




ELSEVIER

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Space proof complexity for random 3-CNFs [☆]Patrick Bennett ^a, Ilario Bonacina ^b, Nicola Galesi ^{e,*}, Tony Huynh ^c,
Mike Molloy ^d, Paul Wollan ^e^a Mathematics Department, Western Michigan University, United States^b School of Computer Science and Communication KTH Royal Institute of Technology, Sweden^c Department of Mathematics, Université Libre de Bruxelles, Boulevard du Triomphe, B-1050 Brussels, Belgium^d Computer Science Department, University of Toronto, 10 Kings College Road, M5S 3G4 Toronto, Canada^e Computer Science Department, Sapienza University of Rome, via Salaria 113, 00198 Rome, Italy

ARTICLE INFO

Article history:

Received 16 October 2015

Received in revised form 5 April 2017

Available online 12 June 2017

Keywords:

Proof complexity
Polynomial calculus
Monomial space
Random CNFs
Resolution
Total space

ABSTRACT

We investigate the space complexity of refuting 3-CNFs in Resolution and algebraic systems. We prove that every *Polynomial Calculus with Resolution* refutation of a random 3-CNF φ in n variables requires, with high probability, $\Omega(n)$ distinct monomials to be kept simultaneously in memory. The same construction also proves that every *Resolution* refutation of φ requires, with high probability, $\Omega(n)$ clauses each of width $\Omega(n)$ to be kept at the same time in memory. This gives a $\Omega(n^2)$ lower bound for the total space needed in Resolution to refute φ . These results are best possible (up to a constant factor) and answer questions about space complexity of 3-CNFs.

The main technical innovation is a variant of *Hall's Lemma*. We show that in bipartite graphs with bipartition (L, R) and left-degree at most 3, L can be covered by certain families of disjoint paths, called *VW-matchings*, provided that L expands in R by a factor of $(2 - \epsilon)$, for $\epsilon < \frac{1}{5}$.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

The *space* of proving a theorem in a given proof system is the minimal memory occupation of an algorithm verifying the correctness of the proof. Since the initial study of space measure for proofs [1,19], its central role in proof complexity has become clear. The reason was initially theoretical, since proof space plays for proofs the analogous role as space complexity does for computations. Therefore, understanding why tautologies require high space led to numerous connections with many other proof complexity measures (like size, length, width, degree) [2,5–7,9,10,20,21,24].

At present, understanding proof space is becoming relevant in more applied algorithmic contexts such as SAT solvers. Using various heuristics, SAT solvers search for proofs in systems often studied in proof complexity. Hence, upper and

[☆] The main part of this work was completed while the authors I. Bonacina and T. Huynh were affiliated to the Computer Science Department of Sapienza University of Rome (Italy). I. Bonacina is supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013)/ERC grant agreement no. 279611. P. Wollan and Tony Huynh were supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013)/ERC grant agreement no. 279558. P. Bennett is supported by a grant from the Simons Foundation (#426894).

* Corresponding author.

E-mail addresses: patrick.bennett@wmich.edu (P. Bennett), ilario@kth.se (I. Bonacina), galesi@di.uniroma1.it (N. Galesi), tony.bourbaki@gmail.com (T. Huynh), molloy@cs.toronto.edu (M. Molloy), wollan@di.uniroma1.it (P. Wollan).




Fig. 1. Matchings, 2-matchings, VW-matchings.

lower bounds for these proof systems give information about the potential and limitations of such algorithms. For example, well-known SAT-solvers used in practice, like CDCL, are based on low-level proof systems such as *Resolution*, which is widely studied theoretically [25].

In this work we focus on two well-known proof systems: *Resolution* [8,26] and *Polynomial Calculus* [15]. Resolution (resolution) is a refutational proof system for unsatisfiable propositional CNF formulas using only one logical rule: $\frac{A \vee x \quad \neg x \vee B}{A \vee B}$. Polynomial calculus is an algebraic refutational proof system for unsatisfiable sets of polynomials (over $\{0, 1\}$ solutions) based on two rules: *linear combination* of polynomials and *multiplication* by variables. In this article, we consider the stronger system *Polynomial Calculus with Resolution* (PCR) which extends both Resolution and Polynomial Calculus [1].

Several different measures for proof space were investigated for these two systems [1,2,5–7,9,10,19,21,24]. In this work we focus on *total space* (for resolution), which is the maximum number of variables (counted with repetitions) to be kept simultaneously in memory while verifying a proof; and *monomial space* (for PCR), which is the maximum number of distinct monomials to be kept simultaneously in memory while verifying a proof. Both measures were introduced in [1], where some preliminary lower and upper bounds were given. In particular, for every unsatisfiable CNF in n variables, there is an easy upper bound of $O(n)$ for monomial space in PCR and $O(n^2)$ for total space in resolution.

Lower bounds for these two measures were initially studied in [1]. Several questions raised in that seminal work have only recently been answered [6,7,20]. In particular, in [6,7] the authors prove that, for $r \geq 4$, random r -CNFs over n variables require $\Theta(n^2)$ total space in resolution and $\Theta(n)$ monomial space in PCR. However, it is not at all obvious how to generalize the techniques in [6,7] to handle 3-CNFs. Indeed, before this work it was an open problem if there exists any family of 3-CNFs requiring large total space in resolution and monomial space in PCR.

1.1. Results

Let φ be a random 3-CNF in n variables. We prove that every PCR refutation of φ requires, with high probability, $\Omega(n)$ distinct monomials to be kept simultaneously in memory (Theorem 5.3). Moreover, every resolution refutation of φ has, with high probability, $\Omega(n)$ clauses each of width $\Omega(n)$ to be kept at the same time in memory (Theorem 5.3). This gives a $\Omega(n^2)$ lower bound for the total space of every resolution refutation of φ .¹ These results resolve questions about space complexity of 3-CNFs mentioned in [6,7,20,21].

Both results follow using the framework proposed in [6], where the construction of suitable families of assignments called *k-winning strategies* (Definition 2.1) leads to monomial space lower bounds in PCR (Theorem 2.2). This construction relies on a modification of Hall’s Lemma [23] from matchings to VW-matchings (Lemma 1.2).

Definition 1.1 (VW-matching). Let G be a bipartite graph with bipartition (L, R) . A VW-matching in G is a subgraph F of G such that each connected component of F is a path with at most 4 edges and both endpoints in R . A VW-matching F covers a set of vertices S if $S \subseteq V(F)$. Define $L(F) = V(F) \cap L$ and $R(F) = V(F) \cap R$.

Fig. 1 compares matchings (Fig. 1(a)), 2-matchings as used in [6,7] (Fig. 1(b)) and VW-matchings (Fig. 1(c)). Note that for technical reasons, we allow 2-matchings and VW-matchings to contain isolated vertices from R . We can now state our variant of Hall’s Lemma. This lemma and its proof are independent from the proof complexity results and might be useful in other contexts.

Lemma 1.2 ($(2 - \epsilon)$ -Hall’s lemma). Let $\epsilon < \frac{1}{5}$. Let G be a bipartite graph with bipartition (L, R) such that each vertex in L has degree at most 3 and no pair of degree 3 vertices in L have the same set of neighbors. If $|N_G(L)| \geq (2 - \epsilon)|L|$, and each proper subset of L can be covered by a VW-matching, then L can be covered by a VW-matching.

Note that the converse of Lemma 1.2 does not hold (unlike in Hall’s Lemma). For instance, the graph shaped as a W in Fig. 1(c) satisfies $|N_G(L)| = 3 < (2 - \epsilon)2$, whenever $\epsilon < 1/2$.

The original proof of Lemma 1.2 is from an earlier version of this work [4], where we showed that it holds for $\epsilon < \frac{1}{23}$. In our earlier version, we also conjectured that Lemma 1.2 holds for $\epsilon \leq \frac{1}{3}$. Since then, Susanna Figueiredo De Rezende

¹ Recently, [11] proved a general inequality between total space and width in resolution which also implies this result. Our proof has the advantage of being more constructive and explicit.

