SIAM J. COMPUT.
Vol. 45, No. 5, pp. 1894–1909

TOTAL SPACE IN RESOLUTION*

ILARIO BONACINA†, NICOLA GALESI‡, AND NEIL THAPEN§

Abstract. We show quadratic lower bounds on the total space used in resolution refutations of random k-CNFs over n variables and of the graph pigeonhole principle and the bit pigeonhole principle for n holes. This answers the open problem of whether there are families of k-CNF formulas of polynomial size that require quadratic total space in resolution. The results follow from a more general theorem showing that, for formulas satisfying certain conditions, in every resolution refutation there is a memory configuration containing many clauses of large width.

Key words. total space, resolution, random CNFs, proof complexity

AMS subject classification. 03F20

DOI. 10.1137/15M1023269

1. Introduction.

The questions most frequently asked in propositional proof complexity concern the *size* of proofs—as is well-known, $\mathsf{NP} = \mathsf{coNP}$ if and only if there is a proof system in which every tautology has a polynomial size proof [16]. There is a natural analogy between the size of a proof and the size of a circuit, or the time taken by a Turing machine. Developing this analogy, [15, 18, 1] introduced a notion of the *space* used by a propositional proof, similar to the notion of space for Turing machines. Since then, space has been investigated in depth in proof complexity, especially for the resolution proof system [18, 1, 4, 2] and in particular concerning trade-offs [5, 21, 24, 7, 22], resolution over k-DNFs [17, 6], and more recently for polynomial calculus [1, 20, 13, 19].

Resolution is a well-studied system for refuting formulas in conjunctive normal form (CNFs). Each line in a resolution refutation is a *clause*, that is, a disjunction of literals, and resolution has only one rule: from two clauses $A \vee x$ and $B \vee \neg x$ we may infer the *resolvent* clause $A \vee B$. A CNF is unsatisfiable if and only if the empty clause can be derived from it using this rule.

Intuitively, the space required by a refutation is the amount of information we need to keep simultaneously in memory as we work through the proof and convince ourselves that the original CNF is unsatisfiable. This was made formal for resolution in [18] as follows. A *memory configuration*, or just *configuration*, is a set of clauses. We assume that a resolution refutation of φ is given in the form of a sequence $M_1, \ldots, M_t$

*Received by the editors May 27, 2015; accepted for publication (in revised form) August 15, 2016; published electronically October 20, 2016. A preliminary version of this paper appeared in *Proceedings of the* 55*th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014.
http://www.siam.org/journals/sicomp/45-5/M102326.html
Funding: Part of this work was done when the first two authors were visiting the Institute of Mathematics of the ASCR, partially supported by grant P202/12/G061 of GAČR. The first author was funded by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement 279611. The third author received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant agreement 339691. The Institute of Mathematics of the Academy of Sciences of the Czech Republic is supported by RVO:67985840.

†School of Computer Science and Communication, KTH Royal Institute of Technology, Stockholm, Sweden (ilario@kth.se).
‡Computer Science Department, Sapienza University of Rome, Rome, Italy (galesi@di.uniroma1.it).
§Institute of Mathematics, Academy of Sciences of the Czech Republic, Prague, Czech Republic (thapen@math.cas.cz).

of configurations, where M_1 is empty, M_t contains the empty clause, and each M_{i+1} is derived from M_i in one of the following three ways:

(AXIOM DOWNLOAD) $M_{i+1} = M_i \cup \{C\}$, where C is a clause from φ.

(ERASURE) $M_{i+1} = M_i \setminus \{D\}$, where D is a clause from M_i.

(INFERENCE) $M_{i+1} = M_i \cup \{D\}$, where D is the resolvent of two clauses in M_i.

This model is 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.

Following [18, 1] the *clause space* used by the refutation is the maximum number of clauses in any configuration M_i in the sequence. The *total space* used is the maximum over i of the total number of symbols needed to write down M_i. In other words, it is the total number of instances of variables[1] occurring in M_i (we ignore punctuation, logical connectives, and the actual labels of the variables, which take $O(\log n)$ bits to write).

Clause space and its relation with proof size are by now well-studied [6, 5, 7, 22, 2]. But much less is known about total space, although it captures more closely the intuitive idea of the memory required by a refutation. As well as being of theoretical interest, total space is also potentially a useful measure for SAT solving. Memory use is a major problem for SAT solvers and a current goal of research is to understand the resources of time and space in resolution proofs, how they are connected to each other, and how they can be optimized in the design of new SAT solvers; see, for example, the recent survey [23].

1.1. Results. Every unsatisfiable CNF φ over n variables can be refuted in resolution in clause space $n + 1$, which is the pebbling number of the brute-force treelike resolution refutation of φ [18]. Since every clause in the refutation has width at most n, this gives an upper bound of $n(n + 1)$ on the total space of refuting φ (where the *width* of a clause is the number of literals in it).

The only previously known lower bounds for total space, other than the linear bounds following trivially from lower bounds on width or clause space, are from [1]. There it is shown that the *complete tree contradiction* CT_n requires $\Omega(n^2)$ total space to refute. CT_n is a CNF formula of exponential size consisting of all 2^n possible clauses of width n over the variables $x_1, \ldots, x_n$. This is the only explicit CNF we are aware of which was previously known to require superlinear total space (in the number of variables). A lower bound of $\Omega(n^2)$ on the total space to refute the *pigeonhole principle* PHP_n also follows from [1], but this lower bound is linear in terms of the number of variables of the CNF, as PHP_n has $\Theta(n^2)$ variables.

Improving these results, by finding a polynomial size CNF requiring at least superlinear total space in the number of variables, is a problem posed in [1] which has since appeared in many other works in proof complexity [4, 21, 6, 5, 7, 22]. We are able to solve it in essentially an optimal way, showing that some standard families of constant-width CNF contradictions, defined over n variables and of size $O(n)$, require $\Omega(n^2)$ total space. More precisely, we prove in each case that every refutation of the formula in question must pass through a configuration containing r clauses each of width at least r, where $r = \Omega(n)$. Our main result is the following theorem.

THEOREM 1.1. *Let $k \geq 4$ and $\Delta > 1$. There is a constant $c > 0$ such that, for large n, if φ is a random k-CNF with n variables and Δn clauses, then with high*

[1] In [1] this is called *variable space*, but we follow [6, 5, 21, 7, 22] in calling it *total space* to distinguish it from a different measure in which different occurrences of the same variable are not counted.

probability any resolution refutation of φ passes through a configuration containing cn clauses of width at least cn.

We show similar lower bounds for some other CNFs. In particular, for the graph pigeonhole principle $\mathcal{G}$-PHP (see the beginning of section 6 for definitions) and the bit pigeonhole principle BPHP_n (see the beginning of section 3) we show the following.

THEOREM 1.2. *Let $d \geq 4$ and $\Delta > 1$. There is a constant $c > 0$ such that, for large n, if $\mathcal{G}$ is a random graph chosen from the set of bipartite graphs with left-degree d with Δn pigeon nodes on the left and n hole nodes on the right, then with high probability any resolution refutation of $\mathcal{G}$-PHP passes through a configuration containing cn clauses of width at least cn.*

THEOREM 1.3. *Let $n = 2^k$ for k an integer. Any resolution refutation of BPHP_n passes through a configuration containing $n/4$ clauses of width at least $n/4$.*

The random formulas and the instances of $\mathcal{G}$-PHP in Theorems 1.1 and 1.2 are k-CNFs with $O(n)$ variables, so in both cases our lower bound matches the quadratic upper bound on total space, up to a constant factor. The bit pigeonhole principle BPHP_n is a $(\log n)$-CNF with $(n+1)\log n$ variables, so our lower bound is only $\Omega(m^2/(\log m)^2)$ in terms of the number m of variables (but the proof is much simpler than for the other two principles).

1.2. Outline of paper. The next section contains a general theorem (Theorem 2.4) from which our results follow. We define the notion of an *r-free family* of assignments and show that if a CNF has such a family, then every resolution refutation of it has a configuration containing $r/2$ clauses each of width at least $r/2$.

