Towards an Understanding of Polynomial Calculus: New Separations and Lower Bounds*

Yuval Filmus
University of Toronto

Massimo Lauria
KTH Royal Institute of Technology

Mladen Mikša
KTH Royal Institute of Technology

Jakob Nordström
KTH Royal Institute of Technology

Marc Vinyals
KTH Royal Institute of Technology

May 6, 2013, at 19:13

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 NP $\neq$ co-NP, and hence P $\neq$ NP, and the study of proof complexity was initiated by Cook and Reckhow [CR79] 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

*This is the full-length version of the paper [FLM+13] to appear in *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP '13).*

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.

1.1 Previous Work

The *resolution* proof system was introduced in [Bla37], and is at the foundation of state-of-the-art SAT solvers based on so-called conflict-driven clause learning (CDCL) [BS97, MS96]. 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 [Hak85]. Truly exponential lower bounds—i.e., bounds $\exp(\Omega(n))$ in the size n of the formula—were later established in [CS88, Urq87] and other papers.

Ben-Sasson and Wigderson [BW01] 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 [ET01], 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. [ABRW02] 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 [ABRW02, BG03, ET01].

Atserias and Dalmau [AD08] 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 [BN08], exhibiting a formula family with constant width complexity but almost linear space complexity. Also, dramatic space-width trade-offs have been shown in [Ben09], 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 [AD08] 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, [ET01] showed that length upper bounds also imply space upper bounds, but for general resolution [BN08] established that this is false in the strongest possible sense. Strong trade-offs between length and space were proven in [BN11, BBI12].

This paper focuses on the more powerful *polynomial calculus (PC)*[1] proof system introduced by Clegg et al. [CEI96], 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.

[1] Strictly speaking, to get a stronger proof system than resolution we need to look at the generalization *PCR* as defined in [ABRW02], but for simplicity we will be somewhat sloppy in this introduction in distinguishing between PC and PCR.

Proof size[2] in PC is measured as the total number of monomials in a proof and the analogue of resolution space is the number of monomials needed in memory during proof verification. Clause width in resolution translates into polynomial degree in PC. While length, space and width in resolution are fairly well understood as surveyed above, our understanding of the corresponding complexity measures in PC is much more limited.

Impagliazzo et al. [IPS99] showed that strong degree lower bounds imply strong size lower bounds. This is a parallel to the length-width relation in [BW01], and in fact this latter paper can be seen as a translation of the bound in [IPS99] from PC to resolution. This size-degree relation has been used to prove exponential lower bounds on size in a number of papers, with [AR03] perhaps providing the most general setting.

The first lower bounds on space were reported in [ABRW02], but only sublinear bounds and only for formulas of unbounded width. The first space lower bounds for k-CNF formulas were presented in [FLN+12], and asymptotically optimal (linear) lower bounds were finally proven by Bonacina and Galesi [BG13]. 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 (which would be 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, [BNT13] recently proved a space-degree trade-off analogous to the resolution space-width trade-off in [Ben09] (in fact for the very same formulas). This could be interpreted as indicating that there should be a space-degree separation analogous to the space-width separation in resolution, and the authors of [BG13] 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 [AD08].

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 [HN12, BNT13], but these trade-offs are weaker than the corresponding results for resolution.

1.2 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 briefly 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.)

[2] 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.

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 [BG13] 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 [BG13] provably cannot establish such lower bounds. Unfortunately, this seems to indicate that we are further from characterizing degree in PC/PCR than previously hoped.

1.3 Organization of This Paper

The rest of this paper is organized as follows. We briefly review preliminaries in Section 2. Section 3 presents a overview of our results and provides some proof sketches outlining the main technical ideas that go into the proofs.

In Section 4, we prove that resolution width lower bounds plus substitutions with XOR or other suitable Boolean functions yields PCR space lower bounds. We use this in Section 5 to separate size and degree from space in PC and PCR. In Section 6, we show PCR space lower bounds for Tseitin formulas over graphs with edge sets decomposable into partitions of small cycles. The proof that random d-regular graphs for $d \geq 4$ (almost) decompose into cycles of length $\mathrm{O}(\sqrt{n})$ is given in Section 7. The fact that PCR space lower bounds cannot be obtained for the functional pigeonhole principle formulas with current techniques is proven in Section 8, and in the same section we show that a larger class of formulas containing FPHP formulas have essentially the same space complexity for PC and PCR (so that when proving lower bounds, one can without loss of generality ignore the complementary formal variables for negative literals in PCR and focus only on PC).

We make some concluding remarks and discuss some of the (many) open questions remaining in Section 9. For completeness, in Appendix A we provide a full description of our version of the techniques in [BG13] and provide proofs that the same claims still hold in this slightly different setting.

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[3] (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. We think of clauses and CNF formulas as sets so that order is irrelevant and there is no repetitions.

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 employ the standard notation $[n] = \{1, \ldots, n\}$.

Definition 2.1 (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 $\overline{x}$).

[3] Note that this notational convention 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.

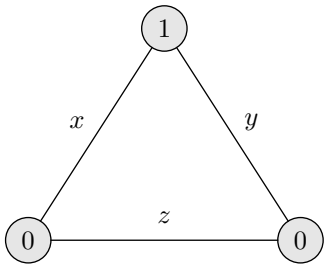

(a) Labelled triangle graph.

$$
\begin{aligned}
&(x \vee y)\\
\wedge\ &(\overline{x} \vee \overline{y})\\
\wedge\ &(x \vee \overline{z})\\
\wedge\ &(\overline{x} \vee z)\\
\wedge\ &(y \vee \overline{z})\\
\wedge\ &(\overline{y} \vee z)
\end{aligned}
$$

(b) Corresponding Tseitin formula.

Figure 1: Example Tseitin formula.

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)[4] 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.

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 family of CNF formulas which will be the main focus of our study.

Definition 2.2 (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, $PARITY_{v,\chi}$ has at most 2^{d-1} clauses, all of width at most d, and hence $Ts(G, \chi)$ is a d-CNF formula with at most $2^{d-1}|V|$ clauses. Figure 1(b) gives an example Tseitin formula generated from the graph in Figure 1(a). We say that a set of vertices 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.

[4] We note that in [ABRW02], space was defined *without* counting repetitions of monomials. All our lower bounds hold in this more stringent setting as well.

Definition 2.3 (Connectivity expansion [ABRW02]). 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 2.4 (Non-authoritarian function [BN11]). 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 [BG13], which we use to prove our PCR space lower bounds.[5] 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 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 $\mathcal{H}'\restriction_{\mathcal{Q}} = \mathcal{H}$, where $\mathcal{H}'\restriction_{\mathcal{Q}} = \prod_{Q_i \in \mathcal{Q}} H'_i$. 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 [BG13] is as follows.

Definition 2.5 (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 $(\mathcal{Q}', \mathcal{H}\restriction_{\mathcal{Q}'})$ is in $\mathcal{F}$.

Extendibility If $|\mathcal{Q}| < r$ then for every clause $C \in F \setminus F'$ 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. $|\mathcal{Q}'| \leq |\mathcal{Q}| + 1$.

When $F' = \emptyset$, we simply say that $\mathcal{F}$ is r-extendible for F.

To prove PCR space lower bounds for a formula F, it is sufficient to find an extendible family for F.

Theorem 2.6 ([BG13]). *Suppose that F is a CNF formula which has an r-extendible family $\mathcal{F}$ with respect to some $F' \subseteq F$. Then $Sp_{\mathcal{PCR}}(F \vdash \bot) \geq r/4$.*

All space lower bounds presented in this paper are obtained in this manner, where in addition we always have $F' = \emptyset$.

[5] The actual definitions that we use are slightly different but essentially equivalent. We provide the full details including proofs in Section A for completeness.

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 polynomial calculus resolution (PCR). In fact, even more can be said: just as is the case in [ABRW02, FLN+12, BG13], all our lower bounds hold also for *functional calculus*, where proof lines are arbitrary Boolean functions over clauses/monomials and anything that follows semantically from the current configuration can be derived in a single step. We do not discuss this further below but instead refer to Appendix A for the details.

3.1 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 2.4). The formal statement is as follows.

Theorem 3.1. *Let F be a k-CNF formula and let f be any non-authoritarian function. Then it holds over any field that $Sp_{PCR}(F[f] \vdash \bot) \geq (W_{\mathcal{R}}(F \vdash \bot) - k + 1)/4$.*

Proof sketch. In one sentence, the proof of Theorem 3.1 is by combining the concept of extendible families in Definition 2.5 with the combinatorial characterization of resolution width in [AD08]. 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 [AD08].

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 [AD08] 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 and is given in Section 4). 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 2.5 and then apply Theorem 2.6. □

3.2 Separation of Size and Degree from Space

An almost immediate consequence of Theorem 3.1 is that there are formulas which have small PC refutations in constant degree but nevertheless require maximal space in PCR.

Theorem 3.2. *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 $p = 2$, deferring the general proof to Section 5. 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 [BW01] 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 3.1. 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 3.2 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 3.3. *It is* not true *in PC/PCR that tree-like space complexity is upper-bounded by the logarithm of tree-like size complexity.*

3.3 Space Complexity of Tseitin Formulas

A closer analysis of the proof of Theorem 3.2 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 3.4. *Let $G = (V, E)$ be a connected graph of bounded degree d with connectivity expansion c such that the edge set 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 [ABRW02, ET01], 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, the vertex parities are all the same (meaning that on a higher level we only have to maintain one good assignment after all). The full proof is presented in Section 6. □

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 3.5. *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 [KW01]. 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 3.4. The argument extends straightforwardly to random d-regular graphs for any $d \geq 4$. The full proof, which contains a bit more by way of technical details, is given in Section 7. □

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 3.4 go through we need graph expansion *plus* partitions into small cycles. It seems plausible that expansion alone should be enough to imply PCR space lower bounds, as for resolution, but again we are not able to prove this.

3.4 Limitations of the PCR Space Lower Bound Technique

The framework in [BG13] can also be used to rederive all PCR space lower bounds shown previously in [ABRW02, FLN+12], and in this sense [BG13] sums up what we know about PCR space lower bounds. There are also intriguing similarities between [BG13] and the resolution width characterization in [AD08] (as partly hinted in the proof sketch for Theorem 3.1), 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 [FLN+12].

We show that the problems in applying [BG13] to the functional version of the pigeonhole principle are inherent, in that these techniques provably cannot establish *any* nontrivial space lower bound. We refer to Section 8 for the formal description of the formulas and the proof of the next theorem.

Theorem 3.6. *There is no r-extendible family for $FPHP_n^{n+1}$ for $r > 1$.*

Since by [Raz98] these formulas[6] require PC refutation degree $\Omega(n)$, one way of interpreting Theorem 3.6 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 $FPHP_n^{n+1}$, however, and for a slightly more general class of formulas described in Section 8, it turns out that such PC space lower bounds would immediately imply also PCR space lower bounds.

Theorem 3.7. $Sp_{\mathcal{PCR}}(FPHP_n^{n+1} \vdash \bot) = \Theta(Sp_{\mathcal{PC}}(FPHP_n^{n+1} \vdash \bot))$.

Proof sketch. In $FPHP_n^{n+1}$ we have variables $x_{i,j}$ for $i \in [n+1]$, $j \in [n]$, encoding that pigeon i goes into hole j. The clauses of the formula require that every pigeon is mapped to some hole and that this mapping is one-to-one. Because of this, the negation of $x_{i,j}$ is equivalent to $\bigvee_{j' \neq j} x_{i,j'}$ and so the literal $\overline{x}_{i,j}$ can be encoded as the monomial $\prod_{j' \neq j} x_{i,j'}$ in PC. Since this substitutes a monomial for a monomial the space does not increase. Now we can take any PCR refutation of $FPHP_n^{n+1}$ and apply such substitutions line by line. The inferences remain sound (with some local auxiliary steps added) and so this process gives a PC refutation of $FPHP_n^{n+1}$ in roughly the same space. □

[6] To be precise, the degree lower bound in [Raz98] is proven for the functional pigeonhole principle encoded as linear equations—the standard CNF version has large initial width/degree and so there is nothing to prove. However, the linear-equations encoding of FPHP has axioms of large space, and so for space lower bounds we want to study the CNF version.

4 PCR Space Lower Bounds From Resolution Width

In the rest of this paper, we give formal proofs of the results described in Section 3. We start by considering the question of relating space and degree in PCR. Although we do not know how to prove (or rule out) an analogue of the relation between space and width in resolution, we can use the combinatorial game from [AD08] to prove a weaker relation between PCR space and resolution width. Recall from the informal description of the game in Section 3.1 that we have two players, Spoiler and Duplicator, and that Duplicator needs to be able to provide an answer to any of Spoiler's questions about assignments to some bounded number of variables in order to win the game. Formally, a winning strategy for Duplicator is defined as follows.

Definition 4.1 (Duplicator's strategy [AD08]). A *Duplicator winning strategy* for the Boolean existential ℓ-pebble game on on a CNF formula F is a non-empty family $\mathcal{D}$ of partial truth value assignments to $Vars(F)$ such that every $\alpha \in \mathcal{D}$ satisfies the following conditions:

1. No clause $C \in F$ is falsified by α.
2. The domain of α has size at most $|\mathrm{Dom}(\alpha)| \leq \ell$.
3. For every subassignment $\alpha' \subseteq \alpha$ it holds that $\alpha' \in \mathcal{D}$.
4. If $|\mathrm{Dom}(\alpha)| < \ell$, then for every variable x there exists an $\alpha' \in \mathcal{D}$ that and assigns a value to x and extends α (i.e., $\alpha' \supseteq \alpha$).

In [AD08], Atserias and Dalmau proved the following tight connection between Duplicator winning strategies and resolution refutation width.

Theorem 4.2 ([AD08]). *The CNF formula F has a resolution refutation of width ℓ if and only if Duplicator has no winning strategy for the Boolean existential $(\ell+1)$-pebble game on F.*

The Duplicator strategy in Definition 4.1 has some similarities with the extendible family in Definition 2.5, which can be taken to suggest that there might be a relation between resolution width and PCR space. The main difference is that extendible families consist of sets of assignments in which we must be able to flip every variable, while Duplicator's strategy is built on fixed individual assignments. However, if we substitute every variable in F with a non-authoritarian function as defined in Definition 2.4, then it is straightforward to make the transition from fixed assignments to sets of flippable assignments.

