copies of every edge in the multi-graph above form such cycles, but this works in greater generality.) In particular, for Tseitin formulas over random d-regular graphs for $d \geq 4$ we establish that an $\Omega(\sqrt{n})$ PCR space lower bound holds asymptotically almost surely.
4. On the negative side, we show that the techniques in [13] cannot prove any non-constant PCR space lower bounds for functional pigeonhole principle (FPHP) formulas. That is, although these formulas require high degree and it seems plausible that they are hard also with respect to space, the machinery developed in [13] provably cannot establish such lower bounds. Unfortunately, this seems to indicate that we are further from characterizing degree in PC/PCR than previously hoped.

Organization of This Paper The rest of this paper is organized as follows. We briefly review preliminaries in Section 2. In Section 3, we give a more detailed overview of our results and sketch some proofs. Section 4 contains some concluding remarks. Due to space constraints, most of the proofs are deferred to the full-length version of this paper.

2 Preliminaries

A *literal* over a Boolean variable x is either the variable x itself (a *positive literal*) or its negation $\neg x$ or $\overline{x}$ (a *negative literal*). It will also be convenient to use the alternative notation $x^0 = x$, $x^1 = \overline{x}$, where we identify 0 with true and 1 with false[5] (so that x^b is true if $x = b$). A *clause* $C = a_1 \vee \cdots \vee a_k$ is a disjunction of literals. We denote the empty clause by $\bot$. A clause containing at most k literals is called a *k-clause*. A *CNF formula* $F = C_1 \wedge \cdots \wedge C_m$ is a conjunction of clauses. A *k-CNF formula* is a CNF formula consisting of k-clauses.

Let $\mathbb{F}$ be a field and consider the polynomial ring $\mathbb{F}[x, \overline{x}, y, \overline{y}, \ldots]$ (where x and $\overline{x}$ are viewed as distinct formal variables). We write $[n] = \{1, \ldots, n\}$.

Definition 1 (Polynomial calculus resolution (PCR)). *A* PCR configuration *$\mathbb{P}$ is a set of polynomials in $\mathbb{F}[x, \overline{x}, y, \overline{y}, \ldots]$. A* PCR refutation *of a CNF formula F is a sequence of configurations $\{\mathbb{P}_0, \ldots, \mathbb{P}_\tau\}$ such that $\mathbb{P}_0 = \emptyset$, $1 \in \mathbb{P}_\tau$, and for $t \in [\tau]$ we obtain $\mathbb{P}_t$ from $\mathbb{P}_{t-1}$ by one of the following steps:*

Axiom download $\mathbb{P}_t = \mathbb{P}_{t-1} \cup \{p\}$, *where p is either a monomial $m = \prod_i x_i^b$ encoding a clause $C = \bigvee_i x_i^b \in F$, or a* Boolean axiom $x^2 - x$ *or* complementarity axiom $x + \overline{x} - 1$ *for any variable x (or $\overline{x}$).*

Inference $\mathbb{P}_t = \mathbb{P}_{t-1} \cup \{p\}$, *where p is inferred by* linear combination $\frac{q \quad r}{\alpha q + \beta r}$ *or* multiplication $\frac{q}{xq}$ *from polynomials $q, r \in \mathbb{P}_{t-1}$ for $\alpha, \beta \in \mathbb{F}$ and x a variable.*

Erasure $\mathbb{P}_t = \mathbb{P}_{t-1} \setminus \{p\}$, *where p is a polynomial in $\mathbb{P}_{t-1}$.*

If we drop complementarity axioms and encode each negative literal $\overline{x}$ as the polynomial $(1 - x)$, the proof system is called polynomial calculus (PC)*.*

The size $S(\pi)$ *of a PC/PCR refutation π is the number of monomials (counted with repetitions) in all downloaded or derived polynomials in π, the* (monomial) space $Sp(\pi)$ *is the maximal number of monomials (counted with repetitions)*[6] *in any configuration in π, and the* degree $Deg(\pi)$ *is the maximal degree of any monomial appearing in π. Taking the minimum over all PCR refutations of a formula F, we define the size $S_{\mathcal{PCR}}(F \vdash \bot)$, space $Sp_{\mathcal{PCR}}(F \vdash \bot)$, and degree $Deg_{\mathcal{PCR}}(F \vdash \bot)$ of refuting F in PCR (and analogously for PC).*