In section 3 we give two applications to illustrate the use of Theorem 2.4. One is the total space lower bound for BPHP_n (Theorem 1.3). The other is the observation that from any constant-width CNF φ requiring width w to refute, we can construct a constant-width CNF $\varphi[\oplus]$, the "xorification of φ," which requires $\Omega(w^2)$ total space to refute (Theorem 3.1). In particular, this gives us a lower bound for certain Tseitin formulas.

Section 4 is the only really technical part of the paper. We develop the tools we will need to construct r-free families of assignments for random k-CNFs and $\mathcal{G}$-PHP, namely, certain families of substructures of bipartite graphs which we call *r-covering families*. We show that in a random bipartite graph such a family exists with high probability, that is, with probability $1 - o(1)$.

In sections 5 and 6 we use this to prove our total space lower bounds respectively for random k-CNFs and $\mathcal{G}$-PHP.

In section 7 we discuss *semantic resolution* [1]. We show that resolution can require much more total space than semantic resolution. We prove that if a CNF has an r-free family, then it requires large total space in a weak version of semantic resolution, in which we can derive a new clause if it is implied by some set of d clauses in memory, where d is fixed (Theorem 7.1). We prove that every r-semiwide CNF requires large semantic total space (Theorem 7.3—see Definition 7.2 for the definition of r-semiwide formulas).

The most important parts of the paper are the definitions and main theorem in section 2 and the application of this to give lower bounds for random k-CNFs in section 5, building on technical results about bipartite expanders in section 4. The result about BPHP_n in Theorem 1.3 (which is already a big improvement over previously known lower bounds) provides an example of a total space lower bound that can be read without needing all the technicalities required for random k-CNFs.

Many of our constructions are inspired by recent work on lower bounds on monomial space (analogous to clause space) in the system PCR of polynomial calculus resolution [13]. In particular, the partial assignments defining r-free families come with some extra structure that means that they are not closed under taking subassignments, as would usually be the case with similar constructions (as used, for example, for clause space lower bounds). The definition of *piecewise assignment* is a simplification of an *admissible configuration* from [13]. The definition of an r-free family is new, and a crucial innovation is that we use the r-free family to explicitly pick out a nicely behaved substructure of the resolution refutation and focus on showing a total space lower bound on this substructure.

In the applications of our main theorem, the idea of an r-covering family and its use with random k-CNFs and $\mathcal{G}$-PHP extends a construction from [13]. The use of BPHP_n is inspired by its use in [20] and the observation about xorifications is modeled on an analogous observation in [19].

1.3. Recent developments. During the preparation of the journal version of this work, Theorem 1.1 was shown to hold also for $k = 3$, that is, for random 3-CNFs [9]. In the first author's Ph.D. thesis [11] it was recently proven that given an unsatisfiable k-CNF formula φ, if W and T are respectively the minimal width and the minimal total space needed to refute φ, then $T \geq \Omega((W-k)^2)$.

The improvement in [9] relies on the constructions we build here and in particular on Theorem 2.4 as it appeared in the conference version of this work [14]. The result in [11] and [12] relies on simplifications of Definition 2.3 and Theorem 2.4 and on a characterization of *asymmetric* width via families of assignments[2] (Definition 21 and Theorem 22 in [10]).

1.4. Open problems. A natural question is whether these lower bounds can be extended to stronger proof systems such as bounded depth Frege, where very little is known about space, or PCR. For unrestricted Frege systems a linear upper bound (in the size of the CNF being refuted) on total space was shown in [1].

Finally, all of our lower bounds are for formulas which are already known to be hard for resolution, in that they have no subexponential size refutations. It is open whether there is a family of CNFs which have short refutations but which still require quadratic, or at least superlinear, total space. By a result of [8], if a CNF has a resolution refutation of size S, then it also has a refutation in which every clause has width at most $O(\sqrt{n \log S})$. Hence we cannot hope to use our arguments, which show large space by finding many clauses of large width.

2. Main theorem. A *partial assignment to* X has the usual meaning of an assignment of 0/1 values to a subset D of X, leaving the rest of the variables in X unassigned. The *domain* of the partial assignment is the set D.

DEFINITION 2.1. *A* piecewise assignment α *to a set of variables* X *is a set of nonempty partial assignments to* X*, with pairwise disjoint domains.*

A piecewise assignment α to X naturally gives rise to a partial assignment to X, namely $\bigcup \alpha$, the union of all the partial assignments in α. It also gives rise to a partition of the domain of $\bigcup \alpha$, into the set of domains of all the members of α. Therefore an alternative, but notationally less convenient, way to define a piecewise assignment would be as such a pair of a partial assignment and a partition of its

[2]The family of assignments used in [10] to characterize asymmetric width share some properties with our Definition 2.3 but the two concepts were introduced independently.

domain, and we will often write α when our intended meaning is the partial assignment $\bigcup \alpha$. For example, we will write $\alpha(\varphi)$ for the evaluation of φ under $\bigcup \alpha$ and $\mathrm{dom}(\alpha)$ for the domain of $\bigcup \alpha$.

We call the elements of α the *pieces* of α. 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 "$\sqsubseteq$" and "$\|\cdot\|$" are formally exactly the same as $\alpha \subseteq \beta$ and $|\alpha|$, using the definition of α and β as sets of partial assignments. We will use the following simple fact about piecewise assignments, which we record as a lemma.

LEMMA 2.2. *Let α, β be piecewise assignments with $\alpha \sqsubseteq \beta$. Let $Y \subseteq \mathrm{dom}(\beta)$. Then there exists a piecewise assignment β' with $\alpha \sqsubseteq \beta' \sqsubseteq \beta$ such that $Y \subseteq \mathrm{dom}(\beta')$ and $\|\beta'\| \leq \|\alpha\| + |Y|$.*

DEFINITION 2.3. *A nonempty family $\mathcal{H}$ of piecewise assignments is* r-free for a CNF φ *if it has the following properties:*

(CONSISTENCY) *No $\alpha \in \mathcal{H}$ falsifies any clause from φ.*

(DOWNWARD CLOSURE) *If $\alpha \in \mathcal{H}$, β is a piecewise assignment, and $\beta \sqsubseteq \alpha$, then $\beta \in \mathcal{H}$.*

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

THEOREM 2.4. *Let φ be a CNF formula and let $r \geq 2$ be an integer. If there is a family of piecewise assignments which is r-free for φ, then any resolution refutation of φ must pass through a memory configuration containing at least $r/2$ clauses each of width at least $r/2$. In particular, the refutation requires total space at least $r^2/4$.*

Proof. Suppose that φ is unsatisfiable and that $\mathcal{H}$ is a family of piecewise assignments which is r-free for φ. Let $\Pi = (M_1, \ldots, M_s)$ be a resolution refutation of φ, given as a sequence of memory configurations.

Let S be the set of all clauses which are falsified by some member of $\mathcal{H}$. There is at least one clause in $\bigcup_{i=1}^{s}(M_i \cap S)$ with width strictly less than $r/2$, namely, the empty clause. Let M_t be the first configuration in Π in which a clause of width strictly less than $r/2$ occurs in $M_t \cap S$ and let C be such a clause. Let $\alpha \in \mathcal{H}$ falsify C. By Lemma 2.2 and the downward closure of $\mathcal{H}$ we may assume that $\|\alpha\| < r/2$. Our goal now is to show that there is some $i < t$ such that $|M_i \cap S| \geq r/2$. Since, by our choice of t, for every $i < t$ every clause in $M_i \cap S$ has width at least $r/2$, this will give the theorem.

Suppose for a contradiction that $|M_i \cap S| < r/2$ for each $i < t$. We will inductively construct a sequence of piecewise assignments $\beta_1, \ldots, \beta_t$ in $\mathcal{H}$ such that for each $i \leq t$ we have that $\alpha \sqsubseteq \beta_i$ and that β_i satisfies every clause in $M_i \cap S$. This will give a contradiction when we reach β_t, since α falsifies the clause $C \in M_t \cap S$.