Lemma 4.3. *Let F be a k-CNF formula and let f be a non-authoritarian function. If Duplicator wins the Boolean existential ℓ-pebble game on F, then there exists an $(\ell-k+1)$-extendible family for $F[f]$.*

Proof. Let $\mathcal{D}$ be a winning Duplicator strategy for F. We will use $\mathcal{D}$ to construct an $(\ell-k+1)$-extendible family $\mathcal{F}$ for the substituted formula $F[f]$. In what follows, let us denote by $Vars^d(x)$ the set of variables that we get when we substitute x by $f(x_1,\ldots,x_d)$ in F for some non-authoritarian function f of arity d.

For $x \in Vars(F)$, define $Q_x = Vars^d(x)$ and let $H_{x,\alpha} = \{\beta \mid \mathrm{Dom}(\beta) = Q_x \text{ and } f(\beta) = \alpha(x)\}$ be the set of all assignments over Q_x for which f evaluates to the value that α assigns to x. For any partial assignment $\alpha \in \mathcal{D}$ we let the corresponding structured assignment set $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha)$ be the pair consisting of $\mathcal{Q}_\alpha = \{Q_x \mid x \in \mathrm{Dom}(\alpha)\}$ and $\mathcal{H}_\alpha = \prod_{x\in\mathrm{Dom}(\alpha)} H_{x,\alpha}$. We define $\mathcal{F}$ to encompass all structured assignment sets $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha)$ corresponding to partial assignments $\alpha \in \mathcal{D}$ with $|\mathrm{Dom}(\alpha)| \leq \ell-k+1$. We need to prove that $\mathcal{F}$ constructed in this way is an $(\ell-k+1)$-extendible family with respect to $F' = \emptyset$.

By construction, for every $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha) \in \mathcal{F}$ we have that $\mathcal{Q}_\alpha$ is a partial partition and that the partial assignments $H_{x,\alpha} \in \mathcal{H}_\alpha$ assign to $Q_x \in \mathcal{Q}_\alpha$. Furthermore, $H_{x,\alpha}$ is flippable on Q_x. This is so since f is a non-authoritarian function, which means that for very variable in $x_i \in Q_x$ there exist assignments β_b, $b \in \{0,1\}$, to Q_x such that $\beta_b(x_i) = b$ and $f(\beta_b) = \alpha(x)$. Hence, all $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha) \in \mathcal{F}$ are structured assignment sets.

The size condition $|\mathcal{Q}_\alpha| \leq \ell-k+1$ in Definition 2.5 is clearly satisfied for all $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha) \in \mathcal{F}$, and respectfulness is vacuously true. To see that the restriction property also holds, consider any $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha) \in \mathcal{F}$

obtained from $\alpha \in \mathcal{D}$. For any subset $\mathcal{Q}' \subseteq \mathcal{Q}_\alpha$, let α' be the subassignment of α restricted to $\{x \mid Q_x \in \mathcal{Q}'\}$ and let $\mathcal{H}' = \prod_{Q_x \in \mathcal{Q}'} H_{x,\alpha} = \prod_{x \in \mathrm{Dom}(\alpha')} H_{x,\alpha'}$. Then since $\alpha' \in \mathcal{D}$ by Definition 4.1, it follows by the construction of $\mathcal{F}$ that $(\mathcal{Q}', \mathcal{H}\restriction_{\mathcal{Q}'}) = (\mathcal{Q}', \mathcal{H}') \in \mathcal{F}$ as required.

It remains to prove that $\mathcal{F}$ has the extension property. Let $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha) \in \mathcal{F}$ be such that $|\mathcal{Q}_\alpha| < \ell - k + 1$ and let C be a clause in $F[f]$. We need to argue that $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha)$ can be extended to satisfy C. Let $A \in F$ be the clause such that $C \in A[f]$, i.e., C is one of the clauses obtained when substituting f in A. If $\alpha \in \mathcal{D}$ satisfies A, it follows by construction that $\mathcal{H}_\alpha$ satisfies all of $A[f]$ and hence, in particular, C, and we are done. Otherwise, it follows from the definition of a winning Duplicator strategy and the fact that $|\alpha| \leq \ell - k$ that α can be extended to an assignment α' that queries all of the (at most k) variables in A without falsifying the clause. Such an α' must satisfy A. Fix some variable $x^* \in \mathrm{Dom}(\alpha') \setminus \mathrm{Dom}(\alpha)$ such that α' satisfies A by assigning to x^*, and let α^* be the subassignment of α' with domain $\mathrm{Dom}(\alpha) \cup \{x^*\}$. This α' must be in $\mathcal{D}$ by Definition 4.1, and analogously to what was argued above it must hold that $\mathcal{H}_{\alpha^*}$ satisfies $C \in A[f]$. It is clear that $(\mathcal{Q}_\alpha, \mathcal{H}_\alpha) \preccurlyeq (\mathcal{Q}_{\alpha^*}, \mathcal{H}_{\alpha^*})$, and that $|\mathcal{Q}_{\alpha^*}| \leq |\mathcal{Q}_\alpha| + 1$. Hence, $\mathcal{F}$ satisfies extendibility, and the lemma follows. □

Combining Lemma 4.3 with the combinatorial characterization of width in Theorem 4.2 and the lower bound on space in terms of extendible families in Theorem 2.6, we obtain the first theorem claimed in Section 3.

Theorem 3.1 (restated). *Let F be a k-CNF formula and let f be any non-authoritarian function. Then*

$$Sp_{\mathcal{PCR}}(F[f] \vdash \bot) \geq \frac{W_{\mathcal{R}}(F \vdash \bot) - k + 1}{4} \ .$$

While it can be argued that this theorem might be interpreted as an indication that degree could be a lower bound for space in PCR, a more immediate and concrete consequence is that it gives us a way to prove the existence of formulas which have very small PCR refutations, but for which any refutation must have essentially maximal space. For polynomial calculus over fields of characteristic 2, we already have all the tools needed to argue this. In particular, the space lower bound needed follows immediately from Theorem 3.1 as described next.

Corollary 4.4. *Let G be an expander graph of bounded degree over n vertices, let f be an odd-charge function on $V(G)$, and let G' be the multi-graph obtained by adding two copies of each edge in G. Then*

$$Sp_{\mathcal{PCR}}(Ts(G', f) \vdash \bot) = \Omega(n) \ .$$

Proof. As shown in [BW01], refuting Tseitin formulas over expander graphs requires linear width in resolution. It is not hard to see that substituting with XOR in a Tseitin formula over G is the same as considering the formula over the multi-graph with two copies of every edge. Thus $Ts(G', f)$ requires monomial space $\Omega(n)$ by Theorem 3.1, which is linear in the formula size if G is a constant-degree expander. □

As briefly discussed in Section 3.2, it is not hard to show that Tseitin formulas have small refutations in PCR (and even PC) over fields of characteristic 2, which yields Corollary 3.3 for this characteristic. However, this upper bound does not hold for characteristics distinct from 2. Therefore, we need to work with generalized version of Tseitin formulas and prove our results for such formulas instead. We do so in the next section.

5 Formulas With Small Proofs May Require Large Space

In Section 2 we defined Tseitin formulas as the CNF encoding of particular linear systems over $\mathbb{F}_2$. Here we consider a generalization over fields of any positive characteristic. Any such formula essentially defines an unsatisfiable linear system over $\mathbb{F}_p$ for some prime p. In order to efficiently encode this linear system as a CNF it is important that each equation mentions a small (for instance constant) number of variables: any equation over d variables can be encoded as a set of at most 2^d clauses with d literals each. In particular, Tseitin formulas are defined on directed graph as follows.

Definition 5.1. Let $G = (V, E)$ be a directed graph and $f\colon V \to \{0, 1, \ldots, p-1\}$ be a function. Identify every directed edge $(u, v) \in E$ with a variable $x_{(u,v)}$ and let $Mod^p_{v,f}$ denote the CNF encoding of the constraint that the number of *incoming* edges $x_{(u,v)}$ incident to a vertex $v \in V$ that are set to true, minus the number of *outgoing* edges $x_{(v,w)}$ set to true is equal to $f(v) \pmod p$. Then the *Tseitin formula* over G with respect to f is $Ts^p(G, f) = \bigwedge_{v\in V} Mod^p_{v,f}$.

This formula is unsatisfiable when $\sum_v f(v) \not\equiv 0 \pmod p$. Compare Definition 2.2 with Definition 5.1: for $p = 2$ the definitions coincide because is such characteristic there is no difference between the contribution of the incoming and the outgoing edges. For $p = 2$ it is natural to define the formula in terms of undirected graphs, indeed. Not surprisingly, polynomial calculus over a field of characteristic p efficiently refutes unsatisfiable Tseitin formulas defined on sums modulo p.

Lemma 5.2. *Consider a directed graph $G = (V, E)$ with n vertices and constant degree, and a function $f\colon V \to [0, p-1]$ with $\sum_v f(v) \not\equiv 0 \pmod p$. The formula $Ts^p(G, f)$ has a tree-like polynomial calculus refutation of constant degree, size $\mathrm{O}(n \log n)$, and monomial space $\mathrm{O}(n)$.*

Furthermore, given any boolean function h on a constant number of variables, the result holds for the substituted formula $Ts^p(G, f)[h]$.

Proof. Let us first consider the case without substitution. Recall that true value is encoded as 0 and false as 1. In this encoding formula $Mod^p_{v,f}$ is equivalent to

$$\sum_{u\colon (u,v)\in E} (1 - x_{uv}) - \sum_{w\colon (v,w)\in E} (1 - x_{vw}) \equiv f(v) \pmod p \ . \tag{5.1}$$

The proof is based on the natural intuition that summing the equations (5.1) for all vertices in the graph results in a contradiction, since in the sum each variable appears twice: once with positive and once with negative sign. Fix an enumeration of $V = \{v_1, \ldots v_n\}$, and fix the following notation for partial sums:

$$S_{a,b} := \sum_{i=a}^{b} \left[\sum_{u:(u,v_i)\in E} (1 - x_{uv_i}) - \sum_{w:(v_i,w)\in E} (1 - x_{v_iw}) \right] \equiv \sum_{i=a}^{b} f(v_i) \pmod p \ . \tag{5.2}$$

We fix $t = 2^{\lceil \log n \rceil} < 2n$ and consider $S_{i,i}$ to be the equation "$0 = 0$" for all $n < i \leq t$. We set up a tree of height $\lceil \log n \rceil$, where leaves are labeled by equations $S_{i,i}$ and internal nodes are labeled by the sum of the two children labels (i.e., a node at level k is labeled by the equation $S_{i,i+2^k-1}$ for some i).

Each equation $S_{i,i}$ is derived from the encoding of $Mod^p_{v_i,f}$. This equation mention only a constant number of variables, so by implicational completeness of polynomial calculus (see Lemma 5.3) we have a derivation of constant space and size.

Equations in internal nodes are derived by summing the equations of the children. We derive all the equations on the tree in a bottom-up fashion. This concludes the refutation since the equation $S_{1,t}$ at the root is

$$\sum_{i=1}^{n} \left[\sum_{u:(u,v_i)\in E} (1 - x_{uv_i}) - \sum_{w:(v_i,w)\in E} (1 - x_{v_iw}) \right] \equiv \sum_{i=1}^{n} f(v_i) \pmod p \tag{5.3}$$

$$\sum_{(u,v)\in E} (1 - x_{uv}) - \sum_{(v,w)\in E} (1 - x_{vw}) \equiv \sum_{i=1}^{n} f(v_i) \pmod p \tag{5.4}$$

$$0 \equiv \sum_{i=1}^{n} f(v_i) \pmod p \tag{5.5}$$

Which is the end of the refutation, since $\sum_{i=1}^{n} f(v_i)$ is non-zero.

The size of the proof accounts $\mathrm{O}(1)$ for the deduction of each $S_{i,i}$, and $\mathrm{O}(n)$ for the total number of monomial at each level of the tree: at level k there are $\frac{t}{2^k}$ equations with at most $\mathrm{O}(2^k)$ monomials. So the total size is as claimed.

Regarding the monomial space, notice that we need to keep simultaneously in memory only the equations of two adjacent levels, which have at most $\mathrm{O}(n)$ monomials.

The degree of the refutation is $\mathrm{O}(1)$ for the inference of each equation $S_{i,i}$. The rest of the proof has degree 1.

The case with substitution is similar: consider a substituting function h on a constant number of variables. There is a multilinear polynomial p_h which evaluates exactly as h on all $\{0,1\}$ inputs, and which mentions a constant number of monomials.

The substituted linear forms $S_{i,i}[h]$ are linear combinations of copies of p_h, so they have a constant number of variables each and their inference from $Mod^p_{v_i,f}[h]$ is doable in constant space, size and degree because of Lemma 5.3.

Once the equations $S_{i,i}[h]$ are derived, the refutation goes exactly as shown for the case with no substitution. From this point on the original refutation is linear; applying the trivial substitution to these proof lines increases the space, degree and size only by constant factors. □

For the sake of self-containment, we give a proof of the implicational completeness of polynomial calculus. This completes the proof of Lemma 5.2.

Lemma 5.3. *Consider a polynomial implication $p_1 \dots p_l \models p$ which is valid over $\{0,1\}$ assignments. Assume all involved polynomials collectively mention d variables and have degree $\mathrm{O}(d)$; then there is a PC proof of this implication in degree $\mathrm{O}(d)$, space $2^{\mathrm{O}(d)}$, and length $2^{\mathrm{O}(d)}$.*

Proof. Without loss of generality we assume that all polynomials are in multilinear form. So each of them has size at most 2^d and degree d. Let $\alpha = \{x_1 \mapsto v_1, \dots, x_d \mapsto v_d\}$ be an assignment; we define C_α as $\prod_i (v_i x_i + (1-v_i)(1-x_i))$, the polynomial which evaluates to 1 exactly on the assignment α. We list some useful observations:

Observation (1) is that given the axioms $\{x_i = v_i\}_{i\in[d]}$ and any polynomial q on variables $x_1, \dots, x_d$, it is possible to efficiently infer $q - \alpha(q) = 0$. We prove this by induction on the number of variables. If $d = 0$ then $q = \alpha(q)$. Now assume that $q - \alpha(q) = s + xt - \alpha(q)$. If we have deduced $q\restriction_{x=0} = s - \alpha(q)$ and we have the axiom x, we can easily infer xt and then $s + xt - \alpha(q)$. If we have deduced $q\restriction_{x=1}$ (which is $s + t - \alpha(q)$) and we have the axiom $x - 1$, we can easily infer $(x-1)t$ and then $s + t + (x-1)t - \alpha(q) = s + xt - \alpha(q)$. This derivation requires $\mathrm{O}(d)$ steps, one per variable, and both size and space are proportional to the number of monomials in q. The degree is equal to the degree of q plus d.