[18] simplified our proof and showed that $\epsilon < \frac{1}{5}$ suffices. Finally, very recently [27] proved our conjecture by showing that Lemma 1.2 does indeed hold for $\epsilon \leq \frac{1}{3}$.

In this paper we present the proof given in [18], as it is the shortest of the three, and the precise value of ϵ is not relevant for our results. We thank Susanna Figueiredo De Rezende for kindly allowing us to include it. In Proposition 3.1, we complement Lemma 1.2 by showing that Lemma 1.2 does not hold for $\epsilon > \frac{1}{3}$. Therefore, the result in [27] is best possible.

1.2. Outline of the paper

In section 2 we recall some preliminary notions about proof complexity. In particular, we recall the formal definitions of Resolution and Polynomial Calculus with Resolution, the model of space, and the formal definitions of total space and monomial space. Families of partial assignments, called k -winning strategies, were used in [6] to prove monomial space lower bounds for PCR. Here we use the same k -winning strategies² to prove, not only space lower bounds for PCR but also total space lower bounds for resolution.

In section 3, we prove our $(2 - \epsilon)$ -Hall's Lemma (Lemma 1.2). We also prove that Lemma 1.2 does not hold for $\epsilon > \frac{1}{3}$ (Proposition 3.1).

In section 4, we define a two player covering game CoverGame, whose aim is to dynamically build a VW-matching inside a fixed bipartite graph G (Definition 4.1). Informally, a player, Choose, queries nodes in the graph G and the other player, Cover, attempts to extend the current VW-matching to also cover the node queried (if not already covered). The main result of section 4 is Theorem 4.3, where we prove that if the graph G has large left-expansion (large enough to apply Lemma 1.2 to sufficiently large subgraphs of G), then there is a winning strategy for Cover to force Choose to query a very large portion of the graph G . In the analysis of the game, we use the $(2 - \epsilon)$ -Hall's Lemma and VW-matchings in a similar manner to how matchings and 2-matchings were used in [3,6,7,13]. A key difference is that we are looking for winning strategies of Cover for the CoverGame only on graphs G where the number of high degree vertices is suitably bounded (Theorem 4.3). This additional information allows us to identify a VW-matching covering all high degree vertices in G but preserving expansion properties of the remaining graph. Cover will use this information to obtain a winning strategy.

In section 5, we prove (Lemma 5.1), that if Cover wins CoverGame on the adjacency graph of a CNF φ (see section 2 for the definition of adjacency graph) guaranteeing VW-matchings of a maximum of μ connected components, then there exists a μ -winning strategy for the polynomial encoding of φ . Finally, the monomial space in PCR and the total space in resolution for random 3-CNFs (Theorem 5.3) follow from well-known results about expansion of its adjacency graph [12–14, 16]. In order to get optimal lower bounds, we show in Lemma 5.2 that the number of variables appearing in many clauses of a random CNF is w.h.p. suitably bounded as required in the conditions of Theorem 4.3.

2. Preliminaries

Let X be a set of variables. A *literal* is a boolean constant, 0 or 1, or a variable $x \in X$, or the negation $\neg x$ of a variable x . A *clause* is a disjunction of literals: $C = (\ell_1 \vee \dots \vee \ell_k)$. The *width* of a clause is the number of literals in it. A formula φ is in *Conjunctive Normal Form* (CNF) if $\varphi = C_1 \wedge \dots \wedge C_m$ where C_i are clauses. It is a k -CNF if each C_i contains at most k literals. Let φ be a CNF and X be the set of variables appearing in φ . The *adjacency graph* of φ is a bipartite graph G_φ with bipartition (L, R) such that L is the set of clauses of φ , $R = X$, and $(C, x) \in E$ if and only if x or $\neg x$ appears in C . If φ is a k -CNF, then G_φ has left-degree at most k .

A *partial assignment* over a set of variables X is a map $\alpha : X \rightarrow \{0, 1, \star\}$. The *domain* of α is $\text{dom}(\alpha) = \alpha^{-1}(\{0, 1\})$. Given a partial assignment α and a CNF φ we can apply α to φ , obtaining a new formula $\alpha(\varphi)$ in the standard way, i.e. substituting each variable x of φ in $\text{dom}(\alpha)$ with the value $\alpha(x)$ and then simplifying the result. We say that α *satisfies* φ , and we write $\alpha \models \varphi$, if $\alpha(\varphi) = 1$. Similarly, for a family F of partial assignments, $F \models \varphi$ means that for each $\alpha \in F$, $\alpha \models \varphi$.

Resolution [8,26] is a propositional proof system for refuting unsatisfiable CNFs. Starting from an unsatisfiable CNF φ , resolution allows us to derive the empty clause \perp using the following inference rule:

$$\frac{C \vee x \quad D \vee \neg x}{C \vee D}.$$

Following [1], we define $\bar{X} = \{\bar{x} : x \in X\}$, which we regard as a set of formal variables with the intended meaning of \bar{x} as $\neg x$. Given a field \mathbb{F} , the ring $\mathbb{F}[X, \bar{X}]$ is the ring of polynomials in the variables $X \cup \bar{X}$ with coefficients in \mathbb{F} . We use the following *standard encoding* (tr) of CNF formulas over X into a set of polynomials in $\mathbb{F}[X, \bar{X}]$: $\text{tr}(\varphi) = \{\text{tr}(C) : C \in \varphi\} \cup \{x^2 - x, x + \bar{x} - 1 : x \in X\}$, where

$$\text{tr}(x) = \bar{x}, \quad \text{tr}(\neg x) = x, \quad \text{tr}\left(\bigvee_{i=1}^n \ell_i\right) = \prod_{i=1}^n \text{tr}(\ell_i).$$

² For simplicity, the definition we present here is only a special case of the definition of a k -winning strategy from [6].

A set of polynomials P in $\mathbb{F}[X]$ is *contradictory* if and only if 1 is in the ideal generated by P . Notice that a CNF φ is unsatisfiable if and only if $\text{tr}(\varphi)$ is a contradictory set of polynomials.

For each partial assignment α over $X \cup \bar{X}$ we assume that it respects the intended meaning of the variables; that is, $\alpha(\bar{x}) = 1 - \alpha(x)$ for each $x, \bar{x} \in \text{dom}(\alpha)$. Given a partial assignment α and a polynomial p in $\mathbb{F}[X, \bar{X}]$, we can apply α to p , obtaining a new polynomial $\alpha(p)$ in the standard way, similarly as before. The notation $\alpha \models p$ means that $\alpha(p) = 0$. If F is a family of partial assignments and P a set of polynomials, we write $F \models P$ if $\alpha \models p$ for each $\alpha \in F$ and $p \in P$. Notice that if φ is a CNF and α is a partial assignment then $\alpha \models \varphi$ if and only if $\alpha \models \text{tr}(\varphi)$.

Polynomial Calculus with Resolution (PCR) [1] is an algebraic proof system for polynomials in $\mathbb{F}[X, \bar{X}]$. Starting from an initial set of contradictory polynomials P in $\mathbb{F}[X, \bar{X}]$, PCR allows us to derive the polynomial 1 using the following inference rules: for all $p, q \in \mathbb{F}[X, \bar{X}]$

$$\frac{p}{\alpha p + \beta q} \forall \alpha, \beta \in \mathbb{F}, \quad \frac{p}{vp} \forall v \in X \cup \bar{X}.$$

To force 0/1 solutions, we always include the *boolean axioms* $\{x^2 - x, x + \bar{x} - 1\}_{x \in X}$ among the initial polynomials, as in the case of the polynomial encoding of CNFs.

In order to study space of proofs we follow a model inspired by the definition of space complexity for Turing machines, where a machine is given a read-only input tape from which it can download parts of the input to the working memory as needed [19].

