

Supercritical Space-Width Trade-offs for Resolution*

Christoph Berkholz
Humboldt-Universität zu Berlin

Jakob Nordström
KTH Royal Institute of Technology

December 21, 2016

Abstract

We show that there are CNF formulas which can be refuted in resolution in both small space and small width, but for which any small-width proof must have space exceeding by far the linear worst-case upper bound. This significantly strengthens the space-width trade-offs in [Ben-Sasson '09], and provides one more example of trade-offs in the “supercritical” regime above worst case recently identified by [Razborov '16]. We obtain our results by using Razborov’s new hardness condensation technique and combining it with the space lower bounds in [Ben-Sasson and Nordström '08].

1 Introduction

Propositional proof complexity studies the problem of how to provide concise, polynomial-time checkable certificates that formulas in conjunctive normal form (CNF) are unsatisfiable. Research in this area was initiated in [CR79] as a way of attacking the problem of showing that $\text{NP} \neq \text{coNP}$, and hence $\text{P} \neq \text{NP}$, and it is therefore natural that the main focus has been on proving upper and lower bounds on proof length/size. More recently, however, other complexity measures have also been investigated, and this study has revealed a rich and often surprising web of connections.

1.1 Resolution Length, Width, and Space

Arguably the most thoroughly studied proof system in proof complexity is *resolution*, which appeared in [Bla37] and began to be investigated in connection with automated theorem proving in the 1960s [DLL62, DP60, Rob65]. Because of its simplicity—there is only one derivation rule—and because all lines in a proof are clauses, this proof system is well suited for proof search, and it lies at the heart of current state-of-the-art SAT solvers based on so-called *conflict-driven clause learning* [BS97, MS99, MMZ⁺01].

It is not hard to show that any unsatisfiable CNF formula over n variables can be proven unsatisfiable, or *refuted*, by a resolution refutation containing $\exp(O(n))$ clauses, and this holds even in the restricted setting of *tree-like resolution*, where each intermediate clause in the refutation has to be rederived from scratch every time it is used. In the breakthrough paper [Hak85], Haken obtained a length lower bound on the form $\exp(\Omega(n^\delta))$ for general resolution refutations of so-called pigeonhole principle formulas, and this result was later followed by truly exponential lower bounds $\exp(\Omega(n))$ for other formula families in [Urq87, CS88, BKPS02] and many other papers.

In a seminal paper [BW01], Ben-Sasson and Wigderson identified *width*, measured as the largest size of any clause appearing in a refutation, as another interesting complexity measure for resolution. Clearly, any unsatisfiable CNF formula over n variables can be refuted in width at most n . Moreover, any resolution refutation in width w need never be longer than $n^{O(w)}$, since this is an upper bound on the number of distinct clauses of width w (and this naive counting argument is essentially tight [ALN16]).

*This is the full-length version of the paper with the same title that appeared in *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP '16)*.

What Ben-Sasson and Wigderson showed is that strong enough *lower* bounds on width also imply lower bounds on length; in particular that linear $\Omega(n)$ width lower bounds imply exponential $\exp(\Omega(n))$ length lower bounds for CNF formulas of bounded width. This connection can be used to rederive almost all currently known resolution length lower bounds.

Motivated by questions in SAT solving, where efficient memory management is a major concern, a more recent line of research in proof complexity has examined a third complexity measure on proofs, namely *space*. This study was initiated by Esteban and Torán [ET01], who defined the (*clause*) *space* of a resolution proof as the maximal number of clauses needed to be kept in memory during verification of the proof,¹ a definition that was generalized to other proof systems by Alekhovich et al. [ABRW02]. It should be noted that although the original impetus for investigating proof space came from the applied SAT solving side, space complexity is of course a well-studied measure in its own right in computational complexity, and the study of space in proof complexity has turned out to be of intrinsic interest in that it has uncovered intriguing connections to proof length and width. It can be shown that a CNF formula over n variables can always be refuted in space $n + O(1)$ even in tree-like resolution [ET01], although the refutation thus obtained might have exponential length. Linear space lower bounds matching the worst-case upper bound up to constant factors were obtained for a number of formula families in [ET01, ABRW02, BG03].

The space lower bounds obtained in the papers just discussed turned out to match closely known lower bounds on width, and in a strikingly simple and beautiful result Atserias and Dalmau [AD08] showed that in fact the resolution width of refuting a k -CNF formula F is a lower bound on the clause space required,² minus an additive term k adjusting for the largest width of any clause in F . This allows to recover the space lower bounds mentioned above as immediate consequences of width lower bounds shown in [BW01]. Furthermore, it follows from [AD08] that for $k = O(1)$ any k -CNF formula that can be refuted by just keeping a constant number of clauses in memory can also be refuted in polynomial length and constant width. These connections go only in one direction, however—in the sequence of papers [Nor09, NH13, BN08] it was shown that there are formula families that have high space complexity although they possess refutations in linear length and constant width.

1.2 Resolution Trade-offs

As was discussed above, a resolution proof in sufficiently small width will by necessity also be short, whereas the linear worst-case upper bound on space is achieved by a proof in exponential length. It is natural to ask, therefore, whether for a given formula F there exists a single resolution refutation of F that can simultaneously optimize these different complexity measures. The question of trade-offs between proof complexity measures was first raised by Ben-Sasson [Ben09], who gave a strong negative answer for space versus width. More precisely, what was shown in [Ben09] is that there are formulas which are refutable separately in constant width and in constant space, but for which any resolution proof minimizing one of the measures must exhibit almost worst-case linear behaviour with respect to the other measure.