Observation (2) is that for any q on variables $x_1, \dots, x_d$, we can infer from Boolean axioms the polynomial $C_\alpha(q - \alpha(q))$, for every assignment α on such variables. The inference is in degree $O(d)$, and length and space are $2^{\mathrm{O}(d)}$. It is immediate for the simple case $q = x_i$: each $C_\alpha(x_i - v_i)$ contains the factor $x_i^2 - x_i$ by construction. For any non-trivial q we apply the inference in Observation (1), with the caveat that each line is multiplied by C_α. The resulting polynomial is $C_\alpha(q - \alpha(q))$.

Observation (3) is that $\sum_{\alpha\in\{0,1\}^d} C_\alpha = 1$, and this is an easy induction over d (it also follows from the semantic of polynomials C_α).

We now see how to deduce $C_\alpha p$ for every assignment α. For α which satisfy p we derive $C_\alpha(p - 0)$ using observation (2). For α which falsify p, pick any falsified p_i and deduce both $C_\alpha(p_i - \alpha(p_i))$ and $C_\alpha p_i$, using observations (2) and multiplication rule, respectively. The sum is $C_\alpha \alpha(p_i)$, and since $\alpha(p_i)$ is a non-zero field element, we can multiply by $\frac{p}{\alpha(p_i)}$ to get $C_\alpha p$.

Having deduced all $C_\alpha p$ we can use observation (3) to infer p. Notice that we did 2^d inferences (one for each α), each of them of degree $\mathrm{O}(d)$ and each of them in space $2^{\mathrm{O}(d)}$. □

Now we have seen that (substituted) Tseitin formulas are easy to polynomial calculus under determined conditions. Nevertheless we can use the tools from Section 4 to show that even under these conditions, any such refutations require large space.

Theorem 5.4 (restatement of Theorem 3.2). *For $\mathbb{F}$ any field 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. The formula family we consider is based on Tseitin formulas over a family of Ramanujan graphs of constant degree. This is a family of simple graphs with good expansion properties; a construction is given in [Mor94]. Consider such a graph G on m vertices: set an arbitrary orientation on the edges, and consider any $f\colon [m] \to \{0, \ldots, p-1\}$ with $\sum_i f(i) \neq 0 \mod p$.

In Corollary 4.5 of [AR03], it is claimed that if G is a d-regular graph for d at least some constant value d_p, then $Ts^p(G, f)$ requires refutations of degree $\Omega(m)$ in polynomial calculus over any field of characteristic different from p.

Polynomial calculus simulates resolution over any characteristic, and the degree of the simulation is exactly the width of the simulated resolution proof. This implies that resolution requires width $\Omega(m)$ to refute the formula.

Fix $k = 2d$. We apply a XOR substitution on formula $Ts^p(G, f)$, and we get a k-CNF formula on $n = dm$ variables. Theorem 3.1 implies that any polynomial calculus (or PCR) refutation requires monomial space $\Omega(n)$, under any characteristic.

If the characteristic of the underlying field is p the upper bound follows by Lemma 5.2. □

6 PCR Space Lower Bounds for Tseitin Formulas

In the following exposition we assume that $G = (V, E)$ is a graph with connectivity expansion c and $f\colon V \to \{0, 1\}$ is a Boolean function. We call a pair (G, f) a *charged graph*, and we say that a set of vertices U is even (odd) charged if $\sum_{v \in U} f(v)$ is even (odd). We denote the set of edges *incident* to a vertex v by $E(v)$ and extend the notation to sets of vertices. We write $\overline{\alpha}$ to denote the complementary assignment of α obtained by flipping the value of all variables in the domain $\mathrm{Dom}(\alpha)$.

Definition 6.1. The *charged graph induced by a partial assignment* α is $((V, E \setminus \mathrm{Dom}(\alpha)), g)$, where $g(v) = f(v) + \sum_{e \ni v}(1 - \alpha(e))$.

Observation 6.2. *The formulas $Ts((V, E \setminus \mathrm{Dom}(\alpha)), g)$ and $Ts(G, f)\restriction_\alpha$ are equivalent. An assignment α satisfies the clauses $PARITY_{v,g}$ if and only if the vertex v is isolated and even (as a singleton set) in the charged graph induced by α. In that case, we say that the assignment α satisfies the vertex v.*

Definition 6.3 (non-splitting assignment). A charged graph is *non-splitting* if all its connected components of size at most $n/2$ are even. A partial assignment α is *non-splitting* if the charged graph induced by α is non-splitting .

Observation 6.4. *The empty assignment is non-splitting for the charged graph (G, f) if and only if (G, f) is non-splitting. A connected graph is always non-splitting.*

Observation 6.5. *Suppose α is a partial assignment extending a partial assignment β (or conversely, $\beta = \alpha\restriction_D$ for some $D \subseteq \mathrm{Dom}(\alpha)$). If α is non-splitting, then so is β. In other words, "unsubstituting" an edge cannot result in an odd component that has size less than or equal to $n/2$ because component sizes can only increase.*

The key idea in the resolution space lower bound is that if a proof does not mention many edges, then it is possible to maintain a satisfiable assignment to the edges the proof mentions. This satisfiable assignment shifts the charged in the graph so that a contradiction only arises in vertices that the proof does not mention and leaves enough freedom to keep adding edges to the assignment unless the proof reaches a space threshold. Thus the proof is unable to derive a contradiction unless it mentions many edges at once.

The following lemma implements the charge shifting idea.

Lemma 6.6. *Let α be a non-splitting assignment. Let e be an edge. Let $D = \mathrm{Dom}(\alpha) \cup \{e\}$. If $|D| \leq c$ then we can extend α to some non-splitting assignment β such that $\mathrm{Dom}(\beta) = D$.*

Proof. Let (G', g) be the charged graph induced by α. Let $e = (u, v)$. Let C be the connected component in G' that contains the vertices u and v. Let $\alpha_0 = \alpha \cup \{e \mapsto 0\}$ and $\alpha_1 = \alpha \cup \{e \mapsto 1\}$. Let (G'', h_0) and (G'', h_1) be the charged graph induced by α_0 and α_1 respectively. Observe that $h_0(C) = h_1(C) = g(C)$.

If e is not a bridge, i.e., removing the edge e from G' does not disconnect C, then we can extend α to either α_0 or α_1. In this case there is no new component.

If e is a bridge, let C' and C'' be the components in G'' that e disconnects C into. If $g(C)$ is even, either both $h_0(C')$ and $h_0(C'')$ are even, in which case we can extend α to α_1, or both $h_0(C')$ and $h_0(C'')$ are odd, in which case we can extend α to α_0 reversing both parities. In this case all new components are even.

Otherwise, since α is non-splitting, $|C| > n/2$. Since $|D| \leq c$, the graph G'' has a connected component larger than $n/2$. The graph G' cannot have two disjoint components both larger than $n/2$, so this large component is a subset of C; either C' or C''. Assume it is C' without loss of generality. Since C is odd, either $h_0(C')$ is odd and $h_0(C'')$ is even, in which case we can extend α to α_1, or $h_0(C')$ is even and $h_0(C'')$ is odd, in which case we can extend α to α_0 reversing both parities. In this case there is one new odd component, but it is larger than $n/2$. □

Corollary 6.7. *Let α be a non-splitting assignment. Let E be a set of edges. Let $D = \mathrm{Dom}(\alpha) \cup E$. If $|D| \leq c$ then we can extend α to some non-splitting assignment β such that $\mathrm{Dom}(\beta) = D$.*

To extend this idea to a PCR lower bound for space, and in particular to the framework of [BG13], we need to use assignments that are not only non-splitting but also resilient to flips of the values of some variables.

Observe that if all the edges along a cycle change their value, the graph induced by the cycle stays the same. The following definition will let us formalize this property. Recall the cartesian product notation for sets of assignments.

Definition 6.8 (Flipped assignments). Let α be a partial assignment and let $\mathcal{Q}$ be a (total) partition of $\mathrm{Dom}(\alpha)$. The set of *flipped assignments of α with respect to $\mathcal{Q}$* is the set of assignments given by

$$\mathit{Flip}(\mathcal{Q}, \alpha) = \prod_{Q \in \mathcal{Q}} \{\alpha\restriction_Q, \overline{\alpha}\restriction_Q\} \ .$$

Observation 6.9. *If α is an assignment over a cycle C, then α and $\overline{\alpha}$ induce the same charged graph. Therefore, if $\mathcal{Q}$ is a set of disjoint cycles, all the flipped assignments of some assignment α with respect to $\mathcal{Q}$ induce the same charged graph.*

Theorem 6.10 (Strengthening of Theorem 3.4). *Let (G, f) be non-splitting charged graph of maximal degree d with connectivity expansion c such that a partition M of E into edge-disjoint cycles of length at most b exists. Then*

$$Sp_{\mathcal{PCR}}(Ts(G, f) \vdash \bot) \geq c/4b - d/8 \ .$$

Note that this is a strengthening of Theorem 3.4 since if G is connected then (G, f) is trivially non-splitting for every f.

Proof. By Theorem 2.6, it is sufficient to build an r-extendible family for $r = c/b - d/2$. Let $\mathcal{F}$ be the set of all pairs $(\mathcal{Q}, \mathcal{H}^\alpha)$ satisfying:

1. $\mathcal{Q} \subseteq M$ and $|\mathcal{Q}| \leq r$.
2. $\mathcal{H}^\alpha = \mathit{Flip}(\mathcal{Q}, \alpha)$, where α is any non-splitting assignment over $\bigcup \mathcal{Q}$.

Note that $\mathcal{Q}$ is a collection of edge-disjoint cycles and every $\mathcal{H}^\alpha$ consists of the some non-splitting assignment α and its flips over cycles. Each $(\mathcal{Q}, \mathcal{H}^\alpha) \in \mathcal{F}$ has many different representations, since $\mathcal{H}^\alpha = \mathcal{H}^\beta$ whenever $\beta \in \mathit{Flip}(\alpha, \mathcal{Q})$.

Let us show that $\mathcal{F}$ is an extendible family. First, pairs $(\mathcal{Q}, \mathcal{H}^\alpha)$ are $\mathcal{Q}$-structured by construction.

The empty assignment is non-splitting by Observation 6.4. So the family $\mathcal{F}$ is not empty because $(\emptyset, \mathcal{H}^\emptyset) \in \mathcal{F}$, where $\emptyset$ is the empty assignment.

Let us show that the family is closed under restriction. Consider any $(\mathcal{Q}, \mathcal{H}) \in \mathcal{F}$ and $\mathcal{Q}' \subseteq \mathcal{Q}$. Let $\alpha \in \mathcal{H}$, and let β be the restriction of α to $\bigcup \mathcal{Q}'$. By construction α is non-splitting, and restriction preserves the property of being non-splitting as noted in Observation 6.5, so $(\mathcal{Q}', \mathcal{H}^\beta) \in \mathcal{F}$. Finally $\mathcal{H}\restriction_{\mathcal{Q}'} = \mathit{Flip}(\mathcal{Q}, \alpha)\restriction_{\mathcal{Q}'} = \mathit{Flip}(\mathcal{Q}', \beta) = \mathcal{H}^\beta$.

Let us show that the family is closed under extension. Let $(\mathcal{Q}, \mathcal{H}) \in \mathcal{F}$ with $|\mathcal{Q}| < r$ and let $p \in \mathit{PARITY}_{v,f}$ for some vertex $v \in V$.

If $\mathcal{H}$ satisfies p we are done; otherwise we will extend a non-splitting assignment associated with $\mathcal{H}$.

Let $\alpha \in \mathcal{H}$ be a non-splitting assignment that does not satisfy p. Let $\mathcal{Q}_v = \{C \in M \mid v \in C\}$ be the cycles adjacent to v, and let $\mathcal{Q}_+ = \mathcal{Q}_v \setminus \mathcal{Q}$; we will see that $\mathcal{Q}_+$ is not empty, but we do not need to assume it now. Let $D = \mathrm{Dom}(\alpha) \cup \bigcup \mathcal{Q}_+$. By hypothesis $|\mathcal{Q} \cup \mathcal{Q}_+| < r + d/2$, and it follows that $|D| < c$. Thus we can apply Corollary 6.7 on α and $\bigcup \mathcal{Q}_+$ to extend α to a non-splitting assignment β over D.

The assignment β disconnects the component $\{v\}$ and is non-splitting, so it makes the component $\{v\}$ even. By Observation 6.2, β satisfies the vertex v. Note that β falsifies $p \cap \bigcup \mathcal{Q}$, the subclause of p with variables $\bigcup \mathcal{Q}$. If for all $C \in \mathcal{Q}_+$ the assignment β supersatisfies or falsified the subclause $p \cap C$, then there would be a non-splitting assignment in $\mathit{Flip}(\mathcal{Q}_+, \beta)$ that falsified p.

Let $C \in \mathcal{Q}_+$ be a cycle that contains one literal of p that β satisfies and one literal that β falsifies. Let $\mathcal{Q}' = \mathcal{Q} \cup \{C\}$ and let $\mathcal{H}' = \mathcal{H}^\beta$. By construction $(\mathcal{Q}', \mathcal{H}') \in \mathcal{F}$, and assignments in $\mathcal{H}'$ restricted to C satisfy p. □

Theorem 3.4 is somewhat restrictive, in that it requires us to partition *all* edges in the graph into short cycles. However, as the following corollary shows, it is enough to partition *most* of the edges.