We can also define *resolution* in this framework, where proof lines are always clauses (i.e., single monomials) and new clauses can be derived by the *resolution rule* inferring $C \vee D$ from $C \vee x$ and $D \vee \overline{x}$. The *length* of a resolution refutation π is the number of downloaded and derived clauses, the *space* is the maximal number of clauses in any configuration, and the *width* is the size of a largest clause appearing in π (or equivalently the degree of such a monomial). Taking the minimum over all refutations as above we get the measures $L_{\mathcal{R}}(F \vdash \bot)$, $Sp_{\mathcal{R}}(F \vdash \bot)$, and $W_{\mathcal{R}}(F \vdash \bot)$. It is not hard to show that PCR can simulate resolution efficiently with respect to all these measures.

[5] Note that this is the opposite of what is found in many other papers, but as we will see shortly it is the natural choice in the context of polynomial calculus.

[6] In [1], space is defined *without* repetitions. All our results hold in this setting as well.

We say that a refutation is *tree-like* if every line is used at most once as the premise of an inference rule before being erased (though it can possibly be rederived later). All measures discussed above can also be defined for restricted subsystems of resolution, PC and PCR admitting only tree-like refutations.

Let us now describe the formulas which will be the main focus of our study.

Definition 2 (Tseitin formula). *Let $G = (V, E)$ be an undirected graph and $\chi \colon V \to \{0, 1\}$ be a function. Identify every edge $e \in E$ with a variable x_e and let $PARITY_{v,\chi}$ denote the CNF encoding of the constraint that the number of true edges x_e incident to a vertex $v \in V$ is equal to $\chi(v) \pmod 2$. Then the* Tseitin formula *over G with respect to f is $Ts(G, \chi) = \bigwedge_{v \in V} PARITY_{v,\chi}$.*

When the degree of G is bounded by d, $Ts(G, \chi)$ is a d-CNF formula with at most $2^{d-1}|V|$ clauses. We say that a vertex set U has *odd (even) charge* if $\sum_{u \in U} \chi(u)$ is odd (even). By a simple counting argument one sees that $Ts(G, \chi)$ is unsatisfiable if $V(G)$ has odd charge. Lower bounds on the hardness of refuting such unsatisfiable formulas $Ts(G, \chi)$ can be proven in terms of the expansion of G as defined next.

Definition 3 (Connectivity expansion [1]). *The* connectivity expansion *of $G = (V, E)$ is the largest c such that for every $E' \subseteq E$, $|E'| \leq c$, the graph $G' = (V, E \setminus E')$ has a connected component of size strictly greater than $|V|/2$.*

If F is a CNF formula and $f \colon \{0, 1\}^d \to \{0, 1\}$ is a Boolean function, then we can obtain a new CNF formula by substituting $f(x_1, \ldots, x_d)$ for every variable x and expanding out to conjunctive normal form. We write $F[f]$ to denote the resulting *substituted formula*, where we will be interested in substitutions with a particular kind of functions defined as follows.

Definition 4 (Non-authoritarian function [10]). *We say that a Boolean function $f(x_1, \ldots, x_d)$ is* non-authoritarian *if for every x_i and for every assignment α to x_i there exist α_0, α_1 extending α such that $f(\alpha_b) = b$ for $b \in \{0, 1\}$.*

By way of example, exclusive or (XOR), denoted $\oplus$, is clearly non-authoritarian, since regardless of the value of one variable, the other one can be flipped to make the function true or false, but standard non-exclusive or $\vee$ is not.