The first configuration M_1 is empty, so we can put $\beta_1 = \alpha$. Supposing that $1 \leq i < t$ and that we already have a suitable β_i, we distinguish three cases:

(AXIOM DOWNLOAD) $M_{i+1} = M_i \cup \{D\}$, where D is a clause from φ. By the consistency property of $\mathcal{H}$, D is not in S and we can simply put $\beta_{i+1} = \beta_i$.

(ERASURE) $M_{i+1} = M_i \setminus \{D\}$, where D is a clause from M_i. We put $\beta_{i+1} = \beta_i$.

(INFERENCE) $M_{i+1} = M_i \cup \{D \vee E\}$, where $D \vee E$ follows by resolution on some variable x from two clauses $D \vee x$ and $E \vee \neg x$ in M_i. Using Lemma 2.2, since we have $\|\alpha\| < r/2$ and $|M_i \cap S| < r/2$ we may assume that $\|\beta_i\| \leq \|\alpha\| + |M_i \cap S| < r$.

If $D \vee E$ contains a variable outside $\mathrm{dom}(\beta_i)$, then by the extension property of $\mathcal{H}$ we can extend β_i to some $\beta_{i+1} \in \mathcal{H}$ which satisfies $D \vee E$, as required.

Suppose that all variables in $D \vee E$ are set by β_i. If $x \in \mathrm{dom}(\beta_i)$ let $\beta_{i+1} = \beta_i$, and otherwise let $\beta_{i+1} \in \mathcal{H}$ be any extension of β_i which assigns a value to x. Then β_{i+1} sets all variables in both $D \vee x$ and $E \vee \neg x$. It cannot falsify either clause, since that would imply that that clause is in S and thus is already satisfied by β_i. Therefore it must satisfy both clauses and thus also satisfy $D \vee E$. ∎

Informally, we can think of each element C of S as identified with a minimal assignment α_C in $\mathcal{H}$ which falsifies it. Then S contains the empty assignment and, by the extension property of $\mathcal{H}$, has a rich structure. In particular, if a clause C in $M_i \cap S$ for some i has width less than r and was derived by resolution on a variable outside $\mathrm{dom}(\alpha_C)$, then *both* parents of C in Π are in S. The proof of Theorem 2.4 then uses an idea from [1], taking the first clause C in S with small width and applying the usual clause space lower-bound argument to the substructure of S which derives C.

3. Two simple applications. Let $n = 2^k$ for $k \in \mathbb{N}$. The formula BPHP_n, the *bit pigeonhole principle on n holes*, is an unsatisfiable CNF with variables $\{x_j^u : u \in [n+1], j \in [k]\}$. It asserts that for all distinct $u, v \in [n+1]$, the length-k binary strings $x_1^u \dots x_k^u$ and $x_1^v \dots x_k^v$ are distinct. We think of each element of $[n+1]$ as a pigeon and of the string $x_1^u \dots x_k^u$ as the address, in binary, of the hole in $[n]$ that pigeon u is mapped to. Understood in this way, BPHP_n asserts that there is an injective mapping of $n+1$ pigeons into n holes. Formally the principle consists of the clauses

$$\bigvee_{j=1}^{k} (x_j^u \neq h_j) \vee \bigvee_{j=1}^{k} (x_j^v \neq h_j)$$

for each $u, v \in [n+1]$ with $u < v$ and each binary string $h_1 \dots h_k \in \{0,1\}^k$. We recall that the notation $x_j^u \neq 0$ stands for x_j^u and $x_j^u \neq 1$ stands for $\neg x_j^u$.

Restated Theorem 1.3. *Let $n = 2^k$ for k an integer. Any resolution refutation of BPHP_n passes through a configuration containing $n/4$ clauses of width at least $n/4$.*

Proof. By Theorem 2.4 it is enough to exhibit a family of piecewise assignments which is $n/2$-free.

For any partial matching f of pigeons into holes, let α_f be the piecewise assignment that, for each pigeon u in $\mathrm{dom}(f)$, assigns to the variables $x_1^u \dots x_k^u$ the binary string corresponding to the hole $f(u)$. The pieces of α_f correspond to the sets of variables $\{x_1^u, \dots, x_k^u\}$ belonging to each pigeon. Let $\mathcal{H}$ be the family of all piecewise assignments arising in this way for all partial matchings f.

Clearly $\mathcal{H}$ is nonempty and satisfies consistency and downward closure. For the extension property, suppose we are given $\alpha_f \in \mathcal{H}$ and a variable x_j^u, with $\|\alpha_f\| < n/2$ and $x_j^u \notin \mathrm{dom}(\alpha_f)$. Then $|\mathrm{ran}(f)| < n/2 = 2^{k-1}$ and $u \notin \mathrm{dom}(f)$, and it is sufficient to find two holes $h_1 \dots h_k$ and $h'_1 \dots h'_k$ in $\{0,1\}^k \setminus \mathrm{ran}(f)$ with $h_j = 0$ and $h'_j = 1$. But there are exactly 2^{k-1} holes h with $h_j = 0$, so there must be at least one such hole outside $\mathrm{ran}(f)$. A similar argument works for h'. ∎

As a second application, we show that a CNF requiring large total space in resolution can be constructed from any CNF which requires large width. This is modeled on a similar result in [19] for monomial space in PCR.

Let φ be a CNF over a set of variables X. Let X' be a new set of variables containing a disjoint pair $\{x^1, x^2\}$ of variables for each $x \in X$. Following [19], for each clause C in φ, let $C[\oplus]$ be the formula over X' obtained by replacing each occurrence of x_i in C with the expression $(x_i^1 \oplus x_i^2)$ and then converting the result back into CNF. Let $\varphi[\oplus]$ be the conjunction of all the CNFs $C[\oplus]$.

The *width* of a resolution refutation is the maximum width of any clause in it. The *refutation width* of a CNF φ in resolution is the minimal width of any refutation of φ.

THEOREM 3.1. *Let φ be a CNF and let w the minimal refutation width of φ in resolution. Then any resolution refutation of $\varphi[\oplus]$ passes through a configuration containing $w/2$ clauses of width at least $w/2$.*

Proof. Using the characterization of width in resolution by Atserias and Dalmau [2], we know that there is a w-winning strategy for the duplicator in the spoiler-duplicator game on φ. That is, there is a nonempty family $\mathcal{K}$ of partial truth assignments such that

1. if $f \in \mathcal{K}$, then f does not falsify any clause from φ;
2. if $f \in \mathcal{K}$ and $g \subseteq f$, then $g \in \mathcal{F}$;
3. if $f \in \mathcal{K}$, $|\operatorname{dom}(f)| < w$, and x is any variable, then there is some $g \in \mathcal{K}$ such that $f \subseteq g$ and $x \in \operatorname{dom}(g)$.

We will use $\mathcal{K}$ to build a w-free family $\mathcal{H}$ of piecewise assignments for $\varphi[\oplus]$. The result then follows by our main theorem.

Consider an assignment $f \in \mathcal{K}$. For each variable $x \in \operatorname{dom}(f)$, let α_x^0 be the partial assignment mapping $(x^1, x^2) \mapsto (0, f(x))$ and let α_x^1 be the partial assignment $(x^1, x^2) \mapsto (1, f(x) \oplus 1)$, so that for $b = 0, 1$ we have $\alpha_x^b(x^1) \oplus \alpha_x^b(x^2) = f(x)$ and for $i = 1, 2$ at least one of the partial assignments α_x^0, α_x^1 sets x^i to 0 and at least one sets x^i to 1. For any map $\delta : \operatorname{dom}(f) \to \{0, 1\}$ let α_f^δ be the piecewise assignment $\{\alpha_x^{\delta(x)} : x \in \operatorname{dom}(f)\}$. Notice that for each clause C in φ, α_f^δ falsifies $C[\oplus]$ if and only if f falsifies C.