Corollary 6.11. *Let (G, f) be a non-splitting charged graph of maximal degree d with connectivity expansion c such that a partition M of E into edge-disjoint cycles of length at most b and an additional number of $t < c$ edges exist. Then*

$$\mathit{Sp}_{\mathcal{PCR}}(\mathit{Ts}(G, f) \vdash \bot) \geq (c - t)/4b - d/8$$

Proof. Let H be the graph obtained by removing the t extra edges. Note that the connectivity expansion of H is at least $c - t$. Corollary 6.7 on the preceding page shows that there exists a non-splitting assignment α on $G \setminus H$. Observation 6.2 on page 14 implies that for some g, (H, g) is a non-splitting charged graph. By a restriction argument, any PCR refutation of a non-splitting Tseitin formula on G in space S can be translated to a PCR refutation of a non-splitting Tseitin formula on H in space at most S. Theorem 3.4 shows that $S \geq (c - t)/4b - d/8$. □

6.1 Application: Grid Graphs

There are families of graphs where we actually get matching upper and lower bounds for PCR space. One such family is square grids. For the following subsection let n be an even integer and denote $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$, the integers modulo n. The following defines a grid over a torus.

Definition 6.12 (Grid graph). The *grid graph* (or discrete torus) $T(n)$ is a 4-regular graph with vertices $V = \mathbb{Z}_n \times \mathbb{Z}_n$ and edges

$$E = \left\{ \big((i,j),(i+1,j)\big), \big((i,j),(i,j+1)\big) \middle| i,j \in V \right\} .$$

We order the vertices of $T(n)$ lexicographically: $(i,j) < (k,l)$ if $i < k$ or $i = k$ and $j < l$. The *predecessor* of a vertex $(i,j) \neq (1,1)$, denoted $\mathit{pred}(i,j)$, is the vertex immediately preceding (i,j) in this order.

We will explicitly refer to the edges we need to disconnect a set of vertices from a graph. This notion is known as edge boundary.

Definition 6.13. Let $G(V, E)$ be a graph and $U \subseteq V$ be a subset of vertices. The *edge boundary* of U is the set of edges $\partial_e(U) = \{(x, y) \in E : x \in U,\ y \notin U\}$.

We can find an upper bound on PC space by mentioning all the vertices in lexicographical order.

Lemma 6.14. *The space of refuting a Tseitin formula over the $n \times n$ grid graph for an odd charge function f over characteristic 2 is $Sp_{\mathcal{PC}}(Ts(T(n), f) \vdash \bot) = \mathrm{O}(n)$.*

Proof sketch. Observe that for every set of vertices U it holds that $\sum_{e \in E(U)} e \equiv \sum_{e \in \partial_e(U)} e \pmod 2$, and that in PC over characteristic 2 this expression corresponds to the polynomial $\sum_{e \in \partial_e(U)} e$. Thus, we can express $\sum_{e \in E(U)} e \equiv f(U)$ in space $\partial_e(U)$. If we let $U_{ij} = \{(a, b) \in V \mid (a, b) \leq (i, j)\}$, the edge boundary of any U_{ij} is at most $2n + 1$, so the monomial space of each of the polynomials $p_{ij} = \sum_{e \in \partial_e(U_{ij})} e - f(U_{ij})$ is at most $2n + 1 = \mathrm{O}(n)$.

If we show how to derive the polynomials p_{ij} in lexicographical order in $\mathrm{O}(n)$ space, we will be done. And indeed, for any vertex (i, j) we can infer the polynomial $q_{ij} = \sum_{e \ni (i,j)} e - f(v)$ by downloading the 2^{d-1} axioms $PARITY_{(i,j),f}$ and adding all of them in constant space. To derive p_{ij} from $p_{pred(ij)}$ it is enough to add the polynomials $p_{pred(ij)}$ and q_{ij}. The maximum space is $Sp(p_{pred(ij)}) + Sp(p_{ij}) + \mathrm{O}(1) = \mathrm{O}(n)$. □

The connectivity expansion follows from the following isoperimetric inequality.

Theorem 6.15 ([BL91]). *Let U be a subset of vertices of $T(n)$ with $|U| \leq n^2/2$. Then*

$$|\partial_e(U)| \geq \min\{2n, 4|U|^{1/2}\} \ .$$

Corollary 6.16. *The connectivity expansion of $T(n)$ is $2n - 1$.*

Proof. If we erase $2n - 1$ or less edges from $T(n)$, then by Theorem 6.15 the largest region we can disconnect has size $|U| \leq \lfloor (2n - 1)/4 \rfloor^2 < n^2/2$, so $c \geq 2n - 1$. If we erase the $2n$ edges $\{((i, 0), (i, 1)) \mid i \in \mathbb{Z}_n\} \cup \{((i, n/2), (i, n/2 + 1)) \mid i \in \mathbb{Z}_n\}$ we obtain two connected components of size $n^2/2$, so $c < 2n$. □

The lower bound on PCR space follows.

Corollary 6.17. *The space of refuting a Tseitin formula over the $n \times n$ grid graph (over any characteristic) is $Sp_{\mathcal{PCR}}(Ts(T(n), f) \vdash \bot) = \Omega(n)$.*

Proof. Let us find a partition of the edges of $T(n)$. Let $C(i, j)$ be the set of edges of the cycle $\big((i, j), (i + 1, j), (i + 1, j + 1), (i, j + 1)\big)$. Then the set $M = \{C(i, j) \mid i + j \equiv 0 \pmod 2\}$ is a partition of the edges of $T(n)$ into edge-disjoint cycles of length 4. By Theorem 2.6, $Sp_{\mathcal{PCR}}(Ts(T(n), f) \vdash \bot) \geq (2n - 9)/16$. □

Theorem 6.18. *The space of refuting a Tseitin formula over the $n \times n$ grid graph for an odd charge function f over characteristic 2 is $Sp_{\mathcal{PCR}}(Ts(T(n), f) \vdash \bot) = \Theta(n)$.*

6.2 Application: Triangulations

Given a graph with good expansion, we can add a few edges to it and obtain a new graph whose Tseitin formula we can prove to be hard for PCR space. We already showed in Section 4 how to use a XOR substitution to obtain such a multi-graph; the following subsection shows how to obtain a simple graph. The proposed method is to convert every edge into a triangle, and a greedy strategy is enough as the following lemma shows.

Lemma 6.19. *Let G be a graph of order n, size m and maximal degree d. If T is an integer such that $T(n-4d-4(T+1)) \geq m$ then there exists a simple graph H of maximal degree at most $2d+2T$ which is a supergraph of G whose edges can be partitioned into disjoint triangles.*

Proof. Consider the algorithm that iteratively chooses any edge (x, y) not yet handled, chooses a vertex z not adjacent to any of the endpoints of minimal degree, and adds the two remaining edges (x, z) and (y, z) from the endpoints to the vertex.

We consider the new edges to be directed (from x and y to z) and the *indegree* and *outdegree* to refer to new edges only. The degree of a vertex is thus the sum of its initial degree, its indegree and its outdegree. Observe that at every step the outdegree of every vertex is at most its initial degree, which is at most d. When choosing the vertex z, we will choose the vertex of minimal *in*degree.

Assume that at some state S of the execution of the algorithm the maximal indegree is $2t$. We claim that the algorithm handles at least the next $n - 4d - 4(t + 1)$ edges without the indegree exceeding $2(t + 1)$.

Indeed, consider the k-th edge (x, y) the algorithm visits after state S. Its endpoints are connected to at most $d + 2(t + 1) + d$ vertices each, which we discard as candidates for z, and at most $k - 1$ vertexes increased their indegree to $2(t + 1)$. There remain at least $n - 4d - 4(t + 1) - k + 1 \geq 1$ potential vertexes of indegree at most $2t$, and the greedy algorithm chooses one of these.

The initial indegree of all vertexes is 0. After handling all m edges, the maximal indegree increases at most T times, where T is such that

$$m \leq \sum_{t=0}^{T-1} n - 4d - 4(t + 1) = T(n - 4d - 4(T + 1)) \ . \tag{6.1}$$

□

In particular, if $d \leq n/4 - \sqrt{m} - 1$ such a T exists, and if $d = o(n)$ the inequality (6.1) holds asymptotically for $T = \lceil \frac{d+1}{2} \rceil$. The lower bound on space follows by applying theorem Theorem 2.6 to the resulting supergraph and noting that the connectivity expansion cannot decrease.

Theorem 6.20. *Let G be a graph of maximal degree $d = o(n)$ and connectivity expansion c. There exists a simple graph H of maximal degree at most $3d + 2$ which is a supergraph of G such that the space of refuting a Tseitin formula over H is at least $Sp_{\mathcal{PCR}}(Ts(H, f) \vdash \bot) \geq c/12 - (3d + 2)/8$.*

7 Cycle Partitions of Random Regular Graphs

7.1 Models of Random Regular Graphs

Let P_n be a sequence of probability spaces. A sequence of events E_n on P_n holds *asymptotically almost surely* if $\Pr[E_n] \longrightarrow 1$. In the sequel, we often abuse notation and say that an event is true asymptotically almost surely in a probability space, when we actually mean sequences of both. The probability space will depend on a parameter n.

Two probability spaces are *contiguous* if every event which holds asymptotically almost surely in one also holds asymptotically almost surely in the other; we will use the notation $A \approx B$ to denote that A and B are contiguous. Let $\mathcal{D}_d$ be the probability space of random d-regular graphs on n vertices, $\mathcal{H} + \mathcal{H}$ be the probability space of unions of (not necessarily disjoint) random Hamilton cycles on n vertices, and $\mathcal{H} \oplus \mathcal{H}$ be the probability space of unions of disjoint random Hamilton cycles on n vertices; $\mathcal{H} \oplus \mathcal{H}$ is obtained by conditioning $\mathcal{H} + \mathcal{H}$ upon the event that the two random Hamilton cycles are disjoint. Note that $\mathcal{H} + \mathcal{H}$ is a probability space on multi-graphs. Kim and Wormald [KW01] proved the following theorem (see also Wormald's survey [Wor99] and [JŁR00, §9.3–9.6]).

Theorem 7.1. *We have $\mathcal{D}_4 \approx \mathcal{H} \oplus \mathcal{H}$.*

We will need one more fact from [KW01], whose proof we only sketch.

Lemma 7.2. *If $G \sim \mathcal{H} + \mathcal{H}$ then $\Pr[G \text{ is simple}] \longrightarrow e^{-2}$.*

sketch. Fix the first Hamilton cycle H_1. Let e_i be the (random) ith edge of the second Hamilton cycle H_2. It is easy to see that $\Pr[e_i \in H_1] = 2/(n-1)$, hence $\mathrm{E}[|H_1 \cap H_2|] \longrightarrow 2$. Moreover, one can show using Brun's sieve (for example [AS00, Theorem 8.3.1]) that the distribution of $|H_1 \cap H_2|$ is asymptotically Poisson; the required calculations are sketched in [KW01, §2(iii)]. Hence $\Pr[|H_1 \cap H_2| = 0] \longrightarrow e^{-2}$. □

Putting both facts together, we get the following result which will serve as our vantage point over random 4-regular graphs.

Lemma 7.3. *Suppose E is an event which holds asymptotically almost surely in $\mathcal{H} + \mathcal{H}$. Then E also holds asymptotically almost surely for random 4-regular graphs.*

Proof. Lemma 7.2 shows that E holds asymptotically almost surely in $\mathcal{H} \oplus \mathcal{H}$, and so in $\mathcal{D}_4$ by Theorem 7.1. □

Corollary 7.4. *A random 4-regular graph is connected asymptotically almost surely.*

7.2 Some Properties of Random Regular Graphs

For a graph $G = (V, E)$ and a subset U of the vertices, recall that $N(U)$ is the set of edges connecting U and $V \setminus U$. We say that the graph G is a *δ-expander* if for every set U of at most $|V|/2$ vertices, $|N(U)| \geq \delta|U|$. Note that our definition involves edge expansion. Bollobás [Bol88] proved the following fundamental result.

Theorem 7.5. *There is a constant c_1 such that asymptotically almost surely, a random 4-regular graph is a c_1-expander.*

In fact, we can choose any $c_1 < 2(1-\eta) \approx 0.4401$, where η is the unique positive solution of $(1-\eta)^{1-\eta}(1+\eta)^{1+\eta} = 2$. In particular, asymptotically almost surely a random 4-regular graph is a 0.44-expander.

The following lemma gives a lower bound on the connectivity expansion of a random 4-regular graph, defined in Definition 2.3.

Lemma 7.6. *There is a constant c_2 such that asymptotically almost surely, the connectivity expansion of a random 4-regular graph on n vertices is at least $c_2 n$.*

Proof. Let G be a random 4-regular graph. Theorem 7.5 shows that asymptotically almost surely, G is a c_1-expander. Suppose G has connectivity expansion s. There is a set W of s edges and an edge e such that $G \setminus W$ has a component of size larger than $n/2$, but $G \setminus (W \cup \{e\})$ has no component of size larger than $n/2$. Since e breaks the giant component into two components, $G \setminus (W \cup \{e\})$ must have a component U of size larger than $n/4$. Expansion shows that $|N(U)| \geq c_1|U| > (c_1/4)n$, and so $s = |W| \geq (c_1/4)n$. This shows that we can choose $c_2 = c_1/4$. □

7.3 Simple Lower Bound

In this section we prove that refuting a non-splitting Tseitin formula on a random 4-regular graph on n vertices requires space $\Omega(\sqrt{n/\log n})$, asymptotically almost surely over the choice of the graph.

The idea is to prove that asymptotically almost surely, a random 4-regular graph on n vertices can be partitioned into cycles of length $O(\sqrt{n \log n})$. In order to prove that, it will be useful to consider a model related to $\mathcal{H} + \mathcal{H}$.

Let $[n] = \{1, \ldots, n\}$, and let S_n be the set of all permutations on $[n]$. Every permutation $\pi \in S_n$ determines a Hamilton cycle

$$H(\pi) = (\pi(1), \pi(2)), (\pi(2), \pi(3)), \ldots, (\pi(n-1), \pi(n)), (\pi(n), \pi(1)) \ .$$

(The cycle is undirected.) Let ι denote the identity permutation. We will consider the probability space $\mathcal{H}(\iota) + \mathcal{H}(\pi)$ formed by taking the union of $H(\iota)$ and $H(\pi)$, where π is chosen uniformly at random from S_n.

The idea of the proof is to divide $[n]$ into $\sqrt{n/\log n}$ blocks of length $\sqrt{n \log n}$. We will show that asymptotically almost surely, each block I_k contains a point t_k such that $s_k = \pi(t_k) \in I_k$. For any two adjacent blocks I_k, I_{k+1}, we can form a cycle of length $O(\sqrt{n \log n})$ by pasting together the path from s_k to s_{k+1} in $H(\iota)$ and the path from $\pi(t_k)$ to $\pi(t_{k+1})$ in $H(\pi)$. As a result, the graph decomposes into $\sqrt{n/\log n}$ cycles of length $O(\sqrt{n \log n})$.