Let us finally give a brief overview of the framework developed in [13], which we use to prove our PCR space lower bounds.[7] A *partial partition* $\mathcal{Q}$ of a variable set V is a collection of disjoint sets $Q_i \subseteq V$. We use the notation $\bigcup \mathcal{Q} = \bigcup_{Q_i \in \mathcal{Q}} Q_i$. For two sets of partial assignments H and H' to disjoint domains, we denote by $H \times H'$ the set of assignments $H \times H' = \{\alpha \cup \beta \mid \alpha \in H \text{ and } \beta \in H'\}$. A set of partial assignments H to the domain Q is *flippable* on Q if for each variable $x \in Q$ and $b \in \{0, 1\}$ there exists an assignment $\alpha_b \in H$ such that $\alpha_b(x) = b$. We say that H *satisfies* a formula F if all $\alpha \in H$ satisfy F.

A $\mathcal{Q}$-*structured assignment set* is a pair $(\mathcal{Q}, \mathcal{H})$ consisting of a partial partition $\mathcal{Q} = \{Q_1, \ldots, Q_t\}$ of V and a set of partial assignments $\mathcal{H} = \prod_{i=1}^{t} H_i$, where

[7] The actual definitions that we use are slightly different but essentially equivalent.

each H_i assigns to and is flippable on Q_i. We write $(\mathcal{Q}, \mathcal{H}) \preccurlyeq (\mathcal{Q}', \mathcal{H}')$ if $\mathcal{Q} \subseteq \mathcal{Q}'$ and $\mathcal{H}'\!\upharpoonright_{\mathcal{Q}} = \mathcal{H}$, where $\mathcal{H}'\!\upharpoonright_{\mathcal{Q}} = \prod_{Q_i \in \mathcal{Q}} H'_i$. A structured assignment set $(\mathcal{Q}, \mathcal{H})$ *respects* a CNF formula F' if for every clause $C \in F'$ either $\mathit{Vars}(C) \cap \bigcup \mathcal{Q} = \emptyset$ or there is a set $Q \in \mathcal{Q}$ such that $\mathit{Vars}(C) \subseteq Q$ and $\mathcal{H}$ satisfies C.

Expressed in this language, the key technical definition in [13] is as follows.

Definition 5 (Extendible family). *A non-empty family $\mathcal{F}$ of structured assignment sets $(\mathcal{Q}, \mathcal{H})$ is r-extendible for a CNF formula F with respect to a satisfiable $F' \subseteq F$ if every $(\mathcal{Q}, \mathcal{H}) \in \mathcal{F}$ satisfies the following conditions.*

Size $|\mathcal{Q}| \leq r$.
Respectfulness $(\mathcal{Q}, \mathcal{H})$ *respects* F'.
Restrictability *For every $\mathcal{Q}' \subseteq \mathcal{Q}$ the restriction $(\mathcal{Q}', \mathcal{H}\!\upharpoonright_{\mathcal{Q}'})$ is in $\mathcal{F}$.*
Extendibility *If $|\mathcal{Q}| < r$ then for every clause $C \in F \backslash F'$ there exists $(\mathcal{Q}', \mathcal{H}') \in \mathcal{F}$ such that 1. $(\mathcal{Q}, \mathcal{H}) \preccurlyeq (\mathcal{Q}', \mathcal{H}')$, 2. $\mathcal{H}'$ satisfies C, and 3. $|\mathcal{Q}'| \leq |\mathcal{Q}| + 1$.*

To prove PCR space lower bounds for a formula F, it is sufficient to find an extendible family for F. All space lower bounds presented in this paper are obtained in this manner, where in addition we always have $F' = \emptyset$.

Theorem 6 ([13]). *Suppose that F is a CNF formula which has an r-extendible family $\mathcal{F}$ with respect to some $F' \subseteq F$. Then $Sp_{\mathcal{PCR}}(F \vdash \bot) \geq r/4$.*

3 Overview of Results and Sketches of Some Proofs

In this section, we give a more detailed overview with formal statements of our results, and also provide some proof sketches in order to convey the main technical ideas. As a general rule, the upper bounds we state are for polynomial calculus (PC) whereas the lower bounds hold for the stronger system PCR.