Given an unsatisfiable CNF formula φ , a *resolution (resp. PCR) refutation* of φ is a sequence $\Pi = \langle M_0, \dots, M_\ell \rangle$ of sets of clauses (resp. polynomials), called *memory configurations*, such that: $M_0 = \emptyset$, $\perp \in M_\ell$ (resp. $1 \in M_\ell$), and for all $i \leq \ell$, M_i is obtained by M_{i-1} by applying one of the following rules:

(AXIOM DOWNLOAD) $M_i = M_{i-1} \cup \{C\}$, where C is a clause of φ (resp. a polynomial of $\text{tr}(\varphi)$);

(INFERENCE ADDING) $M_i = M_{i-1} \cup \{O\}$, where O is inferred by the resolution inference rule (resp. PCR inference rules) from clauses (resp. polynomials) in M_{i-1} ;

(ERASURE) $M_i \subset M_{i-1}$.

If in the definition of PCR refutation we substitute the INFERENCE ADDING rule with:

(SEMANTICAL INFERENCE) M_i is contained in the ideal generated by M_{i-1} in $\mathbb{F}[X, \bar{X}]$,

we have what is called a *semantical PCR refutation* of φ [1].

The *total space* of Π is the maximum over i of the number of variables (counted with repetitions) occurring in M_i .

The *monomial space* of a PCRrefutation Π , denoted by $\text{MSpace}(\Pi)$, is the maximum over i of the number of *distinct* monomials appearing in M_i .

2.1. A framework for space lower bounds

Let A be a family of partial assignments, and let $\text{dom}(A)$ be the union of the domains of the assignments in A . We say that a set of partial assignments A is *flippable* if and only if for all $x \in \text{dom}(A)$ there exist $\alpha, \beta \in A$ such that $\alpha(x) = 1 - \beta(x)$. Two families of partial assignments A and A' are *domain-disjoint* if $\text{dom}(\alpha)$ and $\text{dom}(\alpha')$ are disjoint for all $\alpha \in A$ and $\alpha' \in A'$. Given non-empty and pairwise domain-disjoint sets of assignments³ H_1, \dots, H_t , the *product-family* $\mathcal{H} = H_1 \otimes \dots \otimes H_t$ is the following set of assignments

$$\mathcal{H} = H_1 \otimes \dots \otimes H_t = \{\alpha_1 \cup \dots \cup \alpha_t : \alpha_i \in H_i\},$$

or, if $t = 0$, $\mathcal{H} = \{\lambda\}$, where λ is the partial assignment of the empty domain. Note $\text{dom}(\mathcal{H}) = \bigcup_i \text{dom}(H_i)$. We call the H_i the *factors* of \mathcal{H} . For a product-family $\mathcal{H} = H_1 \otimes \dots \otimes H_t$, the *rank* of \mathcal{H} , denoted $\|\mathcal{H}\|$, is the number of factors of \mathcal{H} different from $\{\lambda\}$. We do not count $\{\lambda\}$ in the rank since $\mathcal{H} \otimes \{\lambda\} = \mathcal{H}$. Given two product-families \mathcal{H} and \mathcal{H}' , we write $\mathcal{H}' \sqsubseteq \mathcal{H}$ if and only if each factor of \mathcal{H}' different from $\{\lambda\}$ is also a factor of \mathcal{H} . In particular, $\{\lambda\} \sqsubseteq \mathcal{H}$ for every \mathcal{H} .

A family of flippable product-families is called a *strategy* and denoted by \mathcal{L} . We now present a definition of suitable families of flippable products: the *k-winning strategies* [6].

Definition 2.1 (*k-winning strategy* [6]). Let P be a set of polynomials in the ring $\mathbb{F}[X, \bar{X}]$. A non-empty strategy \mathcal{L} is a *k-winning strategy* for P if and only if for every $\mathcal{H} \in \mathcal{L}$ the following conditions hold:

(restriction) for each $\mathcal{H}' \sqsubseteq \mathcal{H}$, $\mathcal{H}' \in \mathcal{L}$;

(extension) if $\|\mathcal{H}\| < k$, then for each $p \in P$ there exists a flippable product-family $\mathcal{H}' \in \mathcal{L}$ such that $\mathcal{H}' \sqsupseteq \mathcal{H}$ and $\mathcal{H}' \models p$.

³ We always suppose that the partial assignments respect the intended meaning of the variables in \bar{X} . That is, if $x \in \text{dom}(\alpha)$, then $\alpha(\bar{x}) = 1 - \alpha(x)$; hence a variable x is in $\text{dom}(H_i)$ if and only if \bar{x} is in $\text{dom}(H_i)$.

Notice that, by the restriction property, $\{\lambda\}$ is in every k -winning strategy.

Theorem 2.2. *Let φ be an unsatisfiable CNF and $k \geq 1$ an integer. If there exists a non-empty k -winning strategy \mathcal{L} for $\text{tr}(\varphi)$, then for every semantical PCR refutation Π of φ , $\text{MSpace}(\Pi) \geq \frac{k}{4}$. Moreover, every resolution refutation of φ must pass through a memory configuration containing at least $\frac{k-1}{2}$ clauses each of width at least $\frac{k-1}{2}$. In particular, the resolution refutation requires total space at least $\frac{(k-1)^2}{4}$.*

The monomial space lower bound follows directly from the main theorem of [6].

We show now how to use k -winning strategies to construct the combinatorial objects used in [7] to obtain total space lower bound stated in the theorem.

A *piecewise (p.w.) assignment* α of a set of variables X is a set of non-empty partial assignments to X with pairwise disjoint domains. We will sometimes call the elements of α the *pieces* of α . A piecewise assignment gives rise to a partial assignment $\bigcup \alpha$ to X together with a partition of the domain of $\bigcup \alpha$. For piecewise assignments α, β we will write $\alpha \sqsubseteq \beta$ to mean that every piece of α appears in β . We will write $\|\alpha\|$ to mean the number of pieces in α . Note that these are formally exactly the same as $\alpha \subseteq \beta$ and $|\alpha|$, if we regard α and β as sets.

Definition 2.3 (*r-free* [7]). A family \mathcal{F} of p.w. assignments is *r-free* for a CNF φ if it has the following properties:

(Consistency) No $\alpha \in \mathcal{F}$ falsifies any clause from φ ;

(Retraction) If $\alpha \in \mathcal{F}$, β is a p.w. assignment and $\beta \sqsubseteq \alpha$, then $\beta \in \mathcal{F}$;

(Extension) If $\alpha \in \mathcal{F}$ and $\|\alpha\| < r$, then for every variable $x \notin \text{dom}(\alpha)$, there exist $\beta_0, \beta_1 \in \mathcal{F}$ with $\alpha \sqsubseteq \beta_0, \beta_1$ such that $\beta_0(x) = 0$ and $\beta_1(x) = 1$.

Theorem 2.4 ([7]). *Let φ be an unsatisfiable CNF formula. If there is a family of p.w. assignments which is r -free for φ , then any resolution refutation of φ must pass through a memory configuration containing at least $\frac{r}{2}$ clauses each of width at least $\frac{r}{2}$. In particular, the refutation requires total space at least $\frac{r^2}{4}$.*

By this theorem, in order to prove the total space lower bound of Theorem 2.2 we just have to prove that given a k -winning strategy for $\text{tr}(\varphi)$ we can build a $(k-1)$ -free family for φ .

Proposition 2.5. *Let φ be an unsatisfiable CNF. Given a k -winning strategy for $\text{tr}(\varphi)$ there exists a $(k-1)$ -free family for φ .*

Proof. Let \mathcal{L} be the k -winning strategy. Define the $(k-1)$ -free family \mathcal{F} as follows: $\alpha \in \mathcal{F}$ if and only if there exists $H_1 \otimes \cdots \otimes H_t \in \mathcal{L}$ such that $\alpha = \alpha_1 \cup \dots \cup \alpha_t$ with $\alpha_i \in H_i$ and $t \leq k-1$. The p.w. structure of α is inherited from the domain-disjointness of $H_1 \otimes \cdots \otimes H_t$; in particular, $\|\alpha\| = \|H_1 \otimes \cdots \otimes H_t\|$. The *retraction property* of \mathcal{F} is immediate from the corresponding property of \mathcal{L} .