Let m be a parameter depending on n; in this section, we choose $m = \sqrt{n \log n}$, while in the next section, we choose $m = C\sqrt{n}$. For simplicity, we assume that m and n/m are both integers. We partition $[n]$ into n/m blocks $I_1, \ldots, I_{n/m}$ of size m: $I_k = \{(k-1)m + 1, \ldots, (k-1)m + m\}$. Let B_k be the event that $\pi(I_k) \cap I_k = \emptyset$. We think of B_k as a bad event, and our goal in this section is to show that asymptotically almost surely, none of the B_k happen. In order to show this, we estimate the probability that B_k happens.

Lemma 7.7. *For $k \in [n/m]$, $\Pr[B_k] \leq e^{-m^2/n}$.*

Proof. Using $1 - x \leq e^{-x}$, we calculate

$$\Pr[B_k] = \prod_{i=0}^{m-1} \left(1 - \frac{m}{n-i}\right) \leq \left(1 - \frac{m}{n}\right)^m \leq e^{-m^2/n} \ . \qquad \square$$

If $\overline{B_k}$ holds, we define t_k to be the first point in I_k such that $\pi(t_k) \in I_k$, and let $s_k = \pi(t_k)$.

Lemma 7.8. *Suppose $\overline{B_k}$ and $\overline{B_{k+1}}$ both hold (indices taken modulo n/m). Define a cycle C_k by taking two paths P_k^ι, P_k^π from $s_k = \pi(t_k)$ to $s_{k+1} = \pi(t_{k+1})$, one from each of the two Hamilton cycles:*

$$P_k^\iota = (s_k, s_k + 1), (s_k + 1, s_k + 2), \ldots, (s_{k+1} - 1, s_{k+1}) \ ,$$
$$P_k^\pi = (\pi(t_k), \pi(t_k + 1)), (\pi(t_k + 1), \pi(t_k + 2)), \ldots, (\pi(t_{k+1} - 1), \pi(t_{k+1})) \ .$$

The length of C_k is at most $4m$.

Proof. Assume for simplicity that $k \neq n/m$. Then $s_k, t_k \geq (k-1)m + 1$ and $s_{k+1}, t_{k+1} \leq km + m$. The length of C_k is $(s_{k+1} - s_k) + (t_{k+1} - t_k) \leq 4m - 2$. $\square$

If none of the bad events happen, then the cycles $C_1, \ldots, C_{n/m}$ cover all of the graph. Choosing m accordingly, we can ensure that this happens asymptotically almost surely.

Lemma 7.9. *Let $m = \sqrt{n \log n}$. Asymptotically almost surely, a graph chosen according to $\mathcal{H}(\iota)+\mathcal{H}(\pi)$ decomposes into n/m cycles of size at most $4m$.*

Proof. According to Lemma 7.7, for each $k \in [n/m]$, $\Pr[B_k] \leq e^{-\log n} = 1/n$. A union bound shows that asymptotically almost surely, none of the B_k happen. Lemma 7.8 shows that the graph decomposes into n/m cycles of size at most $4m$. $\square$

The lemma easily implies the lower bound.

Theorem 7.10. *Asymptotically almost surely, the space required to refute in PCR any Tseitin formula on a random 4-regular graph on n vertices is $\Omega(\sqrt{n/\log n})$.*

Proof. For reasons of symmetry, Lemma 7.9 implies that asymptotically almost surely, a graph chosen according to $\mathcal{H} + \mathcal{H}$ decomposes into cycles of size at most $4\sqrt{n \log n}$. Lemma 7.6 shows that asymptotically almost surely, the connectivity expansion of the graph is at least $\Omega(n)$. Corollary 7.4 shows that asymptotically almost surely, the graph is connected, and so the Tseitin formula is non-splitting. Hence Theorem 3.4 gives a lower bound of $\Omega(\sqrt{n/\log n})$. $\square$

7.4 Improved Lower Bound

In this section we improve the results of Section 7.3 by showing that refuting a non-splitting Tseitin formula on a random 4-regular graph on n vertices requires space $\Omega(\sqrt{n})$, asymptotically almost surely over the choice of the graph.

We use the general method of Section 7.3, with a different choice of m, namely $m = C\sqrt{n}$ for some constant C to be determined later. Thinking of B_k as an indicator variable, let $B = \sum_{k=1}^{n/m} B_k$. Lemma 7.7 shows that $\mathrm{E}[B] \leq e^{-C^2}(n/m)$. We will show that asymptotically almost surely, $B \leq 2e^{-C^2}(n/m)$. This implies that the cycles C_k together cover most of the graph, and therefore Corollary 6.11 applies. The difficult part of the proof is showing that B is concentrated around its mean.

Let $p = \Pr[B_k]$ (all the probabilities are the same). We need the following strengthening of Lemma 7.7.

Lemma 7.11. *Let $p = \Pr[B_k]$, where B_k is the event that $I_k \cap \pi(I_k) = \emptyset$. As $n \longrightarrow \infty$, we have that $p \longrightarrow e^{-C^2}$.*

In order to show that B is concentrated around its mean, we show that for $k \neq l$, the events B_k and B_l are asymptotically negatively correlated.

Lemma 7.12. *For every $k \neq l \in [n/m]$, $\Pr[B_k \wedge B_l] \leq p^2 + o(1)$.*

We prove both lemmas below, but first, let us see how they imply the desired result. The idea is that since any two bad events are asymptotically negatively correlated, the variance of B is small, and so Chebyshev's inequality shows that B is concentrated around its mean.

Lemma 7.13. *Asymptotically almost surely, $B \leq 2e^{-C^2}(n/m)$.*

Proof. We have $\mathrm{E}[B] = (n/m)p$ and

$$\begin{aligned}\mathrm{Var}(B) &= \mathrm{E}[B^2] - (\mathrm{E}[B])^2 \\ &= (n/m)p + (n/m)(n/m-1)(p^2+o(1)) - (n/m)^2p^2 \\ &= (n/m)p(1-p) + o((n/m)^2) \ ,\end{aligned}$$

using Lemma 7.12. Chebyshev's inequality shows that

$$\Pr[|B - \mathrm{E}[B]| > \mathrm{E}[B]] \leq \frac{\mathrm{Var}(B)}{\mathrm{E}[B]^2} \leq \frac{(n/m)p + o((n/m)^2)}{(n/m)^2p^2} = o(1) \ ,$$

since $p = \Omega(1)$ by Lemma 7.11. Therefore asymptotically almost surely, $B \leq 2\,\mathrm{E}[B] = 2(n/m)p \leq 2e^{-C^2}(n/m)$, using Lemma 7.7. □

The preceding lemma shows that the fraction of bad indices (indices k such that B_k holds) is small. Say that a block I_k is *good* if $\overline{B_k}$ and $\overline{B_{k+1}}$ both hold, and say that it is *supergood* if both I_{k-1} and I_k are good. Lemma 7.8 associates a cycle C_k with each good block I_k. If I_k is supergood, then the cycles C_{k-1} and C_k together cover the entire stretch of I_k, as the following lemma shows.

Lemma 7.14. *Suppose that block I_k is supergood. Then the union of the cycles C_{k-1}, C_k given by Lemma 7.8 contains the path of length m from $\min I_k$ to $\min I_{k+1}$ in $H(\iota)$, as well as the path of length m from $\pi(\min I_k)$ to $\pi(\min I_{k+1})$ in $H(\pi)$.*

Proof. The cycle C_{k-1} contains the path from $s_{k-1} < \min I_k$ to s_k in $H(\iota)$. The cycle C_k contains the path from s_k to $s_{k+1} \geq \min I_{k+1}$ in $H(\iota)$. Both paths together cover the path from $\min I_k$ to $\min I_{k+1}$ in $H(\iota)$. The argument for $H(\pi)$ is identical. □

We can now prove an analogue of Lemma 7.9.

Lemma 7.15. *Let $m = C\sqrt{n}$. Asymptotically almost surely, a graph chosen according to $\mathcal{H}(\iota) + \mathcal{H}(\pi)$ decomposes into cycles of size at most $4m$ and t additional edges, where $t \leq 12e^{-C^2}n$.*

Proof. Lemma 7.13 shows that asymptotically almost surely, all but $6e^{-C^2}(n/m)$ of the n/m blocks $I_1, \ldots, I_{n/m}$ are supergood. Let $\mathcal{C}$ be the (disjoint) union of all cycles C_k constructed using Lemma 7.8 for all good blocks I_k. The lemma shows that each cycle has size at most $4m$. Lemma 7.14 shows that $\mathcal{C}$ contains all but at most $12e^{-C^2}n$ edges of the graph. □

Replacing Theorem 3.4 with its corollary, Lemma 7.15 easily implies the lower bound.

Theorem 7.16. *Asymptotically almost surely, the space required to refute in PCR any Tseitin formula on a random 4-regular graph on n vertices is $\Omega(\sqrt{n})$.*

Proof. For reasons of symmetry, Lemma 7.15 implies that asymptotically almost surely, a graph chosen according to $\mathcal{H} + \mathcal{H}$ decomposes into cycles of size at most $4C\sqrt{n}$ and t additional edges, where $t \leq 12e^{-C^2}n$. For an appropriate choice of C, $t \leq (c_2/2)n$. Lemma 7.6 shows that asymptotically almost surely, the connectivity expansion of the graph is at least $c_2 n$. Corollary 7.4 shows that asymptotically almost surely, the graph is connected, and so the Tseitin formula is non-splitting. Hence Corollary 6.11 gives a lower bound of $\Omega(\sqrt{n})$. □

7.4.1 Technical Lemmas

We now turn to the proofs of Lemma 7.11 and Lemma 7.12. We start with the former.

of Lemma 7.11. It is easy to check that for $0 \leq x \leq 1/2$, $1 - x \geq e^{-x-x^2}$. Therefore for large enough n,

$$p = \prod_{i=0}^{m-1}\left(1 - \frac{m}{n-i}\right) \geq \left(1 - \frac{m}{n-m}\right)^m \geq \exp\left[-\frac{m^2}{n-m} - \frac{m^3}{(n-m)^2}\right] .$$

For large enough n, $m \leq n/2$, and so $m^2/(n-m) = m^2/n + m^3/(n(n-m)) \leq m^2/n + 2m^3/n^2$. Similarly, $m^3/(n-m)^2 \leq 4m^3/n^2$. Therefore, using $e^{-x} \geq 1 - x$,

$$p \geq \exp\left[-\frac{m^2}{n} - 6\frac{m^3}{n^2}\right] = \exp\left[-C^2 - \frac{6C^3}{\sqrt{n}}\right] \geq e^{-C^2}\left(1 - \frac{6C^3}{\sqrt{n}}\right) .$$

Hence $\liminf p \geq e^{-C^2}$. Lemma 7.7 shows that also $\limsup p \leq e^{-C^2}$. □

The proof of Lemma 7.12 is more involved. Recall that the lemma claims that the events B_k and B_l are asymptotically negatively correlated. In fact, they are asymptotically uncorrelated. Recall that $\Pr[B_k]$ is roughly equal to e^{-C^2}. Given the value of π on I_k, the probability $\Pr[B_l]$ depends on $|\pi(I_k) \cap I_l|$. Typically, this intersection will be very small, and so $\Pr[B_l]$ is also roughly equal to e^{-C^2}.

We will show that $|\pi(I_k) \cap I_l|$ is typically small using an extension of the well-known Chernoff bound due to Kabanets and Impagliazzo [IK10, Theorem 1.1], attributed there to Panconesi and Srinivasan [PS97].

Theorem 7.17. *Let $X_1, \ldots, X_r$ be Boolean random variables such that for any set $S \subseteq [r]$, $\Pr[\bigwedge_{i \in S} X_i] \leq \delta^{|S|}$. Then for $\gamma \geq \delta$,*

$$\Pr\left[\sum_{i=1}^{r} X_i \geq \gamma r\right] \leq e^{-2r(\gamma-\delta)^2} .$$

The following lemma applies this bound to our situation (in an abstracted version).

Lemma 7.18. *Let a, b, c be integers such that $a \geq b, c$, and let T be a random subset of $[a]$ of size b. For all $\rho \geq 1$,*

$$\Pr[|T \cap [c]| \geq \rho(bc/a)] \leq e^{-2c(\rho-1)^2(b/a)^2} .$$

Proof. For $i \in [c]$, let X_i be the event that $i \in T$. For $S \subseteq [c]$ such that $|S| \leq b$,

$$\Pr_T[S \subseteq T] = \frac{\binom{a-|S|}{b-|S|}}{\binom{a}{b}} = \prod_{k=0}^{|S|-1} \frac{b-k}{a-k} \leq \left(\frac{b}{a}\right)^{|S|} .$$

Therefore we can apply Theorem 7.17 with $r = c$, $\delta = b/a$ and $\gamma = \rho(b/a)$. □

We can now prove Lemma 7.12.

of Lemma 7.12. We will show that $\Pr[B_l \mid B_k] \leq p + o(1)$. This implies that $\Pr[B_k \wedge B_l] = \Pr[B_k]\Pr[B_l \mid B_k] \leq p(p + o(1)) = p^2 + o(1)$.

Assuming the event B_k happens, $\pi(I_k)$ is a random subset of $[n] \setminus I_k$ of size m. Plugging $a = n - m$ and $b = c = m$ in Lemma 7.18, we deduce that for all $\rho \geq 1$,

$$\begin{aligned}\Pr[|\pi(I_k) \cap I_l| \geq \rho C^2 \mid B_k] &\leq e^{-2(\rho-1)^2 m(m/(n-m))^2} \\ &\leq e^{-2(\rho-1)^2 m^3/n^2} = e^{-2C^3(\rho-1)^2/\sqrt{n}} .\end{aligned}$$

Hence with probability $1 - o(1)$ given B_k, $D \triangleq |\pi(I_k) \cap I_l| \leq \sqrt{m \log m}$. Now

$$\Pr[B_l \mid D = d] = \prod_{i=0}^{m-1} \left(1 - \frac{m-d}{n-i}\right) \leq \left(1 - \frac{m-d}{n}\right)^m \leq e^{-m(m-d)/n} .$$

For $0 \leq x \leq 1$, one can check that $e^x \leq 1 + 2x$. Hence

$$\begin{aligned}\Pr[B_l \mid D \leq \sqrt{m \log m}] &\leq e^{-m(m-\sqrt{m \log m})/n} \\ &= e^{-C^2 + m\sqrt{m \log m}/n} \leq e^{-C^2}\left(1 + \frac{2m\sqrt{m \log m}}{n}\right) .\end{aligned}$$

Using Lemma 7.11, we deduce that $\Pr[B_l \mid D \leq \sqrt{m \log m}] \leq e^{-C^2} + o(1) = p + o(1)$. We conclude that $\Pr[B_l \mid B_k] = p + o(1)$ and so $\Pr[B_k \wedge B_l] = p^2 + o(1)$. □

7.5 Regular Graphs of Degree Larger Than Four

Wormald [Wor99, Corollary 4.17] showed that when $d > 4$, a random d-regular graph can be obtained (up to contiguity) by taking the disjoint union of a random 4-regular graph and a random $(d-4)$-regular graph, a result summarized in the following theorem (see also [JŁR00, Corollary 9.44]).