Relating PCR Space and Resolution Width The starting point of our work is the question of how space and degree are related in polynomial calculus, and in particular whether it is true that degree lower-bounds space. While this question remains wide open, we make partial progress by showing that if the resolution width of refuting a CNF formula F is large (which in particular must be the case if F requires high degree), then by making XOR substitution we obtain a formula $F[\oplus]$ that requires large PCR space. In fact, this works not only for exclusive or but for any non-authoritarian function (as defined in Definition 4). The formal statement is as follows.

Theorem 7. *Let F be a k-CNF formula and let f be any non-authoritarian function. Then $Sp_{\mathcal{PCR}}(F[f] \vdash \bot) \geq (W_{\mathcal{R}}(F \vdash \bot) - k + 1)/4$ holds over any field.*

Proof (sketch). In one sentence, the proof of Theorem 7 is by combining the concept of extendible families in Definition 5 with the combinatorial characterization of resolution width in [3]. We show that the properties of F implied by the width lower bound can be used to construct an extendible family for $F[f]$.

To make this description easier to parse, let us start by describing in somewhat more detail the width characterization in [3].

Consider the following game played on F by two players *Spoiler* and *Duplicator*. Spoiler asks about assignments to variables in F and Duplicator answers true or false. Spoiler can only remember ℓ assignments simultaneously, however, and has to forget some variable when this limit is reached. If Duplicator is later asked about some forgotten variable, the new assignment need not be consistent with the previous forgotten one. Spoiler wins the game by constructing a partial assignment that falsifies some clause in F, and the game is a Duplicator win if there is a strategy to keep playing forever without Spoiler ever reaching this goal. It was proven in [3] that this game exactly captures resolution width in the sense that Duplicator has a winning strategy if and only if $\ell \leq W_{\mathcal{R}}(F \vdash \bot)$.

Let us fix $r = W_{\mathcal{R}}(F \vdash \bot) - k + 1$ and use Duplicator's winning strategy for $\ell = W_{\mathcal{R}}(F \vdash \bot)$ to build an r-extendible family for $F[\oplus]$ (the proof for general non-authoritarian functions is very similar). Consider any assignment α reached during the game. We define a corresponding structured assignment set $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha)$ by adding a block $Q_x = \{x_1, x_2\}$ to $\mathcal{Q}_\alpha$ for every $x \in \mathrm{Dom}(\alpha)$, and let H_x contain all assignments α_x to $\{x_1, x_2\}$ such that $\alpha_x(x_1 \oplus x_2) = \alpha(x)$.

Given these structured assignment sets $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha)$, the family $\mathcal{F}$ is constructed inductively as follows. The base case is that $(\mathcal{Q}_\emptyset, \mathcal{H}_\emptyset) = (\emptyset, \emptyset)$ is in $\mathcal{F}$. To extend $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha)$ to satisfy a clause in $C[\oplus]$, we simulate a Spoiler with memory α who asks about all variables in C. Since Duplicator does not falsify C, when all variables have been queried some literal in C must be satisfied by the assignment. Fix one such variable assignment $\{x = b\}$ and add $\left(\mathcal{Q}_{\alpha \cup \{x=b\}}, \mathcal{H}_{\alpha \cup \{x=b\}}\right)$ as defined above to $\mathcal{F}$. All that remains now is to verify that this yields an extendible family as described in Definition 5 and then apply Theorem 6.

Separation of Size and Degree from Space It follows from Theorem 7 that there are formulas which have small PC refutations in constant degree but nevertheless require maximal space in PCR.

Theorem 8. *For any field $\mathbb{F}$ of characteristic p there is a family of k-CNF formulas F_n (where k depends on p) of size $\mathrm{O}(n)$ for which $Sp_{\mathcal{PCR}}(F_n \vdash \bot) = \Omega(n)$ over any field but which have tree-like PC refutations $\pi_n : F_n \vdash \bot$ over $\mathbb{F}$ of size $S(\pi_n) = \mathrm{O}(n \log n)$ and degree $Deg(\pi_n) = \mathrm{O}(1)$.*

Proof (sketch). Let us focus on $p = 2$. Consider a Tseitin formula $Ts(G, \chi)$ for any constant-degree graph G over n vertices with connectivity expansion $\Omega(n)$ and any odd-charge function χ.