Let $\mathcal{H}$ contain the piecewise assignment α_f^δ for each $f \in \mathcal{K}$ and each possible map $\delta : \operatorname{dom}(f) \to \{0, 1\}$. Consistency and downward closure for $\mathcal{H}$ follow from properties 1 and 2 of $\mathcal{K}$. For the extension property, suppose $\alpha \in \mathcal{H}$ and x^i is a variable in X' such that $\|\alpha\| < r$ and $x^i \notin \operatorname{dom}(\alpha)$. Then α must arise from some $f \in K$, with $|f| < r$ and $x \notin \operatorname{dom}(f)$. By property 3 of $\mathcal{K}$, there is an extension $g \supseteq f$ in $\mathcal{K}$ with $x \in \operatorname{dom}(g)$. By the construction of $\mathcal{H}$ there exist piecewise assignments β_0 and β_1 arising from g and extending α such that $\beta_0(x^i) = 0$ and $\beta_1(x^i) = 1$. ∎

In particular this result is interesting when φ is a Tseitin formula over some graph G. In this case $\varphi[\oplus]$ can be seen as a Tseitin formula over the graph G' formed by replacing each edge in G with a double edge.

We recall briefly what a *Tseitin formula* is. Let $G = (V, E)$ be a connected graph of degree d over n vertices. For each edge $e \in E$ define a variable x_e. Fix an *odd-weight* function $\sigma : V \to \{0, 1\}$, that is, a function σ such that $\sum_{v \in V} \sigma(v) \equiv 1 \pmod 2$. For each $v \in V$ define PARITY_v as a CNF expressing

$$\sum_{e \ni v} x_e \equiv \sigma(v) \pmod 2.$$

The Tseitin formula $T(G, \sigma)$ is then the conjunction $\bigwedge_{v \in V} \text{PARITY}_v$. It is well-known that refutation width of $T(G, \sigma)$ is at least the connectivity expansion of G (see for example [1]).

COROLLARY 3.2. *Let $G = (V, E)$ be a 3-regular expander graph over n vertices. Let G' be G with each edge replaced with a double edge. Then for any odd weight function $\sigma : V \to \{0, 1\}$ the total space needed to refute $T(G', \sigma)$ is at least $\Omega(n^2)$.*

Here $T(G', \sigma)$ is a 6-CNF. This corollary is a partial answer to the question posed in open problem 2 of [1] about the space needed to refute $T(G, \sigma)$ when G is a 3-regular expander graph.

4. Bipartite expanders and 2-matchings. The goal of this section is to define certain families of substructures of bipartite graphs, which we call *r-covering families*, and to show that in a random bipartite graph such a family exists with high probability. See Definitions 4.10 and 4.11 and Corollary 4.14 at the end of the section. We will need such families in our lower bounds for random formulas and for the graph pigeonhole principle. The constructions in this section are adapted from [13], which in turn is based on [4]. Our main innovation is Lemma 4.8, where we show a useful property of the right-hand side of bipartite right-to-left expander graphs, which may also be useful in other applications of expanders. Roughly, when building a family of matchings in such a graph, given a partial matching and any node on the right we can either extend the matching to cover that node or exclude the node from ever being used in an extension of the matching.

We first introduce some notation. Let $\mathcal{G} = (U \cup V, E)$ be a bipartite graph. For a node a in $\mathcal{G}$ we will write $N(a)$ for the set of neighbors of a, and for a set of nodes A in $\mathcal{G}$ we will write $N(A)$ for $\bigcup_{a \in A} N(a)$.

For sets $A \subseteq U$ and $B \subseteq V$, a *2-matching σ of A into B* is a subset of the edge relation E such that each element of A has as neighbors under σ exactly two elements of B, and no two elements of A share a neighbor under σ. We will sometimes use functional notation for 2-matchings, as follows: for $a \in A$ we will write $\sigma(a)$ for the pair of neighbors of a; for $X \subseteq A$ we will write $\sigma(X)$ for the set of all neighbors of X; we will write $\mathrm{dom}(\sigma)$ for A and $\mathrm{ran}(\sigma)$ for $\sigma(A)$. A *fork* in $\mathcal{G}$ is a 2-matching with a domain of size one.

DEFINITION 4.1. *Let $\mathcal{G} = (U \cup V, E)$ be a bipartite graph. For $\gamma > 1$, we say that $\mathcal{G}$ is an* (s, γ)-expander *if*

$$\forall A \subseteq U,\ |A| \leq s \to |N(A)| \geq \gamma |A|.$$

We will usually be interested in $(s, 2+\epsilon)$-expanders for some $\epsilon > 0$. On subgraphs of such graphs we can apply the following corollary of Hall's theorem proved in [1].

LEMMA 4.2. *Let $\mathcal{G} = (U \cup V, E)$ be a bipartite graph. If $|N(A)| \geq 2|A|$ for every set $A \subseteq U$, then there is a 2-matching of U into V.*

For the rest of this section (until Theorem 4.13), fix integers d and s and a real number $\epsilon > 0$. Let $\mathcal{G} = (U \cup V, E)$ be a fixed bipartite graph of left-degree d which is an $(s, 2+\epsilon)$-expander.

DEFINITION 4.3. *Given two sets $A \subseteq U$ and $B \subseteq V$, we say that (A, B)* has the double-matching property *if for every $C \subseteq U \setminus A$, if $|A| + |C| \leq s$, then there exists a 2-matching of C into $V \setminus B$.*

We have the following useful lemma, which applies the expansion property of $\mathcal{G}$ to bound the size of a minimal witness C that the double-matching property fails.

LEMMA 4.4. *Let $A \subseteq U$ and $B \subseteq V$ be such that (A, B) does not have the double-matching property. Then there is a set $C \subseteq U \setminus A$ with $|C| < \frac{1}{\epsilon}|B|$ such that there is no 2-matching of C into $V \setminus B$.*

Proof. Let $C \subseteq U \setminus A$ be minimal such that $|C| \leq s - |A|$ and there is no 2-matching of C into $V \setminus B$. Then for every $D \subsetneq C$, there is a 2-matching of D into

$V \setminus B$, so in particular $|N(D) \setminus B| \geq 2|D|$. Hence we must have $|N(C) \setminus B| < 2|C|$, since otherwise there would be a 2-matching of C into $V \setminus B$ by Lemma 4.2. On the other hand, by expansion, since $|C| \leq s$ we have that $|N(C)| \geq (2+\epsilon)|C|$.

Combining these, we get

$$(2+\epsilon)|C| \leq |N(C)| \leq |N(C) \setminus B| + |B| < 2|C| + |B|$$

and hence $|C| < \frac{1}{\epsilon}|B|$. ∎

LEMMA 4.5. *The pair $(\emptyset, \emptyset)$ has the double-matching property.*

Proof. This follows directly from Lemma 4.2, since $\mathcal{G}$ is an $(s, 2+\epsilon)$ expander. ∎

LEMMA 4.6 (left extension). *Let $A \subseteq U$ and $B \subseteq V$ be such that (A, B) has the double-matching property and $\frac{d(d-1)}{\epsilon}(|B|+2) + |A| + 1 \leq s$. Then for each $u \in U \setminus A$ there is a 2-matching π of u into $V \setminus B$ such that $(A \cup \{u\}, B \cup \pi(u))$ has the double-matching property.*

Proof. Let Π be the set of all 2-matchings π of u into $V \setminus B$. Since $|A| + 1 \leq s$ and (A, B) has the double-matching property, we know that Π is nonempty. Suppose for a contradiction that for every $\pi \in \Pi$, the pair $(A \cup \{u\}, B \cup \pi(u))$ does not have the double-matching property. By Lemma 4.4, for every $\pi \in \Pi$ there is a set $C_\pi \subseteq U \setminus (A \cup \{u\})$ with $|C_\pi| < \frac{1}{\epsilon}|B \cup \pi(u)|$ such that there is no 2-matching of C_π into $V \setminus (B \cup \pi(u))$.