Theorem 7.19. *For $d > 4$ we have $\mathcal{D}_d \approx \mathcal{D}_4 \oplus \mathcal{D}_{d-4}$. Furthermore, the probability that a uniformly random 4-regular graph and a uniformly random $(d-4)$-regular graph do not intersect tends to a positive constant.*

A Tseitin formula on a random d-regular graph generated according to $\mathcal{D}_4 \oplus \mathcal{D}_{d-4}$ is harder to refute than a Tseitin formula on a random 4-regular graph, and so we can generalize Theorem 7.16 to random d-regular graphs for arbitrary $d \geq 4$.

Theorem 7.20 (restatement of Theorem 3.5). *Let $d \geq 4$. Asymptotically almost surely, the space required to refute in PCR any Tseitin formula on a random d-regular graph on n vertices is $\Omega(\sqrt{n})$.*

Proof. If $d = 4$ then Theorem 7.16 already applies, so assume $d > 4$. Let G_1 be a random 4-regular graph, and let G_2 be a random $(d-4)$-regular graph. The graph $G = G_1 + G_2$ is distributed according to $\mathcal{D}_4 + \mathcal{D}_{d-4}$. We show below that asymptotically almost surely, the space required to refute in PCR any Tseitin formula on G is $\Omega(\sqrt{n})$. Since G_1 and G_2 are disjoint with constant probability according to Theorem 7.19, the theorem follows.

Let α be an arbitrary assignment to the edges of G_2. Observation 6.2 on page 14 shows that for every function f, $Ts(G, f)\restriction_\alpha = Ts(G_1, g)$ for some other function g. By a restriction argument, any PCR refutation of $Ts(G, f)$ in space S can be translated to a PCR refutation of $Ts(G_1, g)$ in space at most S. Theorem 7.16 on the preceding page shows that asymptotically almost surely, we must have $S = \Omega(\sqrt{n})$. □

8 Current Techniques and the Functional Pigeonhole Principle

We now discuss the intrinsic limitations of the techniques employed so far. In Section 8.1 we show that Bonacina-Galesi framework does not allow to prove PCR space lower bounds for an interesting formula like functional pigeonhole principle. In Section 8.2 we show that restricting to PC does not make the problem easier.

8.1 FPHP Formulas Do Not Have Extendible Families

One of the limits of the Bonacina-Galesi framework is that we cannot apply it to formulas for which fixing a small set of variables causes a lot of unit clauses propagation. Indeed, most of the lower bound strategies in this paper aim to control this phenomenon (see for example Lemma 4.3). For the functional pigeonhole principle these strategies do not work, as we now prove.

Definition 8.1. The *functional pigeonhole principle* on m pigeons and n holes is the formula defined on variables x_{ij} for $i \in [m]$ and $j \in [n]$, made of the following clauses:

$$\bigvee_{j\in[n]} x_{ij} \qquad \text{for all } i \in [m]; \qquad \text{(pigeon axioms)}$$

$$\neg x_{ij} \vee \neg x_{i'j} \qquad \text{for any } i \neq i' \in [m] \text{ and } j \in [n]; \qquad \text{(hole axioms)}$$

$$\neg x_{ij} \vee \neg x_{ij'} \qquad \text{for any } i \in [m] \text{ and } j \neq j' \in [n]. \qquad \text{(functional axioms)}$$

It is already known that this formula requires large space in resolution [BW01, AD08]. It is natural to suspect that this formula is hard in terms of monomial space as well. However, the Bonacina-Galesi framework is not strong enough to prove it.

Theorem 8.2 (restatement of Theorem 3.6). *There is no r-extendible family for $FPHP_n^m$ for $r > 1$.*

Proof. Assume that there is an r-extendible family $\mathcal{F}$ for the formula $FPHP_n^m$ which respects some satisfiable $F' \subseteq FPHP_n^m$, for $r > 1$.

Let C be any clause in $FPHP_n^m \setminus F'$; such clause exists because $FPHP_n^m$ is a contradiction. The extension property of $\mathcal{F}$ implies that there is a pair $(\{Q_1\}, H_1) \in \mathcal{F}$, where H_1 satisfies C.

Recall that 0 encodes true, and 1 encodes false. Pick a variable x_{ij} in Q_1. In H_1 there is at least one partial assignment for which $x_{ij} = 0$, and for any such assignment it holds that $x_{i'j} = 1$ and $x_{ij'} = 1$ for all $i' \neq i$ and $j' \neq j$, otherwise an initial clause would be false.

Indeed, fix v to be any of these variables (either $x_{i'j}$ or $x_{ij'}$); the clause $\neg x_{ij} \vee \neg v$ is an axiom. If $v \notin Q_1$ then this clause is not in F' because of the respectfulness of $\mathcal{F}$, and furthermore there is at least one assignment in H_1 which does not satisfy it (i.e., any assignment with $x_{ij} = 0$). The extension property of $\mathcal{F}$ guarantees that there is $(\{Q_1, Q_2\}, H_1 \times H_2) \in \mathcal{F}$ with $v \in Q_2$, such that $H_1 \times H_2$ satisfies $\neg x_{ij} \vee \neg v$. But this contradicts the fact that $H_1 \times H_2$ contains the assignment $\{x_{ij} = 1, v = 1\}$, which falsifies $\neg x_{ij} \vee \neg v$.

It follows that $\{x_{i'j}, x_{ij'} \mid i' \neq i \text{ and } j' \neq j\} \subseteq Q_1$, and that H_1 satisfies all axioms involving either pigeon i or hole j. We have just shown that assuming some $x_{ij} \in Q_1$, we get $\{x_{i'j}, x_{ij'} \mid i' \in [m], j' \in [n]\} \subseteq Q_1$. This choice was arbitrary, so it follows that for any $i \in [m], j \in [n]$, the variable x_{ij} is in Q_1. In other words, Q_1 contains all the variables. Since $FPHP_n^m \setminus F'$ is contradictory, every assignment in H_1 falsifies some clause, and so the extension property fails for any such clause. We conclude that $FPHP_n^m$ has no 2-extendible family. □

8.2 Formulas with Equal PC and PCR Space Complexities

Although finding an r-extendible family for the functional pigeonhole principle (and hence proving a space lower bound) is not feasible, we might try and prove a weaker PC space lower bound. However, as we have pointed out in Section 3.4, in the case of functional pigeonhole principle this makes no difference. In this section, we prove formally this result for a broader class of formulas that is captured by the following definition.

Definition 8.3. We say that a CNF formula F is *totally weight constrained* if for every variable x appearing in F there exists a clause $C_x \in F$ with the following properties:

1. All literals in C_x are positive;
2. x is one of the variables appearing in C_x;
3. For every two distinct variables y, z appearing in C_x, clause $\overline{y} \vee \overline{z}$ is in F.

For each variable x we refer to C_x as the *x-neighborhood clause*.

In such formulas each negative literal can be replaced with a clause/monomial consisting of only positive literals that has the same semantic meaning. Thus, we can turn a PCR refutation into a PC refutation without any substantial loss of space. In order for us to be able to show that such a refutation is a valid PC refutation we need to show that there are PC derivations of these monomials that use small space.

Theorem 8.4. *For a totally weight constrained CNF formula F, where each clause has a costant number of negative literals, it holds that $Sp_{PC}(F \vdash \bot) = \Theta(Sp_{PCR}(F \vdash \bot))$.*

Proof. We can easily see that PCR simulates PC with only a constant loss in space. The only problem in the simulation could arise when downloading an axiom that has negative literals. Nevertheless, it is not hard to prove that PCR can expand every axiom to its PC form while respecting the stated space bound.

In the other direction, we prove that PC can simulate a PCR refutation of F. Let π be a PCR refutation of F in space at most s. As F is a totally weight constrained formula, for every variable x we can fix its x-neighborhood clause C_x. Let us denote by $N(x)$ the set of variables from C_x excluding x. We transform the PCR refutation π into a PC refutation by replacing each negative literal $\overline{x}$ with the monomial $\prod_{y \in N(x)} y$. Obviously this transformation preserves space and we need to show that the transformed configurations form a backbone of a valid PC refutation.

If the PCR refutation deletes a polynomial, we delete the appropriate transformed polynomial from the configuration in the PC refutation. Similarly, in the case of linear combination steps we just deduce the linear combination of the transformed polynomials. Hence, these two types of steps can be done without any loss in space. In the case of multiplication with a literal, if the literal is positive we multiply the appropriate transformed polynomial with the same literal. Otherwise, the literal is negative and we multiply the polynomial with all the variables in $N(x)$, where $\overline{x}$ is the literal, while making sure to delete the intermediate polynomials when they are no longer needed. In this way we derive the transformed polynomial in at most $\mathrm{O}(s)$ space.

The axiom download steps are the only ones that remain. In the case of Boolean axiom download, if we downloaded an axiom for a positive literal, we just download the appropriate axiom in the PC refutation. Otherwise, the Boolean axiom corresponds to some negative literal $\overline{x}$ and we need to derive the polynomial $\prod_{y \in N(x)} y^2 - \prod_{y \in N(x)} y$. This is done by downloading the Boolean axioms for each $y \in N(x)$ and combining them to get the transformed polynomial. Let $B^2 - B$ be one of the intermediate polynomials in the derivation of the transformed Boolean axiom, where B is a monomial formed by multiplying the variables in some subset of $N(x)$. Then, for some variable y not mentioned in B, we derive $(By)^2 - By$ by downloading $y^2 - y$ and taking the linear combination of $y(B^2 - B)$ and $B^2(y^2 - y)$. This PC derivation uses $\mathrm{O}(1)$ more monomials than the PCR axiom download.

When the PCR proof downloads the complementarity axiom $1 - x - \overline{x}$, the corresponding PC proof needs to derive the polynomial $1 - x - \prod_{y \in N(x)} y$. Let $N(x) = \{y_1, \ldots, y_l\}$. We derive the transformed polynomial by successively deriving polynomials

$$T(i) = \prod_{k=i+1}^{l} y_k - x \prod_{k=i+1}^{l} y_k - \prod_{k} y_k \ , \tag{8.1}$$

for $i = 1, \ldots, l$. Note that $T(l)$ is our transformed polynomial. The first $T(1)$ in the PC proof can be derived by downloading the axiom $(1 - x)(1 - y_1)$ and multiplying it with variables $y_2, \ldots, y_l$ in order to get $T(1) + x \prod_k y_k$. Subtracting from it the x-neighborhood clause $C_x = x \prod_k y_k$ we get $T(1)$.

We proceed to derive $T(i+1)$ from $T(i)$ for all i. Similarly as before, we start by downloading the axiom $(1-x)(1-y_{i+1})$ and multiplying it with variables $y_{i+2},\ldots,y_l$ in order to get $T(i+1)-T(i)$. Adding this polynomial to $T(i)$ we derive the $(i+1)$st polynomial $T(i+1)$ in our derivation of the transformed complementarity axiom. This PC derivation uses $\mathrm{O}(1)$ more monomials than the PCR proof and all axioms of the form $(1-x)(1-y_i)$ exist because F is totally weight constrained.

In the case of axiom download step for a clause axiom, we again have two cases. If all literals of the axiom are positive we download the corresponding axiom in the PC proof. Otherwise, we can write the axiom as $\overline{x_1}\cdots\overline{x_s}\cdot x_{s+1}\cdots x_l$, where s is the number of its negative literals. Let us denote by $A(i)$ the polynomial

$$A(i)=\prod_{y_1\in N(x_1)} y_1\cdots \prod_{y_i\in N(x_i)} y_i(1-x_{i+1})\cdots(1-x_s)x_{s+1}\cdots x_l \ , \tag{8.2}$$

where i ranges over $0,\ldots,s$. Note that $A(0)$ is the original PC axiom, while $A(s)$ is the transformed axiom that we want to derive. Also, let us denote by $R(i)$ the polynomial

$$R(i)=\prod_{y_1\in N(x_1)} y_1\cdots \prod_{y_{i-1}\in N(x_{i-1})} y_{i-1}\cdot(1-x_{i+1})\cdots(1-x_s)x_{s+1}\cdots x_l \ , \tag{8.3}$$

for i ranging from 1 to s, that is $A(i)=R(i)\prod_{y_i\in N(x_i)} y_i = R(i+1)(1-x_{i+1})$.

We first derive $A(1)$ by deriving the transformed complementarity axiom $1-x_1-\prod_{y_1\in N(x_1)} y_1$ for the variable x_1 and multiplying it with $R(1)$ in order to get $A(0)-A(1)$. Now we can get $A(1)$ by subtracting the derived polynomial from the PC axiom $A(0)$.

We proceed to derive $A(s)$ by deriving $A(i+1)$ from $A(i)$ for all i from 1 to $s-1$. This is again done by first deriving the appropriate complementarity axiom $1-x_{i+1}-\prod_{y_{i+1}\in N(x_{i+1})} y_{i+1}$ and multiplying it by $R(i+1)$ in order to get $A(i)-A(i+1)$. Subtracting the derived polynomial from previously derived $A(i)$, we get the $(i+1)$st polynomial in our derivation. These steps use $\mathrm{O}(2^s)$ monomials, which is constant by the theorem hypothesis, and the PC derivation of the transformed axiom uses at most $\mathrm{O}(1)$ monomials more than the PCR axiom download step.