From [11] we know that $W_{\mathcal{R}}(F \vdash \bot) = \Omega(n)$. It is not hard to see that XOR substitution yields another Tseitin formula $Ts(G', \chi)$ for the multi-graph G' obtained from G by adding double copies of all edges. This formula requires large PCR space (over any field) by Theorem 7. The upper bound follows by observing that the CNF encodes a linear system of equations, which is easily shown inconsistent in PC by summing up all equations in a tree-like fashion.

It follows from Theorem 8 that tree-like space in PC/PCR is not upper-bounded by tree-like size, in contrast to resolution. This is the only example we are aware of where the relations between size, degree, and space in PC/PCR differ from those between length, width, and space in resolution, so let us state this as a formal corollary.

Corollary 9. *It is* not true *in PC/PCR that tree-like space complexity is upper-bounded by the logarithm of tree-like size complexity.*

Space Complexity of Tseitin Formulas A closer analysis of the proof of Theorem 8 reveals that it partitions the edge set of G' into small edge-disjoint cycles (namely, length-2 cycles corresponding to the two copies of each original edge) and uses partial assignments that all maintain the same parities of the vertices on a given cycle. It turns out that this approach can be made to work in greater generality as stated next.

Theorem 10. *Let $G = (V, E)$ be a connected graph of bounded degree d with connectivity expansion c such that E can be partitioned into cycles of length at most b. Then it holds over any field that $Sp_{\mathcal{PCR}}(Ts(G, \chi) \vdash \bot) \geq c/4b - d/8$.*

Proof (sketch). We build on the resolution space lower bound in [1, 17], where the proof works by inductively constructing an assignment α_t for each derived configuration $\mathbb{C}_t$ (which corresponds to removing edges from G and updating the vertex charges accordingly) such that (a) α_t satisfies $\mathbb{C}_t$, and (b) α_t does not create any odd-charge component in G of size less than $n/2$. The inductive update can be performed as long as the space is not too large, which shows that contradiction cannot be derived in small space (since $\mathbb{C}_t$ is satisfiable).

To lift this proof to PCR, however, we must maintain not just one but an exponential number of such good assignments, and in general we do not know how to do this. Nevertheless, some more thought reveals that the only important aspect of our assignments are the resulting vertex parities. And if the edge set is partitioned into cycles, we can always shift edge charges along the cycles so that for all the exponentially many assignments, these parities are all the same (meaning that we only have to maintain one good assignment after all).

Some graphs, such as rectangular grids, can be partitioned into cycles of size $\mathrm{O}(1)$, yielding tight bounds on space. A bit more surprisingly, random d-regular graphs for $d \geq 4$ turn out to (sort of) admit partitions into cycles of size $\mathrm{O}(\sqrt{n})$, which yields the following theorem.

Theorem 11. *Let G be a random d-regular graph on n vertices, where $d \geq 4$. Then over any field it holds almost surely that $Sp_{\mathcal{PCR}}(Ts(G, \chi) \vdash \bot) = \Omega\big(\sqrt{n}\big)$.*

Proof (sketch). As long as we are interested in properties holding asymptotically almost surely, we can replace random 4-regular graphs with unions of two random Hamiltonian cycles [22]. We show that a graph distributed according to the latter model almost surely decomposes into cycles of length $\mathrm{O}(\sqrt{n})$, along with

εn additional edges (which are easily taken care of separately). Since random graphs are also excellent expanders, we can apply Theorem 10. The argument easily extends to random d-regular graphs for any $d \geq 4$.

We believe that the true space bound should actually be $\Theta(n)$, just as for resolution, but such a result seems beyond the reach of our current techniques. Also, note that to make Theorem 10 go through we need graph expansion *plus* partitions into small cycles. It seems plausible that expansion alone should imply PCR space lower bounds, as for resolution, but again we cannot prove this.