To prove the *consistency property* of \mathcal{F} assume, by contradiction, that there is an $\alpha \in \mathcal{F}$ such that α falsifies some clause $C \in \varphi$. Since $\|\alpha\| \leq k-1 < k$, there exists $\mathcal{H} = H_1 \otimes \cdots \otimes H_t \in \mathcal{L}$ such that $\alpha \in \mathcal{H}$ and $\|\alpha\| = \|\mathcal{H}\|$. By the extension property of \mathcal{L} , there is an $\mathcal{H}' \supseteq \mathcal{H}$ such that $\mathcal{H}' \models \text{tr}(C)$. In particular there exists some partial assignment $\beta \supseteq \alpha$ such that $\beta \models \text{tr}(C)$. By construction, for every assignment γ , $\gamma \models \text{tr}(C)$ if and only if $\gamma \models C$. Thus $\beta \models C$, which is impossible since α falsifies C .

For the *extension property* let $\alpha \in \mathcal{F}$, with $\|\alpha\| < k-1$ and let x be a variable of φ not in $\text{dom}(\alpha)$. By construction, there exists some $\mathcal{H} \in \mathcal{L}$ such that $\alpha \in \mathcal{H}$, $\|\alpha\| = \|\mathcal{H}\|$ and $\text{dom}(\alpha) = \text{dom}(\mathcal{H})$. By the extension property of \mathcal{F} there exists some flippable $\mathcal{H}' \in \mathcal{L}$ such that $\mathcal{H}' \supseteq \mathcal{H}$ and $\mathcal{H}' \models x^2 - x$. By taking restrictions in \mathcal{L} we can suppose that $\|\mathcal{H}'\| = \|\mathcal{H}\| + 1$. Hence there exist $\beta_0, \beta_1 \in \mathcal{F}$ extending α , setting x respectively to 0 and 1 and such that $\|\beta_0\| = \|\beta_1\| = \|\alpha\| + 1 \leq k-1$. \square

3. A $(2 - \epsilon)$ -Hall's Lemma for VW-matchings

We now prove our variant of Hall's Lemma. We use the *discharging method*, which is a standard tool in graph theory (see [17]). The idea is to give each vertex some initial *charge* and then to redistribute the charge according to some rules that do not change the total charge. The final distribution of charges typically reveals some structure of the graph (in our case the existence of our required VW-matching).

Restated Lemma 1.2 (*$(2 - \epsilon)$ -Hall's lemma*). *Let $\epsilon < \frac{1}{5}$. Let G be a bipartite graph with bipartition (L, R) such that each vertex in L has degree at most 3 and no pair of degree 3 vertices in L have the same set of neighbors. If $|N_G(L)| \geq (2 - \epsilon)|L|$, and each proper subset of L can be covered by a VW-matching, then L can be covered by a VW-matching.*

Proof. Define a hypergraph $\mathcal{H} = (V, E)$, where $V = N_G(L)$ and $E = \{N_G(x) : x \in L\}$. If $\deg_G(v) = 1$, then $\{v\}$ cannot be covered by a VW-matching, which is a contradiction. If u and v are distinct vertices with $\deg_G(u) = \deg_G(v) = 2$, then $\{u, v\}$ cannot be covered by a VW-matching, which is a contradiction. Also note that by assumption, no degree 3 vertices of

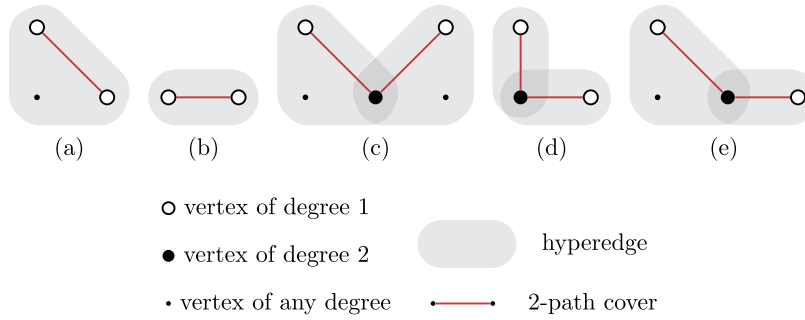


Fig. 2. A set of reducible configurations for \mathcal{H} .

G have the same neighborhood. The above remarks imply that all hyperedges of \mathcal{H} have size 2 or 3, and $N_G : L \rightarrow E$ is a bijection. Moreover, by assumption $|V| \geq (2 - \epsilon)|L| = (2 - \epsilon)|E|$.

The degree of a vertex v in \mathcal{H} , denoted $\text{deg}_{\mathcal{H}}(v)$, is the number of distinct hyperedges which contain v . Let $L' \subseteq L$ and let $E' = \{N_G(x) : x \in L'\}$. Suppose there is a VW-matching M in G covering L' . For each $e = N_G(x) \in E'$ we define $f(e)$ to be $e \cap N_M(x)$. Since each component of M is either a V or a W it follows that

1. $f(e)$ is a subset of size 2 of e , for all $e \in E'$;
2. $f(e_1) \neq f(e_2)$, for all distinct $e_1, e_2 \in E'$;
3. $f(e_1) \cap f(e_2) = \emptyset$ or $f(e_2) \cap f(e_3) = \emptyset$, for all distinct $e_1, e_2, e_3 \in E'$.

Conversely, if there exists a function f satisfying the above three conditions, then L' can be covered by a VW-matching. We call such a function f a 2-path cover of E' .

Observe that all the configurations shown in Fig. 2 have a 2-path cover that is disjoint from all hyperedges of \mathcal{H} not in the configuration. Therefore, if any of these configurations appear in \mathcal{H} , we can by assumption find a 2-path cover f of the remaining hyperedges, and then extend f to a 2-path cover of \mathcal{H} .

We may hence assume that no configuration from Fig. 2 appears in \mathcal{H} . Two vertices u and v of \mathcal{H} are adjacent if there is a hyperedge containing both u and v .

We now give each vertex $v \in V$ a charge of $\text{deg}_{\mathcal{H}}(v)$. We then redistribute the charges according to the following discharging rule. For each vertex $v \in V$ such that $\text{deg}_{\mathcal{H}}(v) = 1$, v receives $\frac{1}{3}$ units of charge from each vertex u adjacent to v (and u loses $\frac{1}{3}$ units of charge). For each $v \in V$ we let $w(v)$ be the charge of v after the discharging process. Note that the following properties hold.

1. $\sum_{v \in V} \text{deg}_{\mathcal{H}}(v) = \sum_{v \in V} w(v)$, since the discharging rule preserves the total charge.
2. If $\text{deg}_{\mathcal{H}}(v) = 1$ and e is the unique hyperedge containing v , then

$$w(v) = \begin{cases} 1 + 1/3 & \text{if } |e| = 2 \\ 1 + 2/3 & \text{if } |e| = 3. \end{cases}$$

Indeed, since the configurations in Figs. 2(a) and (b) do not appear in \mathcal{H} , no two degree 1 vertices of \mathcal{H} are adjacent.

3. If $\text{deg}_{\mathcal{H}}(v) = 2$, then

$$w(v) \geq 2 - 1/3.$$

Indeed, since the configurations in Figs. 2(c), (d) and (e) do not appear in \mathcal{H} , there is at most one vertex u such that $\text{deg}_{\mathcal{H}}(u) = 1$ and u and v are adjacent.