Hence, the theorem follows. Also, although we have ignored the constants involved in the simulation, these constants can be computed explicitly and are small. The only possible exception is the additive constant $\mathrm{O}(2^{s^*})$, where s^* is the largest number of negative literals in a clause of F. □

An obvious example of the totally weight constrained formula is the functional pigeonhole principle.

Corollary 8.5 (Restatement of Theorem 3.7). *It holds that*

$$Sp_{PCR}(FPHP_n^m \vdash \bot) = \Theta(Sp_{PC}(FPHP_n^m \vdash \bot)) \ .$$

Proof. It is easy to see that $FPHP_n^m$ formula is totally weight constrained, as every variable appears in some pigeon axiom that is constrained by the functional axioms. Also, $FPHP_n^m$ has at most 2 negative literals in each clause and hence we have that $Sp_{PCR}(FPHP_n^m \vdash \bot) = \Theta(Sp_{PC}(FPHP_n^m \vdash \bot))$. □

Actually, we can say even more about the space complexity of the functional pigeonhole principle formulas. In [FLN+12], the authors prove that the PCR space complexity of $FPHP_n^m$ is equal (up to constant factors) to the PCR space complexity of the extended formula $\widetilde{FPHP}_n^m$, where $\widetilde{FPHP}_n^m$ is the canonical equivalent 3-CNF version[7] of the formula $FPHP_n^m$. Hence, we have that the PC space complexity lower bound for $FPHP_n^m$ would actually lower bound the PCR space complexity of $\widetilde{FPHP}_n^m$ and give us the first PCR space lower bound for some family of 3-CNF formulas.

This holds in greater generality for totally weight constrained formulas that also fulfill the following technical condition: F is a *weight-constrained* CNF formula if for each clause $l_1\vee l_2\vee\ldots\vee l_m$ of F with

[7] We substitute every clause $l_1\vee l_2\vee\ldots\vee l_k$, which has more than three literals, with the formula $(l_1\vee y_1)\wedge(\neg y_1\vee l_2\vee y_2)\wedge\ldots\wedge(\neg y_{i-1}\vee l_i\vee y_i)\wedge\ldots\wedge(\neg y_{k-1}\vee l_k)$ where for each substituted clause all variables y_i are new. The substituted formula is a 3-CNF and it is satisfiable if and only if the original one is. It is also easy to deduce the original clause from the substituting formula.

more than three literals, the formula also contains clauses $\neg l_i \vee \neg l_j$ for all $1 \leq i < j \leq m$. We stress the fact that the conditions of being weight-constrained and totally weight constrained are incomparable.

Corollary 8.6. *For a simultaneously weight-constrained and a totally weight constrained formula F, where each clause has a costant number of negative literals, it holds that*

$$Sp_{PCR}(\widetilde{F} \vdash \bot) = \Theta(Sp_{PCR}(F \vdash \bot)) = \Theta(Sp_{PC}(F \vdash \bot)) \ .$$

9 Concluding Remarks

In this paper, following up on recent work in [BNT13, BG13, FLN+12, HN12], 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 [BG13] 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 [BG13] 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 [AD08].

It thus remains a major open problem to understand the relation between degree and space in PC/PCR, and in particular whether degree is a lower bound on space or not (or whether it even holds that resolution width provides a lower bound on PCR space).

Also, our separations of size and degree on the one hand and space on the other depend on the characteristic of the underlying field, in that that the characteristic must be chosen first and the formula family exhibiting the separation works only for this specific characteristic. It would be satisfying to find formulas that provide such separations regardless of characteristic. Natural candidates are (various flavours of) ordering principle formulas or onto function pigeon principle formulas, or, for potentially even stronger separations, pebbling formulas.

Finally, an intriguing question is how (monomial) space in PC/PCR is related to (clause) space in resolution. There are separations known for size versus length and degree versus width, and it would seem reasonable to expect that PCR should be strictly stronger than resolution also with respect to space, but this is completely open.[8] The flipside of this question is to what extent space lower bound techniques for resolution carry over to PC/PCR. Since so far we do not know of any counter-examples, it is natural to ask, for instance, whether *semiwide* CNF formulas as defined in [ABRW02] have high space complexity not only in resolution but also in PCR.

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) /

[8] For completeness, we mention that there is a very weak (constant-factor) separation in [ABRW02], but it crucially depends on a somewhat artificial definition of space where monomials are *not* counted with repetitions.

ERC grant agreement no. 279611. The fourth author was also supported by Swedish Research Council grants 621-2010-4797 and 621-2012-5645.

A PCR Space Lower Bounds from Extendible Families

For the sake of self-containment, in this appendix we give an exposition of the Bonacina-Galesi framework [BG13] for proving space lower bounds in Polynomial Calculus. We show how the existence of a r-extendible family for a large value of r implies such bounds. This framework can actually prove space lower bounds for a proof system that it stronger than PC or PCR.

Definition A.1 (Functional Calculus (FC)). A *functional calculus* configuration is a set of arbitrary Boolean functions over Boolean variables. There is a single derivation rule, *semantic implication*, where g can be inferred from $f_1, \ldots, f_n$ if every assignment that satisfies $f_1 \wedge \cdots \wedge f_n$ also satisfies g.

Verifying a proof in FC is co-NP-complete, and so FC is not a proof system in the sense of Cook and Reckhow [CR79] unless co-NP $=$ NP.

There are many different circuit representations of the same Boolean function, so we need to choose a minimal representation in order to define clause space.

Definition A.2. Let $\mathbb{P}$ be a FC configuration. A set of monomials $U = \{m_1, \ldots, m_s\}$ *defines* $\mathbb{P}$ if for every function $f \in \mathbb{P}$ there is a function g such that $g(m_1, \ldots, m_s) \equiv f(x_1, \ldots, x_n)$. The *monomial space* of $\mathbb{P}$ is the minimum size of a defining set of monomials.

We can interpret polynomials in PCR as Boolean functions if we project them to the Boolean ring $\mathbb{F}[x, \overline{x}, y, \overline{y}, \ldots] / \operatorname{Span}\left(x^2 - x, 1 - x - \overline{x}, \overline{x}^2 - \overline{x}, y^2 - y, \ldots\right)$. Furthermore, the set of monomials in a PCR configuration counted without repetitions is a defining set of monomials for a FC configuration. Therefore we can view every proof in PCR as a proof in FC that uses at most the same space. In particular, $Sp_{\mathcal{FC}}(F \vdash \bot) \leq Sp_{\mathcal{PCR}}(F \vdash \bot)$.

We now prove Theorem 2.6, following Bonacina and Galesi [BG13]. The general plan of the proof is to consider a FC derivation of a formula F in small space, and show that every configuration arising in the derivation is satisfiable. Since a refutation ends with an unsatisfiable configuration, the derivation is not a refutation.

In order to show that every configuration arising in the derivation is satisfiable, we maintain a satisfiability witness, in the form of a structured set of assignments together with a CNF formula. The following definition captures the sense in which a satisfiability witness guarantees that a board configuration is satisfiable. Fix a set of variables V and consider partitions and total assignments with respect to this set. Recall that a total assignment assigns a value to *each* variable in V.

Definition A.3. Let $(\mathcal{Q}, \mathcal{H})$ be a structured set of assignments, G be a CNF formula, and $\mathbb{P}$ be a set of Boolean functions. We write $G \models_{(\mathcal{Q},\mathcal{H})} \mathbb{P}$ if every total assignment that extends some partial assignment in $\mathcal{H}$ and satisfies G also satisfies $\mathbb{P}$.

In the proof, $\mathbb{P}$ is the contents of the board at a given point in the FC refutation, and $(\mathcal{Q}, \mathcal{H}), G$ together form a satisfiability witness. The CNF G is composed of two parts: a satisfiable subset $F' \subset F$, which could be empty, and a 2-CNF M with a very specific form given by the following definition.

Definition A.4. Let M be a 2-CNF formula over the variables V. We say that M is a *transversal* of a partial partition $\mathcal{Q}$ defined on V if M mentions exactly one variable from each block $Q_i \in \mathcal{Q}$. (In particular, $|\mathcal{Q}|$ must be even and the number of clauses in M is $|\mathcal{Q}|/2$.)

A transversal CNF formula is always satisfiable, and so for $F' = \emptyset$, any board configuration $\mathbb{P}$ that has a satisfiability witness of this form must in fact be satisfiable. To handle an arbitrary F', we add the requirement that $(\mathcal{Q}, \mathcal{H})$ respect F'. Finally, we can formally define the concept of satisfiability witness.

Definition A.5. Let $\mathbb{P}$ be a set of Boolean functions. A tuple $(F'; \mathcal{Q}, \mathcal{H}, M)$ is a *satisfiability witness* for $\mathbb{P}$ if:

1. F' is a satisfiable CNF formula.
2. $(\mathcal{Q}, \mathcal{H})$ is a structured assignment set which respects F'.
3. M is a 2-CNF formula which is a transversal of $\mathcal{Q}$.
4. $F' \wedge M \models_{(\mathcal{Q},\mathcal{H})} \mathbb{P}$.

The *size* of a satisfiability witness $(F'; \mathcal{Q}, \mathcal{H}, M)$ is $|M|$.

We single F' out since its value is fixed while $\mathcal{Q}, \mathcal{H}, M$ are dynamic and change throughout the FC refutation.

A FC refutation is composed of three kinds of steps: axiom download, inference and erasure. It turns out that the first two steps are relatively easy to handle, as long as we maintain the invariant that the size of the satisfiability witness is $\mathrm{O}(Sp(\mathbb{P}))$. This invariant allows us to expand the witness in order to accommodate new axioms as long as the monomial space is small enough, using the extension property of extendible families.

Erasure is more difficult, since the monomial space of the configuration could shrink, and in order to maintain the invariant, we need to shrink the witness as well. This is accomplished by the following crucial lemma, which shows that if a configuration has any satisfiability witness, then we can find another satisfiability witness for the configuration whose size is bounded in terms of the monomial space of the configuration.

Because of the multiple representations technical issue we also need to use the locality lemma in axiom download steps, but we could omit it in a proof of a space lower bound for PCR. It is however a key piece in erasure steps.

Lemma A.6 (Locality lemma). *Suppose $(F'; \mathcal{Q}, \mathcal{H}, M)$ is a satisfiability witness for some set of Boolean functions $\mathbb{P}$. There is another satisfiability witness $(F'; \mathcal{Q}', \mathcal{H}', M')$ for $\mathbb{P}$ such that $\mathcal{Q}' \subseteq \mathcal{Q}$, $\mathcal{H}' = \mathcal{H}\restriction_{\mathcal{Q}'}$ and $|M'| \leq 2Sp(\mathbb{P})$.*

Proof. In this proof $\mathcal{Q}[x]$ denotes the (unique) class in $\mathcal{Q}$ that contains variable x.

The starting point of the proof is understanding the relation between monomials in a defining set of monomials U of $\mathbb{P}$ and clauses in M which underlies the property $F' \wedge M \models_{(\mathcal{Q},\mathcal{H})} \mathbb{P}$. A clause $C \in M$ affects a monomial $m \in U$ whenever the two mention variables belonging to the same partition in $\mathcal{Q}$. If a clause C does not affect a monomial m, then the clause C puts no constraints on the value of m.

Formally, we construct a bipartite graph between a minimal defining set of monomials U and the set of clauses in M (which we identify with M itself). We draw an edge between $m \in U$ and $C \in M$ whenever for some $Q \in \mathcal{Q}$, both m and C mention some variable in Q.

We break U into two parts: one part which is collectively affected by a small number of clauses, and another part in which we can associate with each monomial two clauses affecting it. To this end, let U_1 be an inclusion-maximal set under the constraint $|N(U_1)| \leq 2|U_1|$, and let $U_2 = U \setminus U_1$. We partition M accordingly into $M_1 = N(U_1)$ and $M_2 = M \setminus M_1$. As a slight modification of Hall's marriage theorem shows, the maximality of U_1 implies that we can associate with each monomial in U_2 two *unique* clauses in M_2 (that is, each clause in M_2 is associated with at most one monomial). In other words, there is a double matching from U_2 to M_2. (For more details on this step, see [ABRW02, FLN+12, BG13].)

We construct the new 2-CNF M' out of two parts: $M' = M_1 \cup M_2'$. The first part M_1, taken verbatim from M, takes care of U_1. The other part M_2', which we construct from the double matching, takes care of U_2.

The 2-CNF M_2' consists of one clause C_m for every monomial $m \in U_2$. In order to define C_m, let $x^a \vee y^b$ and $z^c \vee w^d$ be the two clauses in M_2 that are matched to m in the double matching. Assume without loss of generality that $m = r^e s^f m'$, where $r \in \mathcal{Q}[x]$ and $s \in \mathcal{Q}[z]$. The clause C_m is defined as $C_m = r^e \vee s^f$.

By construction, $|M'| \leq 2|U_1| + |U_2| \leq 2|U| = 2Sp(\mathbb{P})$. Having defined M', we complete the definition of the new satisfiability witness as follows. First, let $\mathcal{Q}' = \{\mathcal{Q}[x] \mid x \in Vars(M')\}$; this guarantees that M' is a transversal of $\mathcal{Q}'$. Observe that $\mathcal{Q}' \subseteq Q$. Second, let $\mathcal{H}' = \mathcal{H}\restriction_{\mathcal{Q}'}$. It is easy to check that $(F'; \mathcal{Q}', \mathcal{H}', M'))$ satisfies the first three properties of a satisfiability witness. It remains to prove that $F' \wedge M' \models_{(\mathcal{Q}',\mathcal{H}')} \mathbb{P}$.

In order to show that $F' \wedge M' \models_{(\mathcal{Q}',\mathcal{H}')} \mathbb{P}$, we consider an arbitrary total assignment α extending some partial assignment in $\mathcal{H}'$ and satisfying $F' \wedge M'$. We will modify α to another total assignment β that extends some partial assignment in $\mathcal{H}$ and satisfies $F' \wedge M$, and furthermore has the property that $\beta(m) = \alpha(m)$ for every $m \in U$. By assumption, $F' \wedge M \models_{(\mathcal{Q},\mathcal{H})} \mathbb{P}$, and so $\beta(\mathbb{P}) = 0$. Since $\beta(m) = \alpha(m)$ for every $m \in U$, we conclude that $\alpha(\mathbb{P}) = 0$ as well.

We proceed to define β. For each clause $x^a \vee y^b$ in M_2, we will define β on $\mathcal{Q}[x], \mathcal{Q}[y]$ using partial assignments from $\mathcal{H}$, distinguishing two cases: the clause is matched to some monomial in U_2, or it is *unmatched*. The values of all the other variables are taken directly from α.