Limitations of the PCR Space Lower Bound Technique The framework in [13] can also be used to rederive all PCR space lower bounds shown previously in [1, 18], and in this sense [13] sums up what we know about PCR space lower bounds. There are also intriguing similarities between [13] and [3] (as partly hinted in the proof sketch for Theorem 7), which raises the question whether extendible families could perhaps be a step towards characterizing degree and showing that degree lower-bounds space in PC/PCR.

Even more intriguingly, however, there are CNF formulas for which it seems reasonable to expect that PCR space lower bounds should hold, but where extendible families seem very hard to construct. Such formulas include ordering principle formulas, functional pigeonhole principle (FPHP) formulas, and random 3-CNF formulas. In fact, no PCR space lower bounds are known for *any* 3-CNF formula—it is consistent with current knowledge that all 3-CNF formulas could have constant space complexity in PCR (!), though this seemingly absurd possibility can be ruled out for PC [18].

We show that the problems in applying [13] to the functional version of the pigeonhole principle are inherent, in that these techniques provably cannot establish *any* nontrivial space lower bound.

Theorem 12. *There is no r-extendible family for $FPHP^{n+1}_n$ for $r > 1$.*

Since by [24] these formulas require PC refutation degree $\Omega(n)$, one way of interpreting Theorem 12 is that the concept of r-extendible families is very far from providing the hoped-for characterization of degree.

One step towards proving PCR space lower bounds could be to obtain a weaker PC space lower bound—as noted above in the discussion of 3-CNF formulas, this can sometimes be easier. For $FPHP^{n+1}_n$, however, and for a slightly more general class of formulas described in the full-length version of this paper, it turns out that such PC space lower bounds would immediately imply also PCR space lower bounds.

Theorem 13. $Sp_{\mathcal{PCR}}(FPHP^{n+1}_n \vdash \perp) = \Theta(Sp_{\mathcal{PC}}(FPHP^{n+1}_n \vdash \perp))$.

4 Concluding Remarks

In this paper, following up on recent work in [6, 13, 18, 20], we report further progress on understanding space complexity in polynomial calculus and how the space measure is related to size and degree. Specifically, we separate size and degree from space, and provide some circumstantial evidence for the conjecture that degree might be a lower bound on space in PC/PCR. We also prove space lower bounds for a large class of Tseitin formulas, a well-studied formula family for which nothing was previously known regarding PCR space.

We believe that our lower bounds for Tseitin formulas over random graphs are *not* optimal, however. And for the functional pigeonhole principle, we show that the technical tools developed in [13] cannot prove any non-constant PCR space lower bounds. Although we have not been able to prove this, we believe that similar impossibility results should hold also for ordering principle formulas and for the canonical 3-CNF version of the pigeonhole principle. Since all of these formulas require large degree in PCR and large space in resolution, it is natural to suspect that they should be hard for PCR space as well. The fact that arguments along the lines of [13] do not seem to be able to establish this suggests that we are still far from a combinatorial characterization of degree analogous to the characterization of resolution width in [3]. It thus remains a major open problem to understand the relation between degree and space in PC/PCR, and in particular whether degree (or even width) is a lower bound on space or not.

Also, our separations of size and degree on the one hand and space on the other depend on the characteristic of the underlying field. It would be satisfying to find formulas that provide such separations regardless of characteristic. Natural candidates are ordering principle formulas or onto function pigeon principle formulas, or, for potentially even stronger separations, pebbling formulas.

Acknowledgements The authors wish to thank Ilario Bonacina and Nicola Galesi for numerous and very useful discussions.

The research of the first author has received funding from the European Union's Seventh Framework Programme (FP7/2007–2013) under grant agreement no. 238381. Part of the work of the first author was performed while visiting KTH Royal Institute of Technology. The other authors were funded by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement no. 279611. The fourth author was also supported by Swedish Research Council grants 621-2010-4797 and 621-2012-5645.

References

[1] M. Alekhnovich, E. Ben-Sasson, A. A. Razborov, and A. Wigderson. Space complexity in propositional calculus. *SIAM J. Comput.*, 31(4):1184–1211, 2002.