Let $C = \bigcup_{\pi \in \Pi} C_\pi$. Then $|C| < \frac{d(d-1)}{\epsilon}(|B|+2)$, since $|\Pi| \leq d(d-1)$. Hence, by our assumption about the sizes of $|A|$ and $|B|$, we have that $|C \cup \{u\}| \leq s - |A|$. Furthermore $C \cup \{u\} \subseteq U \setminus A$, so by the double-matching property for (A, B) there is a 2-matching σ of $C \cup \{u\}$ into $V \setminus B$.

There must be some $\pi \in \Pi$ such that $\pi(u) = \sigma(u)$. Let σ' be σ with the fork $u \mapsto \pi(u)$ removed. Then σ' is a 2-matching of C into $V \setminus (B \cup \pi(u))$ and in particular contains a 2-matching of C_π into $V \setminus (B \cup \pi(u))$, contradicting the choice of C_π. ∎

LEMMA 4.7 (left retraction). *Let $A \subseteq U$ and $B \subseteq V$ be such that (A, B) has the double-matching property and $\frac{1}{\epsilon}|B| + |A| \leq s$. Suppose that $u \in A$ and there is a 2-matching π of u into B. Then $(A \setminus \{u\}, B \setminus \pi(u))$ has the double-matching property.*

Proof. Let $C \subseteq (U \setminus A) \cup \{u\}$ with $|C| \leq s - |A \setminus \{u\}|$. We want to show that there is a 2-matching of C into $(V \setminus B) \cup \pi(u)$. By Lemma 4.4, it is enough to consider only sets C with $|C| < \frac{1}{\epsilon}|B \setminus \pi(u)|$.

If $u \in C$, then $|C \setminus \{u\}| \leq s - |A|$ so by the double-matching property for (A, B) there is a 2-matching σ of $C \setminus \{u\}$ into $V \setminus B$. Hence $\sigma \cup \pi$ is a 2-matching of C into $(V \setminus B) \cup \pi(u)$.

If $u \notin C$, then $|C| \leq s - |A|$ by our assumption about the sizes of $|A|$ and $|B|$, so by the double-matching property for (A, B) there is a 2-matching of C into $V \setminus B$. ∎

LEMMA 4.8 (right extension). *Let $A \subseteq U$ and $B \subseteq V$ be such that (A, B) has the double-matching property. Let $v \in V \setminus B$ have degree e, and suppose that $\frac{d(d-1)}{\epsilon}(|B| + 2e) + |A| + e \leq s$. Then either*

1. *for some $u \in U \setminus A$ there is a 2-matching π of u into $V \setminus B$ such that $v \in \pi(u)$ and $(A \cup \{u\}, B \cup \pi(u))$ has the double-matching property, or*
2. *$(A, B \cup \{v\})$ has the double-matching property.*

Proof. Let D be $N(v) \setminus A$, so that $|D| \leq e$. By applying Lemma 4.6 $|D|$ many times, we can find a 2-matching σ of D into $V \setminus B$ such that $(A \cup D, B \cup \sigma(D))$ has the double-matching property. Notice that $\frac{1}{\epsilon}(|B| + |\sigma(D)|) + |A| + |D| \leq s$ so that,

by Lemma 4.7, the double-matching property is preserved if we remove any number of elements from D and the corresponding forks from σ.

There are now two cases. In the first case, there is $u \in D$ and a corresponding fork π in σ such that $v \in \pi(u)$. In this case we may remove all other elements from D and all other forks from σ and thus satisfy condition 1 of the lemma.

In the second case, $v \notin \sigma(D)$. Then the double-matching property for $(A \cup D, B \cup \sigma(D))$ implies the double-matching property for $(A \cup D, B \cup \sigma(D) \cup \{v\})$, since no neighbors of v remain in $U \setminus (A \cup D)$. As in the previous case, it follows by Lemma 4.7 that $(A, B \cup \{v\})$ has the double-matching property, satisfying condition 2. ∎

LEMMA 4.9 (right retraction). *Let $A \subseteq U$ and $B \subseteq V$ be such that (A, B) has the double-matching property. For each $v \in V$, the pair $(A, B \setminus \{v\})$ has the double-matching property.*

Proof. This is trivial from the definition of the double-matching property. ∎

The following combinatorial objects are central to our lower bounds. They allow us to extend the space lower bound argument on standard matchings in bipartite graphs (see, for example, [4]) to a generalization of 2-matchings, in which we also allow single, unmatched points on the right which we call "singletons." We will use families of these objects to define our families of r-free assignments. The extension and downward closure properties will follow from respectively the extension and retraction properties in Definition 4.11.

DEFINITION 4.10. *A* 2-structure *κ in $\mathcal{G}$ is a pair (σ, S), where σ is a 2-matching and $S \subseteq V \setminus \operatorname{ran}(\sigma)$. We think of κ as consisting of a set of* forks *(the forks in σ) and a disjoint set of* singletons *(the elements of S).*

The size *of a 2-structure κ is defined to be $|\kappa| = |\operatorname{dom}(\sigma)| + |S|$, that is, the number of forks plus the number of singletons. Given two 2-structures $\kappa = (\sigma, S)$ and $\lambda = (\sigma', S')$ we say that λ* extends *κ, written $\kappa \subseteq \lambda$, if $\sigma \subseteq \sigma'$ and $S \subseteq S'$. We say that the 2-structure κ* covers *a node w in $U \cup V$ if $w \in \operatorname{dom}(\sigma) \cup \operatorname{ran}(\sigma) \cup S$.*

DEFINITION 4.11. *A nonempty set $\mathcal{F}$ of 2-structures in $\mathcal{G}$ is called an* r-covering family *if it has the following two properties:*

(RETRACTION) *If $\kappa \in \mathcal{F}$ and λ is a 2-structure in $\mathcal{G}$ with $\lambda \subseteq \kappa$, then $\lambda \in \mathcal{F}$.*

(EXTENSION) *If $\kappa \in \mathcal{F}$ with $|\kappa| < r$ and w is any node of $\mathcal{G}$, then κ can be extended to a 2-structure in $\mathcal{F}$ which covers w.*

LEMMA 4.12. *Let $r = s\epsilon/6d^2$. Suppose that no node in V has degree more than r. Then an r-covering family $\mathcal{F}$ of 2-structures exists on $\mathcal{G}$.*

Proof. For a 2-structure κ, let $A_\kappa = \operatorname{dom}(\sigma)$ and $B_\kappa = \operatorname{ran}(\sigma) \cup S$. We take $\mathcal{F}$ to be the set consisting of all 2-structures κ in $\mathcal{G}$ for which (A_κ, B_κ) has the double-matching property and $\frac{1}{\epsilon}|B_\kappa| + |A_\kappa| \leq s$.

This family is nonempty by Lemma 4.5 and has the retraction property by Lemmas 4.7 and 4.9. For the extension property, suppose that $|\kappa| < r$, that is, $|\operatorname{dom}(\sigma)| + |S| < r$. Then $|A_\kappa| < r$ and $|B_\kappa| = 2|\operatorname{dom}(\sigma)| + |S| < 2r$. Since $\mathcal{G}$ is an $(s, 2+\epsilon)$-expander of left-degree d, we must have $\epsilon < d$, so $r < s/6$. Thus

$$\frac{d(d-1)}{\epsilon}(|B_\kappa| + 2r) + |A_\kappa| + r < \frac{4d^2 r}{\epsilon} + 2r < \frac{4s}{6} + \frac{2s}{6} = s.$$

Hence the requirements on the sizes of A_κ and B_κ for Lemmas 4.6 and 4.8 are satisfied. Now given $v \in V$, applying Lemma 4.8 we can extend κ to a 2-structure κ' which covers v, by adding either one more fork or one more singleton. In either case,

$(A_{\kappa'}, B_{\kappa'})$ still has the double-matching property and $\frac{1}{\epsilon}|B_{\kappa'}|+|A_{\kappa'}| \leq s$, so we remain within $\mathcal{F}$. Similarly, given $u \in U$ we can apply Lemma 4.6 to extend κ to $\kappa' \in \mathcal{F}$ covering u. ∎

We will say that a graph $\mathcal{G}$ is a (n, d, Δ)-*random bipartite graph* if it is chosen uniformly at random from the set of bipartite graphs $(U \cup V, E)$ of left-degree d with $|U| = \Delta n$ and $|V| = n$. The next result is standard and can be found, for example, in [3] and Lemma 5.1 in [4].