Suppose $m \in U_2$ is matched to the clauses $x^a \vee y^b$ and $z^c \vee w^d$ and $C_m = r^e \vee s^f$, where $\mathcal{Q}[x] = \mathcal{Q}[r]$ and $\mathcal{Q}[z] = \mathcal{Q}[s]$. (In other words, we are in exactly the same situation described above while constructing M'.) Define β on $\mathcal{Q}[x], \mathcal{Q}[y], \mathcal{Q}[z], \mathcal{Q}[w]$ using partial assignments from $\mathcal{H}$ satisfying r^e, y^b, s^f, w^d. As a result, β satisfies the clauses $x^a \vee y^b$ and $z^c \vee w^d$ and the monomial m.

For each unmatched clause $x^a \vee y^b$ in M_2, we define β on $\mathcal{Q}[x]$ and $\mathcal{Q}[y]$ using partial assignments from $\mathcal{H}$ satisfying x^a and y^b. As a result, β satisfies the clause $x^a \vee y^b$. Finally, complete the definition of β by defining $\beta(x) = \alpha(x)$ for any hitherto undefined variable x. From the construction it is clear that β extends some partial assignment in $\mathcal{H}$.

In order to complete the proof, we need to show that β satisfies $F' \wedge M$, and that β agrees with α on all the monomials in U. We start by showing that β satisfies $F' \wedge M$. By construction, β satisfies the clauses in M_2. Since β agrees with α on variables mentioned in M_1, β satisfies M_1. Finally, let $C \in F'$. Since $(\mathcal{Q}, \mathcal{H})$ respects F', either the variables in C are disjoint from $\bigcup \mathcal{Q}$, or the variables in C all belong to some $Q_i \in \mathcal{Q}$, and all assignments in the respective $H_i \in \mathcal{H}$ satisfy C. In the former case, β agrees with α on variables mentioned in C, and so β satisfies C. In the latter case, β satisfies C since β extends some partial assignment in $\mathcal{H}$.

It remains to show that $\beta(m) = \alpha(m)$ for all monomials $m \in U$. In short, this is true for monomials in U_1 since α and β agree on all the relevant variables, and for monomials in U_2 since in both assignments they are reduced to zero. We proceed to show this formally.

Suppose first that $m \in U_1$. We claim that $\alpha(v) = \beta(v)$ for all variables v mentioned in m. Indeed, if $\alpha(v) \neq \beta(v)$ then $v \in \mathcal{Q}[x]$ for some clause $C = x^a \vee y^b$ in M_2. Yet this implies that m is connected to C, contradicting the definition of M_2. We conclude that α and β agree on all variables mentioned in m, and so $\alpha(m) = \beta(m)$ in this case.

Suppose next that $m \in U_2$. We claim that $\alpha(m) = \beta(m) = 0$. Let $C_m = r^e \vee s^f$, and recall that m is of the form $m = r^e s^f m'$. Thus $\alpha(m) = 0$ since α satisfies C_m, and $\beta(m) = 0$ since it satisfies r^e and s^f by construction. □

Theorem A.7 (restatement of Theorem 2.6 [BG13]). *Let F be a CNF formula with an r-extendible family $\mathcal{F}$ with respect to some $F' \subseteq F$. Then $Sp_{\mathcal{FC}}(F \vdash \bot) \geq r/4$.*

Proof. Let $\mathcal{F}$ be an r-extendible family with respect to some satisfiable $F' \subseteq F$. Let π be a derivation from F in space $Sp(\pi) < r/4$. We will show that $1 \notin \pi$ or, even stronger, that every configuration $\mathbb{P}_t$ appearing in π is satisfiable.

We will maintain a satisfiability witness $(F'; \mathcal{Q}_t, \mathcal{H}_t, M_t)$ for every configuration $\mathbb{P}_t$. Our satisfiability witnesses will satisfy two conditions: $(\mathcal{Q}_t, \mathcal{H}_t) \in \mathcal{F}$, and the *size bound* $|M_t| \leq 2Sp(\mathbb{P}_t)$. The existence of a satisfiability witness implies that $\mathbb{P}_t$ is satisfiable. Indeed, let $\alpha \in \mathcal{H}_t$ be some partial assignment that satisfies all the literals in M_t. Since $(\mathcal{Q}_t, \mathcal{H}_t)$ respects F', each clause in F' is either already satisfied by α or is completely disjoint from the domain of α. As F' is satisfiable, we can extend α to a total assignment β which satisfies F'. Hence, from $F' \wedge M_t \models_{(\mathcal{Q}_t,\mathcal{H}_t)} \mathbb{P}_t$ we have that β satisfies $\mathbb{P}_t$, and so $\mathbb{P}_t$ is satisfiable.

We construct the satisfiability witnesses by induction. For $t = 0$, the satisfiability witness is $(F'; \emptyset, \emptyset, \emptyset)$. For the induction step, suppose we are given a satisfiability witness $(F'; \mathcal{Q}, \mathcal{H}, M)$ for $\mathbb{P}_t$. We will construct a satisfiability witness $(F'; \mathcal{Q}', \mathcal{H}', M')$ for $\mathbb{P}_{t+1}$. To simplify the notation, let $\mathbb{P} = \mathbb{P}_t$ and $\mathbb{P}' = \mathbb{P}_{t+1}$. We distinguish three cases, which correspond to the three possible steps in the proof.

Axiom download. Let C be the downloaded clause, which we also regard as a monomial. If $C \in F'$ or every extension α of a partial assignment in $\mathcal{H}$ satisfies C, then in particular $F' \wedge M \models_{(\mathcal{Q},\mathcal{H})} \mathbb{P} \cup \{C\} = \mathbb{P}'$, and $M' = M$, $\mathcal{Q}' = \mathcal{Q}$, $\mathcal{H}' = \mathcal{H}$ form a satisfiability witness.

Otherwise, by hypothesis $Sp(\mathbb{P}') < r/4$ and so $Sp(\mathbb{P}) < r/4 - 1$. Indeed, if U is a defining set of monomials of $\mathbb{P}$, then $U \cup \{C\}$ is a defining set of monomials of $\mathbb{P}'$. By the induction hypothesis, $|\mathcal{Q}| < r - 1$. By the extension property of extendible, there exists a structured set of assignments $(\tilde{\mathcal{Q}}, \tilde{\mathcal{H}}) \in \mathcal{F}$ such that $|\tilde{\mathcal{Q}}| < r$, $(\mathcal{Q}, \mathcal{H}) \preccurlyeq (\tilde{\mathcal{Q}}, \tilde{\mathcal{H}})$ and $\tilde{\mathcal{H}} \models C$. By assumption $\mathcal{H} \not\models C$ and so $\mathcal{Q} \neq \tilde{\mathcal{Q}}$. Let $\tilde{\mathcal{Q}} = \mathcal{Q} \cup \{Q\}$.

The assignments corresponding to Q in $\tilde{\mathcal{H}}$ will ensure that the clause C is satisfied. Since we are going to add a new clause to M', we need to come up with two new parts in $\mathcal{Q}'$, and so we repeat the process. Let D be any axiom in $F \setminus F'$ such that $\tilde{\mathcal{H}} \not\models D$; if no such axiom exists then F is satisfiable and the theorem follows vacuously. Repeat the argument above and obtain a new disjoint set Q' and a structured set of assignments $(\mathcal{Q}', \mathcal{H}') \in \mathcal{F}$.

Choose arbitrary variables $x \in Q$ and $y \in Q'$, and let $M' = M \cup \{x \vee y\}$. By construction, $(F'; \mathcal{Q}', \mathcal{H}', M')$ is a satisfiability witness for $\mathbb{P}'$.

In both cases, by Lemma A.6 there is another satisfiability witness $(F'; \mathcal{Q}'', \mathcal{H}'', M'')$ for $\mathbb{P}'$ satisfying the size bound and with $\mathcal{Q}'' \subseteq \mathcal{Q}'$, $\mathcal{H}'' = \mathcal{H}'\restriction_{\mathcal{Q}''}$. By the restriction property of the extendible family, we have $(\mathcal{Q}'', \mathcal{H}'') \in \mathcal{F}$.

Inference. It is enough to pick $M' = M$, $\mathcal{Q}' = \mathcal{Q}$, $\mathcal{H}' = \mathcal{H}$. The first three properties in the definition of satisfiability witness continue to hold, while the last property follows from the soundness of FC. Finally, the size bound trivially holds since $|\mathbb{P}'| \geq |\mathbb{P}|$.

Erasure. Since FC is sound, $(F'; \mathcal{Q}, \mathcal{H}, M)$ is a satisfiability witness for $\mathbb{P}'$ as well. Hence Lemma A.6 furnishes us with a satisfiability witness $(F'; \mathcal{Q}', \mathcal{H}', M')$ for $\mathbb{P}'$ satisfying the size bound and with $\mathcal{Q}' \subseteq \mathcal{Q}$, $\mathcal{H}' = \mathcal{H}\restriction_{\mathcal{Q}'}$. By the restriction property of extendible, $(\mathcal{Q}', \mathcal{H}') \in \mathcal{F}$.

□

References

[ABRW02] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *SIAM Journal on Computing*, 31(4):1184–1211, 2002. Preliminary version appeared in *STOC '00*.

[AD08] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, May 2008. Preliminary version appeared in *CCC '03*.

[AR03] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003. Available at http://people.cs.uchicago.edu/~razborov/files/misha.pdf. Preliminary version appeared in *FOCS '01*.

[AS00] Noga Alon and Joel Spencer. *The probabilistic method*. Wiley-Interscience, 2nd edition, 2000.

[BBI12] Paul Beame, Chris Beck, and Russell Impagliazzo. Time-space tradeoffs in resolution: Superpolynomial lower bounds for superlinear space. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC '12)*, pages 213–232, May 2012.

[Ben09] Eli Ben-Sasson. Size space tradeoffs for resolution. *SIAM Journal on Computing*, 38(6):2511–2525, May 2009. Preliminary version appeared in *STOC '02*.

[BG03] Eli Ben-Sasson and Nicola Galesi. Space complexity of random formulae in resolution. *Random Structures and Algorithms*, 23(1):92–109, August 2003. Preliminary version appeared in *CCC '01*.

[BG13] Ilario Bonacina and Nicola Galesi. Pseudo-partitions, transversality and locality: A combinatorial characterization for the space measure in algebraic proof systems. In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science (ITCS '13)*, pages 455–472, January 2013.

[BL91] Béla Bollobás and Imre Leader. Edge-isoperimetric inequalities in the grid. *Combinatorica*, 11:299–314, 1991.

[Bla37] Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.

[BN08] Eli Ben-Sasson and Jakob Nordström. 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)*, pages 709–718, October 2008.

[BN11] Eli Ben-Sasson and Jakob Nordström. Understanding space in proof complexity: Separations and trade-offs via substitutions. In *Proceedings of the 2nd Symposium on Innovations in Computer Science (ICS '11)*, pages 401–416, January 2011. Full-length version available at http://eccc.hpi-web.de/report/2010/125/.

[BNT13] Chris Beck, Jakob Nordström, and Bangsheng Tang. Some trade-off results for polynomial calculus. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC '13)*, May 2013. To appear.

[Bol88] Béla Bollobás. The isoperimetric number of random regular graphs. *European Journal of Combinatorics*, 9:241–244, 1988.

[BS97] Roberto J. Bayardo Jr. and Robert Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI '97)*, pages 203–208, July 1997.

[BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version appeared in *STOC '99*.

[CEI96] Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 174–183, May 1996.

[CR79] Stephen A. Cook and Robert Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, March 1979.

[CS88] Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988.

[ET01] Juan Luis Esteban and Jacobo Torán. Space bounds for resolution. *Information and Computation*, 171(1):84–97, 2001. Preliminary versions of these results appeared in *STACS '99* and *CSL '99*.

[FLM+13] Yuval Filmus, Massimo Lauria, Mladen Mikša, Jakob Nordström, and Marc Vinyals. Towards an understanding of polynomial calculus: New separations and lower bounds (extended abstract). In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP '13)*, July 2013. to appear.

[FLN+12] Yuval Filmus, Massimo Lauria, Jakob Nordström, Neil Thapen, and Noga Ron-Zewi. Space complexity in polynomial calculus. In *Proceedings of the 27th Annual IEEE Conference on Computational Complexity (CCC '12)*, pages 334–344, June 2012.

[Hak85] Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985.

[HN12] Trinh Huynh and Jakob Nordström. On the virtue of succinct proofs: Amplifying communication complexity hardness to time-space trade-offs in proof complexity. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC '12)*, pages 233–248, May 2012.

[IK10] Russell Impagliazzo and Valentine Kabanets. Constructive proofs of concentration bounds. In *Proceedings of the 13th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems and 14th International Workshop on Randomization and Computation (APPROX-RANDOM '10)*, pages 617–631, 2010.

[IPS99] Russell Impagliazzo, Pavel Pudlák, and Jiri Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Computational Complexity*, 8(2):127–144, 1999.

[JŁR00] Svante Janson, Tomasz Łuczak, and Andrzej Ruciński. *Random graphs*. Wiley-Interscience, 2000.

[KW01] Jeong Han Kim and Nicholas C. Wormald. Random matchings which induce Hamilton cycles, and hamiltonian decompositions of random regular graphs. *Journal of Combinatorial Theory B*, 81:20–44, 2001.

[Mor94] M. Morgenstern. Existence and explicit constructions of $q + 1$ regular Ramanujan graphs for every prime power q. *Journal of Combinatorial Theory, Series B*, 62(1):44–62, 1994.

[MS96] João P. Marques-Silva and Karem A. Sakallah. GRASP—a new search algorithm for satisfiability. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD '96)*, pages 220–227, November 1996.

[PS97] Alessandro Panconesi and Aravind Srinivasan. Randomized distributed edge coloring via an extension of the Chernoff-Hoeffding bounds. *SIAM Journal on Computing*, 26(2):350–368, 1997.

[Raz98] Alexander A. Razborov. Lower bounds for the polynomial calculus. *Computational Complexity*, 7(4):291–324, December 1998.

[Urq87] Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, January 1987.

[Wor99] Nicholas C. Wormald. Models of random regular graphs. In *Surveys in Combinatorics*, pages 239–298. Cambridge University Press, 1999.