[2] M. Alekhnovich and A. A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proc. Inst. Math.*, 242:18–35, 2003.

[3] A. Atserias and V. Dalmau. A combinatorial characterization of resolution width. *J. Comput. Syst. Sci.*, 74(3):323–334, 2008.

[4] R. J. Bayardo Jr. and R. Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *Proc. 14th National Conference on Artificial Intelligence (AAAI '97)*, pages 203–208, 1997.

[5] P. Beame, C. Beck, and R. Impagliazzo. Time-space tradeoffs in resolution: Superpolynomial lower bounds for superlinear space. In *Proc. 44th Symposium on Theory of Computing (STOC '12)*, pages 213–232, 2012.

[6] C. Beck, J. Nordström, and B. Tang. Some trade-off results for polynomial calculus. In *Proc. 45th Symposium on Theory of Computing (STOC '13)*, 2013.

[7] E. Ben-Sasson. Size space tradeoffs for resolution. *SIAM J. Comput.*, 38(6):2511–2525, 2009.

[8] E. Ben-Sasson and N. Galesi. Space complexity of random formulae in resolution. *Random Struct. Algorithms*, 23(1):92–109, 2003.

[9] E. Ben-Sasson and J. Nordström. Short proofs may be spacious: An optimal separation of space and length in resolution. In *Proc. 49th Symposium on Foundations of Computer Science (FOCS '08)*, pages 709–718, 2008.

[10] E. Ben-Sasson and J. Nordström. Understanding space in proof complexity: Separations and trade-offs via substitutions. In *Proc. 2nd Symposium on Innovations in Computer Science (ICS '11)*, pages 401–416, 2011.

[11] E. Ben-Sasson and A. Wigderson. Short proofs are narrow—resolution made simple. *J. ACM*, 48(2):149–169, 2001.

[12] A. Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.

[13] I. Bonacina and N. Galesi. Pseudo-partitions, transversality and locality: A combinatorial characterization for the space measure in algebraic proof systems. In *Proc. 4th Conference on Innovations in Theoretical Computer Science (ITCS '13)*, pages 455–472, 2013.

[14] V. Chvátal and E. Szemerédi. Many hard examples for resolution. *J. ACM*, 35(4):759–768, 1988.

[15] M. Clegg, J. Edmonds, and R. Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proc. 28th Symposium on Theory of Computing (STOC '96)*, pages 174–183, 1996.

[16] S. A. Cook and R. Reckhow. The relative efficiency of propositional proof systems. *J. Symb. Log.*, 44(1):36–50, 1979.

[17] J. L. Esteban and J. Torán. Space bounds for resolution. *Inf. Comput.*, 171(1):84–97, 2001.

[18] Y. Filmus, M. Lauria, J. Nordström, N. Thapen, and N. Ron-Zewi. Space complexity in polynomial calculus. In *Proc. 27th Conference on Computational Complexity (CCC '12)*, pages 334–344, 2012.

[19] A. Haken. The intractability of resolution. *Theor. Comput. Sci.*, 39(2-3):297–308, 1985.

[20] T. Huynh and J. Nordström. On the virtue of succinct proofs: Amplifying communication complexity hardness to time-space trade-offs in proof complexity. In *Proc. 44th Symposium on Theory of Computing (STOC '12)*, pages 233–248, 2012.

[21] R. Impagliazzo, P. Pudlák, and J. Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Comput. Complex.*, 8(2):127–144, 1999.

[22] J. H. Kim and N. C. Wormald. Random matchings which induce Hamilton cycles, and hamiltonian decompositions of random regular graphs. *J. Comb. Theory B*, 81:20–44, 2001.

[23] J. P. Marques-Silva and K. A. Sakallah. GRASP—a new search algorithm for satisfiability. In *Proc. International Conference on Computer-Aided Design (ICCAD '96)*, pages 220–227, 1996.

[24] A. A. Razborov. Lower bounds for the polynomial calculus. *Comput. Complex.*, 7(4):291–324, 1998.

[25] A. Urquhart. Hard examples for resolution. *J. ACM*, 34(1):209–219, 1987.