THEOREM 4.13. *For any $d \geq 3$, $\Delta \geq 1$ and any real constant $\epsilon \in (0, d-2)$, there is a constant $\gamma = \gamma_{d,\epsilon,\Delta}$ such that, for large n, if G is a (n, d, Δ)-random bipartite graph, then, with high probability, G is a $(\gamma n, 1+\epsilon)$-expander.*

We are interested in $(s, 2+\epsilon')$-expander graphs with $\epsilon' > 0$ since on such graphs we can apply Lemma 4.2 and the constructions of this section. By Theorem 4.13 we can guarantee that a bipartite graph is an $(s, 2+\epsilon')$-expander by supposing that its left-degree d is at least 4 and taking $\epsilon = 1 + \epsilon'$.

LEMMA 4.14. *Choose constants $d \geq 4$ and $\Delta > 1$. There is a constant $\delta > 0$ such that, for large n, if $\mathcal{G}$ is a (n, d, Δ)-random bipartite graph, then with high probability there exists a δn-covering family of 2-structures on $\mathcal{G}$.*

Proof. Fix $\epsilon = 1.5$. Let γ be the constant $\gamma_{d,\epsilon,\Delta}$ from Theorem 4.13 and let $\delta = \gamma\epsilon/6d^2$. With high probability, $\mathcal{G}$ is a $(\gamma n, 2.5)$-expander. To show that $\mathcal{G}$ has a δn-covering family, by Lemma 4.12 it is enough to show that every node in V has degree at most δn. The degree of such a node is the sum of independent Boolean random variables and has expected value Δd, so this is true with high probability by the Chernoff bound (multiplicative version[3]). ∎

5. Random *k*-CNFs. A *random k-CNF* with n variables and clause density Δ is a CNF picked as follows: choose independently uniformly at random Δn clauses from the set of all possible clauses in the variables $\{x_1, \ldots, x_n\}$ containing exactly k literals. As is well-known, there is a constant θ_k such that if $\Delta > \theta_k$, then such a φ is unsatisfiable with high probability for large n.

RESTATED THEOREM 1.1. *Let $k \geq 4$ and $\Delta > 1$. There is a constant $c > 0$ such that, for large n, if φ is a random k-CNF with n variables and Δn clauses, then with high probability any resolution refutation of φ passes through a configuration containing cn clauses of width at least cn.*

Proof. We associate with φ the bipartite graph $\mathcal{G} = (U \cup V, E)$, where U is the set of clauses of φ, V is the set $\{x_1, \ldots, x_n\}$ of variables, and an edge exists between a clause C in U and a variable x in V if x appears in C (either positively or negatively). Then $\mathcal{G}$ is an (n, k, Δ)-random bipartite graph. Hence by Lemma 4.14 there is a constant δ such that with high probability there exists a δn-covering family $\mathcal{F}$ of 2-structures on $\mathcal{G}$. We will show how such a family $\mathcal{F}$ can be used to construct a family $\mathcal{H}$ of piecewise assignments that is δn-free for φ. The theorem then follows by Theorem 2.4, with $c = \delta/2$.

[3]The precise version of the Chernoff bound we use is the following: let $X_1, \ldots, X_m$ be independent random variables such that for each X_i, $0 \leq X_i \leq 1$ and let $\mu = \mathbb{E}[X_1] + \cdots + \mathbb{E}[X_m]$; then for any $\beta \geq 0$, $\Pr[\sum_{i=1}^m X_i \geq (1+\beta)\mu] \leq \exp(-\frac{\beta^2}{2+\beta}\mu)$.

Let $\kappa = (\sigma, S)$ be any 2-structure in $\mathcal{F}$ and consider the following way of labeling the forks and singletons of κ with partial assignments:

- Let $\pi : u \mapsto \{x_i, x_j\}$ be a fork in κ with $i < j$. Label π with an assignment to $\{x_i, x_j\}$ chosen as follows: either set x_i to satisfy the clause u and set x_j arbitrarily or set x_j to satisfy the clause u and set x_i arbitrarily.
- Label each singleton x_i in κ with an arbitrary assignment to x_i.

Notice that for both forks and singletons, for every variable x_i covered there is at least one possible label which sets $x_i \mapsto 1$ and one label which sets $x_i \mapsto 0$.

Let L be an assignment of such a label to every fork and singleton in κ. All the labels in L have disjoint domains. Hence we can use L to define a piecewise assignment α as the set of all labels chosen for the forks in κ together with all labels chosen for the singletons of κ. Then in particular $\|\alpha\| = |\kappa|$ and α satisfies every clause C covered by κ. We take $\mathcal{H}$ to consist of every piecewise assignment α which arises in this way from a 2-structure $\kappa \in \mathcal{F}$ and a labeling L of κ.

We now need to show that $\mathcal{H}$ satisfies Definition 2.1. It is clearly nonempty. For the downward closure, observe that given two piecewise assignments $\beta \sqsubseteq \alpha$, if $\alpha \in \mathcal{H}$, then there is some $\kappa \in \mathcal{F}$ such that α is a labeling of κ. We can obtain β from α by removing some pieces from α. Let κ' be the 2-structure obtained by removing the corresponding forks and singletons from κ. Then β is a labeling of κ' and $\kappa' \in \mathcal{F}$ by the retraction property for $\mathcal{F}$. Hence $\beta \in \mathcal{H}$.

For the consistency property, suppose for a contradiction that some $\alpha \in \mathcal{H}$ falsifies a clause C of φ. By the downward closure of $\mathcal{H}$ proved above, we may assume without loss of generality that $\|\alpha\| \leq k$ by removing any pieces of α which do not mention a variable in C and remembering that $|C| = k$. The piecewise assignment α arises as a labeling of some 2-structure $\kappa \in \mathcal{F}$ which cannot cover C, since otherwise α by construction would satisfy C. Since $|\kappa| = \|\alpha\| \leq k < \delta n$ for large n, by the extension property for $\mathcal{F}$ we can extend κ to a 2-structure κ' in $\mathcal{F}$ which does cover C and thus contains some fork $\pi : C \mapsto \{x_i, x_j\}$. Then in particular the variable x_i appears in C but is not in the domain of α, contradicting the assumption that α falsifies C.

For the extension property, suppose that $\alpha \in \mathcal{H}$ is a labeling of $\kappa \in \mathcal{F}$ with $|\kappa| < \delta n$, and let x_i be any variable not in the domain of α. Then x_i is not covered by κ. By the extension property for $\mathcal{F}$, we can extend κ to a 2-structure $\kappa' \in \mathcal{F}$ by adding either a fork or a singleton which covers x_i, and by the properties of our labelings we can extend α to a labeling α' of κ' which sets x_i to whichever value we choose. □

6. The graph pigeonhole principle. Let $\mathcal{G} = (U \cup V, E)$ be a bipartite graph with $|U| > |V|$. We think of U as a set of pigeons and V as a set of holes. The formula $\mathcal{G}$-PHP, the *graph pigeonhole principle for* $\mathcal{G}$, is an unsatisfiable CNF in variables $\{x_{uv} : (u, v) \in E\}$. It asserts that the variables describe a map, given by a subset of the edges of $\mathcal{G}$, in which each pigeon gets mapped to at least one hole but no hole receives two pigeons. Formally, it is a conjunction of all clauses

1. $\bigvee\{x_{uv} : (u, v) \in E\}$ for each $u \in U$,
2. $\neg x_{uv} \vee \neg x_{u'v}$ for each distinct pair of edges (u, v) and (u', v) in E.