4. If $\text{deg}_{\mathcal{H}}(v) \geq 3$, then

$$w(v) \geq 2.$$

Indeed, suppose $\text{deg}_{\mathcal{H}}(v) = d \geq 3$. In particular, there are at most d degree 1 vertices in \mathcal{H} that are adjacent to v . Thus,

$$w(v) \geq d - \frac{d}{3} = \frac{2}{3}d \geq 2.$$

Let $E = E_2 \cup E_3$, where E_i are the edges of size i in E . Let V_1 be the set of degree 1 vertices of V that belong to some edge in E_2 and let $V_2 = V \setminus V_1$. Since Fig. 2(b) does not appear in \mathcal{H} , we have $|V_1| \leq |E_2|$. Also, by the above observations, $w(v) = \frac{4}{3}$ for all $v \in V_1$ and $w(v) \geq \frac{5}{3}$ for all $v \in V_2$. Therefore

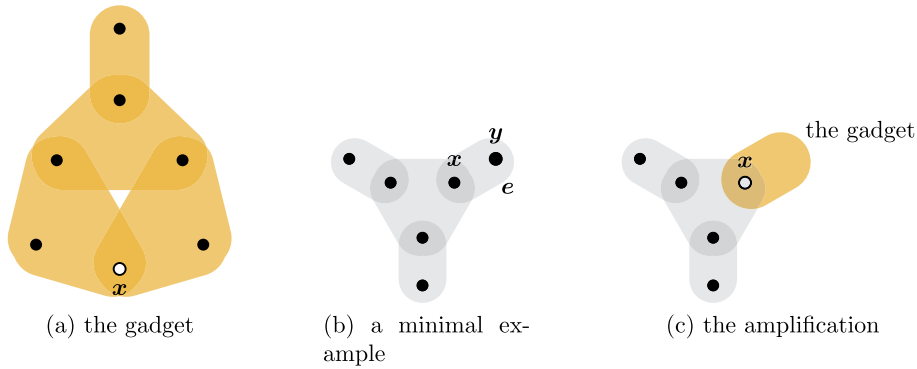


Fig. 3. The construction.

$$\begin{aligned}
 \frac{3|V|}{(2-\epsilon)} &\geq 3|E| = 3|E_3| + 2|E_2| + |E_2| = \sum_{v \in V} \deg_{\mathcal{H}}(v) + |E_2| = \\
 &= \sum_{v \in V} w(v) + |E_2| = \sum_{v \in V_1} w(v) + \sum_{v \in V_2} w(v) + |E_2| \geq \\
 &\geq \frac{5}{3}|V| - \frac{1}{3}|V_1| + |E_2| \geq \frac{5}{3}|V|.
 \end{aligned}$$

It follows that $\epsilon \geq 1/5$. But this contradicts that $\epsilon < 1/5$. \square

We end this section with a comment on the parameter ϵ . Since VW-matchings expand by a factor of at least $\frac{3}{2}$, we certainly require $\epsilon \leq \frac{1}{2}$ in the statement of Lemma 1.2. In Proposition 3.1 we show a stronger upper bound that $\epsilon \leq \frac{1}{3}$ is necessary.

Proposition 3.1. For all $\epsilon > \frac{1}{3}$ there exists a bipartite graph G_ϵ with bipartition (L, R) such that each vertex in L has degree at most 3 and no pair of degree 3 vertices in L have the same set of neighbors. Moreover, $|N_{G_\epsilon}(L)| \geq (2 - \epsilon)|L|$ and each proper subset of L can be covered by a VW-matching but L cannot be covered by a VW-matching.

In an earlier version of this work [4], we conjectured that Lemma 1.2 holds for $\epsilon \leq \frac{1}{3}$. Very recently, [27] proved our conjecture, which is best possible by Proposition 3.1.

We now prove Proposition 3.1, rephrased in terms of hypergraphs.

Proposition 3.2. For every $\epsilon > \frac{1}{3}$, there exists a hypergraph \mathcal{H}_ϵ such that \mathcal{H}_ϵ has no isolated vertices, each hyperedge of \mathcal{H}_ϵ has size 2 or 3, $|V(\mathcal{H})| \geq (2 - \epsilon)|E(\mathcal{H})|$, every proper subset of $E(\mathcal{H}_\epsilon)$ has a 2-path cover, but \mathcal{H}_ϵ does not have a 2-path cover.

Proof. Let $\epsilon > \frac{1}{3}$ and consider the gadget \mathcal{G} shown in Fig. 3(a). It is easy to verify that every 2-path cover of \mathcal{G} must cover the vertex x . Next note that the hypergraph \mathcal{H} shown in Fig. 3(b) is obviously not 2-path coverable, but every proper subset of $E(\mathcal{H})$ is 2-path coverable. We have $\frac{|V(\mathcal{H})|}{|E(\mathcal{H})|} = \frac{6}{4}$. However, we can increase this ratio via the amplification trick shown in Fig. 3(c).

That is, let e be a hyperedge of \mathcal{H} of size 2. Label the vertices of e as x and y , where y has degree 1. Let \mathcal{H}_1 be the hypergraph obtained from \mathcal{H} by deleting y and then gluing \mathcal{G} to $\mathcal{H} - y$ along x . Since every 2-path cover of \mathcal{G} must use the vertex x , \mathcal{H}_1 does not have a 2-path cover. On the other hand, since every proper subset of $E(\mathcal{G})$ has a 2-path cover avoiding x , it follows that every proper subset of $E(\mathcal{H}_1)$ has a 2-path cover. Note that this amplification trick increases the number of vertices of \mathcal{H} by 10 and the number of edges of \mathcal{H} by 6. Moreover, we can repeat this amplification trick arbitrarily many times since \mathcal{G} also has pendent edges of size 2. So, choose n such that

$$\frac{6 + 5n}{4 + 3n} \geq 2 - \epsilon$$

and take \mathcal{H}_ϵ to be the graph obtained from \mathcal{H} by performing the amplification trick n times. \square

4. A cover game over bipartite graphs

As an application, we use the previous result to build a winning strategy for a game played on bipartite graphs.

Definition 4.1 (*Cover game*). The *Cover Game* $\text{CoverGame}_{\text{VW}}(G, \mu)$ is a game between two players, Choose and Cover, on a bipartite graph G with bipartition (L, R) . At each step i of the game the players maintain a VW-matching F_i in G . At step $i + 1$ Choose can

1. remove a connected component from F_i , or
2. if the number of connected components of F_i is strictly less than μ , pick a vertex (either in L or R) and challenge Cover to find a VW-matching F_{i+1} in G such that
 - (a) F_{i+1} extends F_i . That is, each connected component of F_i is also a connected component of F_{i+1} ;
 - (b) F_{i+1} covers the vertex picked by Choose.

Cover loses the game $\text{CoverGame}_{\text{VW}}(G, \mu)$ if at some point she cannot answer a challenge by Choose. Otherwise, Cover wins.

Definition 4.2 ((s, δ) -bipartite expander). Let s be a positive integer and δ be a positive real number. A bipartite graph G with bipartition (L, R) is an (s, δ) -bipartite expander if all subsets $X \subseteq L$ of size at most s satisfy $|N_G(X)| \geq \delta|X|$.

The next theorem shows that Cover has a winning strategy for the game $\text{CoverGame}_{\text{VW}}(G, \mu)$ for expander graphs G with appropriately chosen parameters.

Theorem 4.3. Let G be a bipartite graph with bipartition (L, R) , s, D be integers, and $\epsilon < \frac{1}{5}$ be a real number. For every integer $d \geq D$ let $S_d \subseteq R$ be the set of vertices of R with degree larger than d . Suppose that

1. each vertex in L has degree 3;
2. G is an $(s, 2 - \frac{\epsilon}{2})$ -bipartite expander;
3. for every $d \geq D$, $\frac{72d}{\epsilon}(|S_d| + d) + 1 \leq \frac{s}{2}$.

Then Cover wins the cover game $\text{CoverGame}_{\text{VW}}(G, \mu)$ with $\mu = \frac{\epsilon s}{144D}$.

The proof of this result is similar to constructions that can be found in [3,7,13].

For the rest of this section, fix a bipartite graph G with bipartition (L, R) , an integer s , and a real number $\epsilon < \frac{1}{5}$ such that G is an $(s, 2 - \frac{\epsilon}{2})$ -bipartite expander where each vertex in L has degree 3. Given $A \subseteq L$ and $B \subseteq R$, we let $G_{A,B}$ be the subgraph of G induced by $(L \cup R) \setminus (A \cup B)$.