We will call these clauses respectively the pigeon axioms and the hole axioms. Notice that if $\mathcal{G}$ has left-degree d, then $\mathcal{G}$-PHP is a d-CNF. We will write X_v for the set of variables representing the edges touching the hole v.

Restated Theorem 1.2. *Let $d \geq 4$ and $\Delta > 1$. There is a constant $c > 0$ such that, for large n, if $\mathcal{G}$ is a random graph chosen from the set of bipartite graphs with left-degree d with Δn pigeon nodes on the left and n hole nodes on the right, then with*

high probability any resolution refutation of $\mathcal{G}$-PHP *passes through a configuration containing* cn *clauses of width at least* cn.

Proof. The proof of this result closely follows the pattern of the proof of Theorem 1.1. By Lemma 4.14 there is a constant δ such that with high probability there exists a δn-covering family $\mathcal{F}$ of 2-structures on $\mathcal{G}$. We will construct from such an $\mathcal{F}$ a family $\mathcal{H}$ of piecewise assignments that is δn-free for $\mathcal{G}$-PHP. The result follows by Theorem 2.4.

Let $\kappa = (\sigma, S)$ be any 2-structure in $\mathcal{F}$ and consider the following way of labeling the forks and singletons of κ:

- Label each fork $\pi : u \mapsto \{v, v'\}$ in κ with an assignment α_π to $X_v \cup X_{v'}$ chosen as follows: order the holes v, v' arbitrarily as v_1, v_2. Map pigeon u to hole v_1 and set the remaining variables in X_{v_1} to zero. Either choose any pigeon $u' \in N(v_2)$ and map it to hole v_2 (we allow $u' = u$), setting the remaining variables in X_{v_2} to zero, or simply set all variables in X_{v_2} to zero.
- Label each singleton v in κ with an assignment α_v to X_v chosen as follows: either choose any pigeon $u \in N(v)$ and map it to v, setting all other variables in X_v to zero, or simply set all variables in X_v to zero.

Notice that in both cases, for every pigeon v covered and every variable $x \in X_v$, there is at least one label which sets $x \mapsto 1$ and one label which sets $x \mapsto 0$.

As in the proof of Theorem 1.2, we can label κ with a piecewise assignment α arising from our choice L of labels for the parts of κ. Notice that $\|\alpha\| = |\kappa|$, that α does not violate any hole axiom, and that α satisfies the pigeon axiom for each pigeon u covered by κ. We take $\mathcal{H}$ to consist of every piecewise assignment α which arises in this way from any $\kappa \in \mathcal{F}$ and any labeling L of κ. We now need to show that $\mathcal{H}$ satisfies Definition 2.1.

Clearly $\mathcal{H}$ is nonempty. The downward closure and consistency properties follow exactly as in Theorem 1.1, using the observation that no $\alpha \in \mathcal{H}$ falsifies any hole axiom. For the extension property, suppose that $\alpha \in \mathcal{H}$ is a labeling of some 2-structure $\kappa \in \mathcal{F}$ with $\|\alpha\| = |\kappa| < r$, and let x be any variable not in the domain of α. Then x must be in X_v for some hole v which is not covered by κ. By the extension property for $\mathcal{F}$, we can extend κ to a 2-structure $\kappa' \in \mathcal{F}$ by adding either a fork or a singleton which covers v. By the freedom in our choice of labelings, there is an extension β_0 of α to a labeling of κ' which sets x to zero, and another such extension β_1 which sets x to one. ∎

An alternative version of this theorem would be to show a total space lower bound for $\mathcal{G}$-PHP for all bipartite expanders of left-degree d with a suitable bound on the right-degree (rather than for random graphs), applying Lemma 4.12 directly to get the covering family of 2-structures.

7. Semantic total space. In this section we address a question raised in [1]. The space bounds in that paper hold not only for the usual versions of the proof systems considered but also for *semantic* versions of the systems. In particular a *semantic resolution* refutation of a CNF φ is a sequence of configurations where, at each step in the refutation, we can either add an axiom from φ to the current configuration M_i or replace M_i with *any* configuration M_{i+1} with the property that every clause in M_{i+1} is implied by M_i.

In [1] the authors show that for any unsatisfiable CNF φ, the clause space required to refute φ in resolution is no more than twice the clause space required in semantic resolution, and they ask whether the same thing is true for total space.

It follows from our lower bounds that, for total space, resolution can require quadratically more space than semantic resolution. In particular, let φ be an unsatisfiable random k-CNF with n variables and clause density Δ, where n is large. We can refute φ in semantic resolution by simply writing down all the clauses of φ and then deriving the empty clause in one step. This uses total space Δkn, the size of φ. But by Theorem 1.1, a resolution refutation of φ requires total space $\Omega(n^2)$.

On the other hand, the proof of Theorem 2.4 does not depend very much on the details of the syntax of the resolution rule. The theorem generalizes easily to give lower bounds for a weak form of semantic resolution, with the following inference rule: from a configuration M_i we can move to a configuration $M_i \cup \{C\}$, where the clause C is implied by some set of at most d clauses in M_i, for a fixed integer d. Calling this system *d-bounded semantic resolution*, we have the following theorem.

THEOREM 7.1. *Let φ be an unsatisfiable CNF formula and suppose $d \leq r$. If there is a family of piecewise assignments which is r-free for φ, then any d-bounded semantic resolution refutation of φ must pass through a configuration containing at least $(r-d)/2$ clauses each of width at least $(r-d)/2$.*

Proof. The proof is the same as for Theorem 2.4, except that we replace the bound $r/2$ with $(r-d)/2$ and use a different argument for the inference case, as follows. Suppose $M_{i+1} = M_i \cup \{E\}$, where E is implied by clauses $D_1, \ldots, D_d \in M_i$. Since $\|\alpha\| < (r-d)/2$ and $|M_i \cap S| < (r-d)/2$ we may assume that $\|\beta_i\| \leq \|\alpha\| + |M_i \cap S| < r - d$.

Either D_1 is satisfied by β_i or it is not. If it is, let $\gamma_1 = \beta_i$. If not, then D_1 cannot be in S, since β_i satisfies all members of $M_i \cap S$. It follows that D_1 is not falsified by β_i either and thus must contain some literal not set by β_i. In this case let $\gamma_1 \in \mathcal{H}$ be a minimal extension of β_i which satisfies this literal.

We have found $\gamma_1 \in \mathcal{H}$ which satisfies D_1 with $\beta_i \sqsubseteq \gamma_1$ and $\|\gamma_1\| < r - d + 1$. Applying the same reasoning to $D_2, \ldots, D_d$ in turn, we can build a sequence of extensions $\gamma_1 \sqsubseteq \gamma_2 \sqsubseteq \cdots \sqsubseteq \gamma_d$ in $\mathcal{H}$, finishing with γ_d which satisfies each of $D_1, \ldots, D_d$ and thus also satisfies E. We put $\beta_{i+1} = \gamma_d$. ∎

Finally, in [1] the notion of an *r-semiwide* formula is defined, and it is shown that any such formula requires clause space r in semantic resolution. We can strengthen this, to show that such a formula also requires total space $r^2/4$ in semantic resolution, by a straightforward generalization of the total space lower bounds in [1] for PHP_n and CT_n. For a CNF Z and a partial assignment α, we say that α is *Z-consistent* if α can be extended to satisfy Z.