Definition 4.4 (*VW-matching property*). Given two sets $A \subseteq L$ and $B \subseteq R$, we say that the pair (A, B) has the *VW-matching property*, if for every $C \subseteq L \setminus A$ with $|C| \leq s$, there exists a VW-matching F in $G_{A,B}$ covering C .

Lemma 4.5. Let $A \subseteq L$ and $B \subseteq R$ be such that the pair (A, B) does not have the VW-matching property. Then there exists a set $C \subseteq L \setminus A$ with $|C| < \frac{2}{\epsilon}|B|$, such that no VW-matching in $G_{A,B}$ covers C .

Proof. Take $C \subseteq L \setminus A$ of minimal size such that no VW-matching in $G_{A,B}$ covers C . We have that $|C| \leq s$ and by minimality of C and Lemma 1.2 it follows that

$$|N_{G_{A,B}}(C)| < (2 - \epsilon)|C|.$$

But, by hypothesis G is an $(s, 2 - \frac{\epsilon}{2})$ -bipartite expander; hence $(2 - \frac{\epsilon}{2})|C| \leq |N_G(C)|$. Therefore,

$$(2 - \frac{\epsilon}{2})|C| \leq |N_G(C)| \leq |N_{G_{A,B}}(C)| + |B| < (2 - \epsilon)|C| + |B|.$$

Hence $|C| < \frac{2}{\epsilon}|B|$, as required. \square

Lemma 4.5 is the only place where we directly use the $(2 - \epsilon)$ -Hall's Lemma (Lemma 1.2) from the previous section. However, Lemma 4.5 itself plays a crucial role in proving the following Lemmas.

Lemma 4.6. The pair (\emptyset, \emptyset) has the VW-matching property.

Proof. Apply Lemma 4.5 with $A = \emptyset$ and $B = \emptyset$. \square

Lemma 4.7 (*Component removal*). Let $A \subseteq L$ and $B \subseteq R$ be such that the pair (A, B) has the VW-matching property and $\frac{2}{\epsilon}|B| \leq s$. Then for each VW-matching F contained in the subgraph of G induced by $A \cup B$, $(A \setminus L(F), B \setminus R(F))$ has the VW-matching property.

Proof. Let $A' = A \setminus L(F)$ and $B' = B \setminus R(F)$ and suppose, by contradiction, that (A', B') does not have the VW-matching property. By Lemma 4.5, it is sufficient to prove that for each set $C \subseteq L \setminus A'$ with $|C| < \frac{2}{\epsilon}|B'|$, there is a VW-matching in $G_{A',B'}$ covering C . Let $C' = C \cap L(F)$ and $C'' = C \setminus C'$. By construction, F is a VW-matching such that $L(F) \subseteq A$, $R(F) \subseteq B$ and F covers C' . Moreover, we have that

$$|C''| \leq |C| < \frac{2}{\epsilon}|B'| < \frac{2}{\epsilon}|B| \stackrel{(*)}{\leq} s,$$

where the inequality $(*)$ is by hypothesis. Hence there exists a VW-matching F'' of C'' in $G_{A,B}$, and so $F \cup F''$ is a VW-matching covering C in $G_{A',B'}$. \square

Lemma 4.8 (Covering a vertex in L). Let $A \subseteq L$ and $B \subseteq R$ be such that the pair (A, B) has the VW-matching property and let d be the maximum degree of a vertex in $R \setminus B$. If $\frac{24d}{\epsilon}(|B| + 3) + 1 \leq s$, then for each vertex v in $L \setminus A$, there is a VW-matching F in $G_{A,B}$ covering v and such that $(A \cup L(F), B \cup R(F))$ has the VW-matching property.

Proof. Let $A' = A \setminus L(F)$ and $B' = B \setminus R(F)$ and suppose, by contradiction, that (A', B') does not have the VW-matching property. By Lemma 4.5, it is sufficient to prove that for each set $C \subseteq L \setminus A'$ with $|C| < \frac{2}{\epsilon}|B'|$, there is a VW-matching in $G_{A',B'}$ covering C . Let $C' = C \cap L(F)$ and $C'' = C \setminus C'$. By construction, F is a VW-matching such that $L(F) \subseteq A$, $R(F) \subseteq B$ and F covers C' . Moreover, we have that

$$|C''| \leq |C| < \frac{2}{\epsilon}|B'| < \frac{2}{\epsilon}|B| \stackrel{(*)}{\leq} s,$$

where the inequality $(*)$ is by hypothesis. Hence there exists a VW-matching F'' of C'' in $G_{A,B}$, and so $F \cup F''$ is a VW-matching covering C in $G_{A',B'}$. \square

Lemma 4.9 (Covering a vertex in R). Let $A \subseteq L$ and $B \subseteq R$ be such that the pair (A, B) has the VW-matching property and let d be the maximum degree of a vertex in $R \setminus B$. If $\frac{24d}{\epsilon}(|B| + 3d) + 1 \leq s$, then for each vertex v in $R \setminus B$, there is a VW-matching F in $G_{A,B}$ covering v and such that $(A \cup L(F), B \cup R(F))$ has the VW-matching property.

Proof. Fix $v \in R \setminus B$ and let Π be the set of all VW-matchings F in $G_{A,B}$, covering v and such that F is connected.

Since $1 \leq s$ and (A, B) has the VW-matching property, we know that Π is non-empty. For every $F \in \Pi$, let (A_F, B_F) be the pair $(A \cup L(F), B \cup R(F))$, and suppose for a contradiction that for every $F \in \Pi$, (A_F, B_F) does not have the VW-matching property. By Lemma 4.5, for every $F \in \Pi$ there is a set $C_F \subseteq L \setminus A_F$ with $|C_F| < \frac{2}{\epsilon}|B_F|$ and such that there is no VW-matching of C_F in G_{A_F, B_F} .

Let $C = \bigcup_{F \in \Pi} C_F$. Then

$$|C| \leq \sum_{F \in \Pi} |C_F| < |\Pi| \frac{2}{\epsilon}(|B| + 3) \leq 12d \frac{2}{\epsilon}(|B| + 3),$$

since $|\Pi| \leq 3 + 3 \cdot 2 \cdot (d - 1) \cdot 2 \leq 12d$ and $|B_F| \leq |B| + 3$. Hence, by our assumption about the size of $|B|$, we have that $|C \cup \{v\}| \leq s$. Furthermore, $C \cup \{v\} \subseteq L \setminus A$, so by the fact that (A, B) has the VW-matching property, there is a VW-matching F' covering $C \cup \{v\}$ in $G_{A,B}$.

There must be some $F \in \Pi$ such that F is a connected component of F' . Let F'' be F' with the component F removed. Then F'' is a VW-matching in G_{A_F, B_F} and F'' covers C_F , contradicting the choice of C_F . \square

We now have all the preliminary lemmas needed to prove Theorem 4.3.

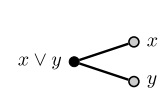
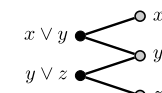
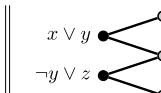
Proof of Theorem 4.3. By the hypothesis on $|S_d|$, for each $d \geq D$, we can repeatedly apply Lemma 4.9 starting from (\emptyset, \emptyset) to cover vertices in R of degree larger than D . By starting from vertices of R of maximum degree and proceeding in decreasing order until reaching the vertices of degree D , we can build a VW-matching M covering S_D such that $(L(M), R(M))$ has the VW-matching property. Moreover, by the choice of S_D , $G_{L(M), R(M)}$ (the subgraph induced by $(L \cup R) \setminus (L(M) \cup R(M))$) has maximum degree at most D . We say that a VW-matching F is compatible with M if each connected component of F is either a connected component of M or disjoint from all connected components of M .

We describe a winning strategy for Cover to win $\text{CoverGame}_{\text{VW}}(G, \mu)$. Take \mathcal{L} to be the set of all VW-matchings F in G compatible with M such that

1. $(L(M) \cup L(F), R(M) \cup R(F))$ has the VW-matching property, and
2. $\frac{2}{\epsilon}|R(M) \cup R(F)| \leq s$.