DEFINITION 7.2. *A CNF formula φ is* r-semiwide *if it is the conjunction of a CNF Z and a CNF W, where Z is satisfiable, and for each Z-consistent partial assignment α and each clause C from W, if $|\alpha| < r$, then α can be extended to a Z-consistent assignment which satisfies C.*

THEOREM 7.3. *Let φ be an unsatisfiable r-semiwide formula. Then every semantic resolution refutation of φ must pass through a configuration containing $r/2$ clauses each of width at least $r/2$.*

Proof. Let $\varphi = Z \wedge W$ as in Definition 7.2 and let $\Pi = (M_1, \ldots, M_s)$ be a refutation of φ. Let $M_i^* = \{C \in M_i \ : \ Z \not\models C\}$. Take the first t such that there exists a clause $C \in M_t^*$ of width strictly less than $r/2$. Fix such a clause C and let α be the minimal partial assignment falsifying α. Then α is Z-consistent and $|\operatorname{dom}(\alpha)| = |C| < r/2$.

It is now enough to show that $|M_i^*| \geq r/2$ for some $i < t$, since for $i < t$ every clause in $|M_i^*|$ has width at least $r/2$. So suppose for a contradiction that $|M_i^*| < r/2$ for all $i < t$. We prove by induction that for each $i = 1, \ldots, t$ there exists some Z-consistent $\beta_i \supseteq \alpha$ such that $\beta_i \models M_i^*$. This leads immediately to a contradiction when $i = t$.

For the erasure case we trivially put $\beta_{i+1} = \beta_i$. For semantic inference, that is, $M_i \models M_{i+1}$, we let β_{i+1} be an extension of β_i which satisfies Z. Then from the fact that $\beta_{i+1} \models M_i^* \wedge Z$ it follows that $\beta_{i+1} \models M_i$ and hence $\beta_{i+1} \models M_{i+1}$. For axiom download, suppose $M_{i+1} = M_i \cup \{D\}$ with D a clause from W. We may assume without loss of generality that $|\operatorname{dom}(\beta)| \leq |\operatorname{dom}(\alpha)| + |M_i^*| < r$. Hence by r-semiwideness there is a Z-consistent $\beta_{i+1} \supseteq \beta_i$ such that $\beta_{i+1} \models D$. ∎

Acknowledgment. The authors are grateful to Jakob Nordström for helpful discussions about this work and about resolution space in general.

REFERENCES

[1] M. Alekhnovich, E. Ben-Sasson, A. A. Razborov, and A. Wigderson, *Space complexity in propositional calculus*, SIAM J. Comput., 31 (2002), pp. 1184–1211, http://dx.doi.org/10.1137/S0097539700366735.

[2] A. Atserias and V. Dalmau, *A combinatorial characterization of resolution width*, J. Comput. System Sci., 74 (2008), pp. 323–334, http://dx.doi.org/10.1016/j.jcss.2007.06.025.

[3] E. Ben-Sasson, *Expansion in Proof Complexity*, Ph.D. thesis, Hebrew University, Jerusalem, 2001.

[4] E. Ben-Sasson and N. Galesi, *Space complexity of random formulae in resolution*, Random Struct. Algorithms, 23 (2003), pp. 92–109, http://dx.doi.org/10.1002/rsa.10089.

[5] E. Ben-Sasson and J. Nordström, *Understanding space in resolution: Optimal lower bounds and exponential trade-offs*, in Computational Complexity of Discrete Problems, P. B. Miltersen, R. Reischuk, G. Schnitger, and D. van Melkebeek, eds., Dagstuhl Seminar Proceedings, 08381, Schloss Dagstuhl, Leibniz-Zentrum für Informatik, Germany, 2008, http://drops.dagstuhl.de/opus/volltexte/2008/1781/.

[6] E. Ben-Sasson and J. Nordström, *A space hierarchy for k-DNF resolution*, Electronic Colloquium on Computational Complexity, 16 (2009), p. 47, http://eccc.hpi-web.de/report/2009/047.

[7] E. Ben-Sasson and J. Nordström, *Understanding space in proof complexity: Separations and trade-offs via substitutions*, in Proceedings of Innovations in Computer Science, ICS 2010, Tsinghua University, Beijing, China, B. Chazelle, ed., Tsinghua University Press, 2011, pp. 401–416, http://conference.itcs.tsinghua.edu.cn/ICS2011/content/papers/3.html.

[8] E. Ben-Sasson and A. Wigderson, *Short proofs are narrow—resolution made simple*, J. ACM, 48 (2001), pp. 149–169, http://doi.acm.org/10.1145/375827.375835.

[9] P. Bennett, I. Bonacina, N. Galesi, T. Huynh, M. Molloy, and P. Wollan, *Space proof complexity for random* 3*-CNFs*, CoRR abs/1503.01613, 2015, http://arxiv.org/abs/1503.01613.

[10] O. Beyersdorff and O. Kullmann, *Unified characterisations of resolution hardness measures*, in Theory and Applications of Satisfiability Testing, C. Sinz and U. Egly, eds., Lecture Notes in Comput. Sci. 8561, Springer, New York, 2014, pp. 170–187, http://dx.doi.org/10.1007/978-3-319-09284-3_13.

[11] I. Bonacina, *Space in Weak Propositional Proof Systems*, Ph.D. thesis, Sapienza University of Rome, Italy, 2015.

[12] I. Bonacina, *Total space in resolution is at least width squared*, in 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016), LIPIcs. Leibniz Int. Proc. Inform 55, I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, and D. Sangiorgi, eds., Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016, pp. 56:1–56:13.

[13] I. Bonacina and N. Galesi, *A framework for space complexity in algebraic proof systems*, J. ACM, 62 (2015), pp. 23:1–23:20, http://doi.acm.org/10.1145/2699438.

[14] I. Bonacina, N. Galesi, and N. Thapen, *Total space in resolution*, in Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, IEEE Computer Society, 2014, pp. 641–650, http://dx.doi.org/10.1109/FOCS.2014.74.

[15] H. K. Büning and T. Lettmann, *Propositional Logic—Deduction and Algorithms*, Cambridge University Press, Cambridge, UK, 1999.

[16] S. A. Cook and R. A. Reckhow, *The relative efficiency of propositional proof systems*, J. Symbolic Logic, 44 (1979), pp. 36–50, http://dx.doi.org/10.2307/2273702.

[17] J. L. Esteban, N. Galesi, and J. Messner, *On the complexity of resolution with bounded conjunctions*, Theoret. Comput. Sci., 321 (2004), pp. 347–370, http://dx.doi.org/10.1016/j.tcs.2004.04.004.

[18] J. L. Esteban and J. Torán, *Space bounds for resolution*, Inform. and Comput., 171 (2001), pp. 84–97, http://dx.doi.org/10.1006/inco.2001.2921.

[19] Y. Filmus, M. Lauria, M. Mikša, J. Nordström, and M. Vinyals, *Towards an understanding of polynomial calculus: New separations and lower bounds*, in Proceedings of Automata, Languages, and Programming — 40th International Colloquium, ICALP 2013, Riga, Latvia, Part I, F. V. Fomin, R. Freivalds, M. Z. Kwiatkowska, and D. Peleg, eds., Lecture Notes in Comput. Sci. 7965, Springer, New York, 2013, pp. 437–448, http://dx.doi.org/10.1007/978-3-642-39206-1_37.

[20] Y. Filmus, M. Lauria, J. Nordström, N. Ron-Zewi, and N. Thapen, *Space complexity in polynomial calculus*, SIAM J. Comput., 44 (2015), pp. 1119–1153, http://dx.doi.org/10.1137/120895950.

[21] J. Nordström, *Narrow proofs may be spacious: Separating space and width in resolution*, SIAM J. Comput., 39 (2009), pp. 59–121, http://dx.doi.org/10.1137/060668250.

[22] J. Nordström, *Pebble games, proof complexity, and time-space trade-offs*, Log. Methods Comput. Sci., 9 (2013), http://dx.doi.org/10.2168/LMCS-9(3:15)2013.

[23] J. Nordström, *On the interplay between proof complexity and SAT solving*, ACM SIGLOG News, 2 (2015), pp. 19–44.

[24] J. Nordström and J. Håstad, *Towards an optimal separation of space and length in resolution*, Theory Comput., 9 (2013), pp. 471–557, http://dx.doi.org/10.4086/toc.2013.v009a014.