This family is non-empty since the empty VW-matching is in \mathcal{L} . Moreover, \mathcal{L} is closed under removing connected components by Lemma 4.7. Suppose now that at step $i + 1$ of the game Choose picks a vertex v in $G_{L(M), R(M)}$ and that F_i

Table 1
Flippable assignments from VW-matchings.

$\circ x$			
$x \mapsto 0$	$(x, y) \mapsto (0, 1)$	$(x, y, z) \mapsto (0, 1, 0)$	$(x, y, z) \mapsto (0, 1, 1)$
$x \mapsto 1$	$(x, y) \mapsto (1, 0)$	$(x, y, z) \mapsto (1, 0, 1)$	$(x, y, z) \mapsto (1, 0, 0)$

has strictly less than $\mu = \frac{\epsilon s}{144D}$ components. Then, $(L(M) \cup L(F_i), R(M) \cup R(F_i))$ satisfies the hypotheses of [Lemma 4.8](#) and [Lemma 4.9](#):

$$\begin{aligned}
 \frac{24D}{\epsilon} (|R(M) \cup R(F_i)| + 3D) + 1 &\leq \frac{24D}{\epsilon} (|R(M)| + |R(F_i)| + 3D) + 1 \\
 &\leq \frac{24D}{\epsilon} (|R(M)| + 3D) + 1 + \frac{24D}{\epsilon} |R(F_i)| \\
 &\stackrel{(\star)}{\leq} \frac{24D}{\epsilon} (3|S_D| + 3D) + 1 + \frac{72D}{\epsilon} \mu \\
 &\stackrel{(\star\star)}{\leq} \frac{s}{2} + \frac{72D}{\epsilon} \mu = \frac{s}{2} + \frac{72D}{\epsilon} \frac{\epsilon s}{144D} = s,
 \end{aligned}$$

where the inequality (\star) follows from the fact that $|R(F_i)| \leq 3\mu$ and $|R(M)| \leq 3|S_D|$, where S_D is the set of vertices in R of degree bigger than D . The inequality $(\star\star)$ follows by the hypothesis on the size of S_D .

Hence, if v is covered by F_i we take $F_{i+1} = F_i$. If v is covered by M we take $F_{i+1} = F_i \cup M_v$, where M_v is the connected component of M covering v . Otherwise, by [Lemma 4.8](#) and [Lemma 4.9](#) applied to $(L(M) \cup L(F_i), R(M) \cup R(F_i))$, there exists a VW-matching F_{i+1} extending $F_i \cup M$ by a new connected component covering v such that $(L(F_{i+1}), R(F_{i+1}))$ has the VW-matching property. From the previous chain of inequalities, it follows easily that the pair $(L(F_{i+1}), R(F_{i+1}))$ satisfies the cardinality condition $\frac{2}{\epsilon} |R(M) \cup R(F_{i+1})| = \frac{2}{\epsilon} |R(F_{i+1})| \leq s$. \square

5. Space lower bounds for random 3CNFs

Lemma 5.1. *Let φ be an unsatisfiable 3-CNF and G_φ its adjacency graph. If Cover wins the cover game $\text{CoverGame}_{\text{VW}}(G_\varphi, \mu)$, then there is a μ -winning strategy \mathcal{L} for $\text{tr}(\varphi)$.*

Proof. First of all we prove that for every VW-matching F in G_φ , there exists a flippable product-family of assignments H_F such that $H_F \models L(F)$, $\text{dom}(H_F) = R(F)$, and $\|H_F\|$ is the number of connected components of F .

We prove the result by induction on the number of connected components of F . If F is the union of two disjoint VW-matchings F', F'' then by hypothesis $H_{F'} \models L(F')$, $\text{dom}(H_{F'}) = R(F')$ and $\|H_{F'}\|$ is the number of connected components of F' . And analogously for F'' . Then, since $R(F')$ and $R(F'')$ are disjoint, $H_F = H_{F'} \otimes H_{F''}$ is well-defined. We immediately see that $H_F \models L(F)$, $\text{dom}(H_F) = R(F)$ and $\|H_F\|$ is the number of connected components of F .

It remains to consider the case when the VW-matching F is just one connected component. It is easy to see that all the possibilities can be reduced to those in [Table 1](#).

It is straightforward to check that a winning strategy for Cover in the game $\text{CoverGame}_{\text{VW}}(G_\varphi, \mu)$ defines, by previous observations, a family \mathcal{L} of flippable product-families such that for all $\mathcal{H} \in \mathcal{L}$

1. for each $\mathcal{H}' \sqsubseteq \mathcal{H}$, $\mathcal{H}' \in \mathcal{L}$;
2. if $\|\mathcal{H}\| < \mu$, then: (a) for each $C \in \varphi$, there exists a flippable product-family $\mathcal{H}' \in \mathcal{L}$ such that $\mathcal{H}' \models C$ and $\mathcal{H}' \sqsupseteq \mathcal{H}$; and (b) for each variable $x \notin \text{dom}(\mathcal{H})$, there exists a flippable family $\mathcal{H}' \in \mathcal{L}$ such that $\mathcal{H}' \sqsupseteq \mathcal{H}$ and $x \in \text{dom}(\mathcal{H}')$.

We claim that \mathcal{L} is a μ -winning strategy. The *restriction property* is immediate. For the *extension property* we use the properties in (2) above: if we have to extend to something in \mathcal{L} that satisfies a boolean axiom we use property 2.(b), otherwise for all other polynomials in $\text{tr}(\varphi)$ we use property 2.(a). \square

Let $n, \Delta \in \mathbb{N}$ and let $X = \{x_1, \dots, x_n\}$ be a set of n variables. The probability distribution $\mathcal{R}(n, \Delta, 3)$ is obtained by the following experiment: choose independently uniformly at random Δn clauses from the set of all possible clauses with 3 literals over X . It is well-known that when Δ exceeds a certain constant θ_3 , φ is almost surely unsatisfiable, see for example [\[22\]](#). Hence we always consider $\varphi \sim \mathcal{R}(n, \Delta, 3)$, where Δ is a constant bigger than θ_3 , which implies that φ is unsatisfiable with high probability.

Lemma 5.2. *Let $\Delta > \theta_3$ and $\varphi \sim \mathcal{R}(n, \Delta, 3)$ a random 3-CNF. For every integer d let S_d be the set of variables of φ appearing in at least d clauses of φ . Then for every constant $c > 0$ and $\epsilon > 0$, with high probability there exists a constant D such that for every $d \geq D$,*

$$\frac{72d}{\epsilon} (|S_d| + d) + 1 \leq cn.$$

Proof. Let G_φ be the adjacency graph of φ . First of all we show that w.h.p. there are at most $\frac{en}{2^d}$ many variable nodes of degree d for every $d \geq 24e\Delta$ and that w.h.p. there is no variable node of degree bigger than $\log n$. First note that the expected number of variable nodes of degree at least $\log n$ is

$$n \binom{\Delta n}{\log n} \left(\frac{3}{n-2}\right)^{\log n} \leq n \left(\frac{e\Delta n}{\log n}\right)^{\log n} \left(\frac{3}{n-2}\right)^{\log n} = o(1).$$

So w.h.p. there are no such nodes. Let $d \geq 24e\Delta$. The probability that there are $\frac{en}{2^d}$ many variable nodes of degree d is at most

$$\begin{aligned} \binom{n}{\frac{en}{2^d}} \left[\binom{\Delta n}{d} \left(\frac{3}{n-2}\right)^d \right]^{\frac{en}{2^d}} &\leq \left(\frac{en}{2^d}\right)^{\frac{en}{2^d}} \left[\left(\frac{e\Delta n}{d}\right)^d \left(\frac{3}{n-2}\right)^d \right]^{\frac{en}{2^d}} \\ &\leq \left(\frac{12e\Delta}{d}\right)^{\frac{edn}{2^d}} \leq \left(\frac{1}{2}\right)^{\frac{edn}{2^d}}, \end{aligned}$$

so, by the union bound, the probability that there exists any d between $24e\Delta$ and $\log n$ such that there are $\frac{en}{2^d}$ many variable nodes of degree d is at most

$$\sum_{24e\Delta \leq d \leq \log n} \left(\frac{1}{2}\right)^{\frac{edn}{2^d}}.$$

To bound this sum, note that the ratio of consecutive terms is

$$2 \frac{edn}{2^d} - \frac{e(d+1)n}{2^{d+1}} = 2 \frac{e(d-1)n}{2^{d+1}} \geq 2$$

for d in this range, and so the sum is of the order of its last term, which is $\left(\frac{1}{2}\right)^{\frac{en \log n}{2^{\log n}}} = o(1)$.

So we have that w.h.p.

$$|S_d| \leq \sum_{d' \geq d} \frac{en}{2^{d'}} \leq \frac{2en}{2^d}$$

and, for (a not yet chosen constant) $D \geq 24e\Delta$, we have that for each d such that $D \leq d \leq \log n$:

$$\frac{72d}{\epsilon} (|S_d| + d) + 1 \leq \frac{72d}{\epsilon} \left(\frac{2en}{2^d} + d\right) + 1 \leq \frac{72D}{\epsilon} \frac{2en}{2^D} + O(\log^2 n) \leq cn,$$

where the last inequality holds if D is a sufficiently large constant. \square

Theorem 5.3. *If $\Delta > \theta_3$ and $\varphi \sim \mathcal{R}(n, \Delta, 3)$, then the following statements hold with high probability. For every semantical PCR refutation Π of $\text{tr}(\varphi)$, $\text{MSpace}(\Pi) \geq \Omega(n)$. Moreover, every resolution refutation of φ must pass through a memory configuration containing $\Omega(n)$ clauses each of width $\Omega(n)$. In particular, each refutation of φ requires total space $\Omega(n^2)$.*

Proof. Let G_φ be the adjacency graph of φ . It is well known that for every $\delta > 0$ there exists a γ such that G_φ is a $(\gamma n, 2 - \delta)$ -bipartite expander [12–14,16]. Hence in particular for $0 < \delta < \frac{1}{5}$ and using Lemma 5.2 with $c = \frac{\gamma}{2}$, we satisfy all the hypotheses of Theorem 4.3. Thus, Cover wins the cover game $\text{CoverGame}_{\text{vw}}(G_\varphi, \mu)$ for $\mu = \Omega(n)$. Lemma 5.1 provides a $\Omega(n)$ -winning strategy and by Theorem 2.2 we have the monomial space lower bound in semantical PCR and the total space lower bound in resolution. \square

Acknowledgments

We thank Susanna Figueiredo De Rezende for allowing us to include her nice simplification of the original proof of Lemma 1.2. We also thank Jakob Nordström and Massimo Lauria for interesting discussions on space complexity in propositional calculus.

References

- [1] Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, Avi Wigderson, Space complexity in propositional calculus, *SIAM J. Comput.* 31 (4) (2002) 1184–1211.
- [2] Albert Atserias, Víctor Dalmau, A combinatorial characterization of resolution width, *J. Comput. Syst. Sci.* 74 (3) (2008) 323–334.
- [3] Albert Atserias, On sufficient conditions for unsatisfiability of random formulas, *J. ACM* 51 (2) (2004) 281–311.
- [4] Patrick Bennett, Ilario Bonacina, Nicola Galesi, Tony Huynh, Mike Molloy, Paul Wollan, Space proof complexity for random 3-cnfs, March 2015.
- [5] Eli Ben-Sasson, Size space tradeoffs for resolution, in: John H. Reif (Ed.), *Proceedings on 34th Annual ACM Symposium on Theory of Computing*, Montréal, Québec, Canada, May 19–21, 2002, pp. 457–464.
- [6] Ilario Bonacina, Nicola Galesi, A framework for space complexity in algebraic proof systems, *J. ACM* 62 (3) (June 2015) 23:1–23:20.
- [7] Ilario Bonacina, Nicola Galesi, Neil Thapen, Total space in resolution, in: *55th Annual IEEE Symposium on Foundations of Computer Science, FOCS, 2014*, pp. 641–650.
- [8] Archie Blake, *Canonical Expressions in Boolean Algebra*, PhD thesis, University of Chicago, 1937.
- [9] Eli Ben-Sasson, Jakob Nordström, Short proofs may be spacious: an optimal separation of space and length in resolution, in: *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008*, October 25–28, 2008, IEEE Computer Society, Philadelphia, PA, USA, 2008, pp. 709–718.
- [10] Eli Ben-Sasson, Jakob Nordström, Understanding space in proof complexity: separations and trade-offs via substitutions, in: Bernard Chazelle (Ed.), *Innovations in Computer Science Proceedings, ICS 2010*, Tsinghua University, Beijing, China, January 7–9, 2011, Tsinghua University Press, 2011, pp. 401–416.
- [11] Ilario Bonacina, Total space in resolution is at least width squared, in: *43rd International Colloquium on Automata, Languages, and Programming, ICALP, 2016*, pp. 56:1–56:13.
- [12] Paul Beame, Toniann Pitassi, Simplified and improved resolution lower bounds, in: *FOCS, IEEE Computer Society, 1996*, pp. 274–282.
- [13] Eli Ben-Sasson, Nicola Galesi, Space complexity of random formulae in resolution, *Random Struct. Algorithms* 23 (1) (2003) 92–109.
- [14] Eli Ben-Sasson, Avi Wigderson, Short proofs are narrow – resolution made simple, *J. ACM* 48 (2) (2001) 149–169.
- [15] Matthew Clegg, Jeff Edmonds, Russell Impagliazzo, Using the Groebner basis algorithm to find proofs of unsatisfiability, in: Gary L. Miller (Ed.), *STOC, ACM, 1996*, pp. 174–183.
- [16] Chvátal Vasek, Endre Szemerédi, Many hard examples for resolution, *J. ACM* 35 (4) (1988) 759–768.
- [17] Daniel W. Cranston, Douglas B. West, A guide to the discharging method, arXiv preprint, arXiv:1306.4434, 2013.
- [18] Susanna F. de Rezende, *Personal communication*, 2017.
- [19] Juan Luis Esteban, Jacobo Torán, Space bounds for resolution, *Inf. Comput.* 171 (1) (2001) 84–97.
- [20] Yuval Filmus, Massimo Lauria, Mladen Mikša, Jakob Nordström, Marc Vinyals, Towards an understanding of polynomial calculus: new separations and lower bounds (extended abstract), in: Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, David Peleg (Eds.), *ICALP*, in: *Lecture Notes in Computer Science*, vol. 7965, Springer, 2013, pp. 437–448.
- [21] Yuval Filmus, Massimo Lauria, Jakob Nordström, Neil Thapen, Noga Ron-Zewi, Space complexity in polynomial calculus, in: *Proceedings of the 27th Conference on Computational Complexity, CCC 2012*, Porto, Portugal, June 26–29, 2012, IEEE, 2012, pp. 334–344.
- [22] Ehud Friedgut, Sharp thresholds of graph properties, and the k-sat problem, *J. Am. Math. Soc.* 12 (1998) 1017–1054.
- [23] P. Hall, On representatives of subsets, *J. Lond. Math. Soc.* s1-10 (1) (1935) 26–30.
- [24] Jakob Nordström, Narrow proofs may be spacious: separating space and width in resolution, *SIAM J. Comput.* 39 (1) (2009) 59–121.
- [25] Jakob Nordström, On the interplay between proof complexity and SAT solving, *SIGLOG News* 2 (3) (2015) 19–44.
- [26] J.A. Robinson, A machine-oriented logic based on the resolution principle, *J. ACM* 12 (1) (January 1965) 23–41.
- [27] Alexander Roberts, Tree matchings, arXiv preprint, arXiv:1612.01694, 2016.