A question that arises in the context of SAT solving is whether it is possible to simultaneously optimize size and space (corresponding to running time and memory usage). In addition to the space-width trade-offs discussed above, in [Ben09] Ben-Sasson also proved a size-space trade-off for the subsystem tree-like resolution, and building on [BN08, Ben09] it was shown in [BN11] for general resolution that there are formulas which have refutations in linear length and also in small space, but for which any space-efficient refutation must have superpolynomial or even exponential length. Beame et al. [BBI16] extended the range of parameters of the trade-offs further by exhibiting formulas over n variables refutable in length polynomial in n where bringing the space down to linear, or even just shaving a

¹For completeness, we want to mention that for resolution there is also a *total space* measure counting the total number of literals in memory (with repetitions), which has been studied in [ABRW02, BGT14, BBG⁺15, Bon16]. In this paper, however, “space” will always mean “clause space” in the sense of [ET01] unless otherwise stated.

²Note that this is a nontrivial connection since a lower bound on width, i.e., the *number of literals* in a clause, is shown to imply essentially the same lower bound on the *number of clauses* needed.

constant factor of the polynomial space bound that follows immediately from the length bound, incurs a superpolynomial penalty in proof length, a result that was generalized and strengthened in [BNT13].

Turning finally to the relation between length and width, what was shown in [BW01] is that a short resolution refutation can be converted to a refutation of small width, but the way this conversion is done in [BW01] blows up the length exponentially. Thapen [Tha14] proved that this is inherent by exhibiting formulas refutable in small width and small length, but for which any small-width refutation has to have exponential length. For the restricted case of tree-like resolution, Razborov [Raz16] recently showed that there are formulas refutable in small width for which any tree-like refutation even doing slightly better than the trivial linear upper bound with respect to width must by necessity have doubly exponential length.

We want to emphasize an intriguing property of the trade-off results in [BBI16, BNT13, Raz16] that was highlighted by Razborov, and that sets these results apart from the other trade-offs surveyed above. Namely, for most trade-off results between complexity measures it is the case that the trade-off plays out in the region between the worst-case upper bounds for the measures, where as one measure decreases the other measure has to approach its critical worst-case value. However, the short resolution proofs in [BBI16, BNT13] require space even polynomially larger than the worst-case upper bound, and the small-width tree-like proofs in [Raz16] require proofs exponentially larger than the exponential upper bound for tree-like length. To underscore the dramatic nature of such trade-off results, Razborov refers to them as *ultimate* in the preliminary version [Raz15] of [Raz16]. In this paper, we will instead use the term *supercritical trade-offs*, which we feel better describes the behaviour that one of the complexity measures is pushed up into the supercritical regime above worst case when the other measure is decreased.

1.3 Our Contribution

Answering Razborov’s call in [Raz16] for more examples of the type of trade-offs discussed above, in this paper we prove a supercritical trade-off between space and width in resolution. As already observed, any refutation in width w of a CNF formula over n variables in general resolution need not contain more than $O(n^w)$ clauses, which is also a trivial upper bound on the space complexity of such a refutation. Our main result is that this bound is essentially tight, and is also somewhat robust. Namely, we show that there are n -variable formulas that can be refuted in width w , but for which any refutation in width even up to almost a multiplicative logarithmic factor larger than this requires space $n^{\Omega(w)}$.

Theorem 1.1. *For any constant $\varepsilon > 0$ and any non-decreasing function $\ell(n)$, $6 \leq \ell(n) \leq n^{\frac{1}{2}-\varepsilon}$, there is a family $\{F_n\}_{n \in \mathbb{N}}$ of n -variable CNF formulas which can be refuted in resolution in width $\ell(n)$ but for which any resolution refutation in width $o(\ell(n) \log n)$ requires clause space at least $n^{\Omega(\ell(n))}$.*

1.4 Techniques

In one sentence, we obtain our results by using Razborov’s hardness condensation technique in [Raz16] and combining it with the space lower bounds in [BN08].

In slightly more detail, our starting point are the so-called *pebbling formulas* defined in [BW01]. These formulas are refutable in constant width, but it was observed in [Ben09] that space lower bounds for pebble games on directed acyclic graphs (DAGs) carry over to lower bounds on the *number of variables* kept simultaneously in memory in resolution refutations of pebbling formulas defined over these DAGs. It was shown in [BN08] that substituting every variable in such formulas by an exclusive or of two new variables and expanding out to CNF produces a new family of formulas which are still refutable in constant width but for which the variable space lower bounds have been amplified to clause space lower bounds.

The result in [BN08] is one of several examples of how *XOR substitution*, or *XORification*, has been used to amplify weak proof complexity lower bounds to much stronger lower bounds. In all of these applications distinct variables of the original formula are replaced by disjoint sets of new variables. The wonderfully simple (with hindsight) but powerful new idea in [Raz16] is to instead do XOR substitution

with overlapping sets of variables from a much smaller variable pool (but with exclusive ors of higher arity).

This recycling of variables has the consequence that hardness amplification as in [BN08] no longer works, since it crucially depends on the fact that all new substitution variables are distinct. What Razborov showed in [Raz16] was essentially that if the pattern of overlapping variable substitutions is described by a strong enough bipartite expander, then locally there are enough distinct new variables to make tree-like amplification lower bounds as in [Ben09] go through over a fairly wide range of the parameter space, yielding supercritical trade-offs between width and tree-like length. Since in addition the number of variables in the formula has decreased significantly, this can be viewed as a kind of *hardness condensation*.

We use Razborov’s idea of XORification with recycled variables, but since we want to obtain results not for tree-like but for DAG-like resolution the technical details of our proofs are somewhat different. At a high level, we start with formulas over N variables that are refutable in constant width but require space $\Omega(N/\log N)$. We modify these formulas by applying w -wise XORification using a much smaller set of n variables, and then show that from any refutation in width $O(w)$ of this new, XORified formula it is possible to recover a refutation of the original formula with comparable space complexity. But this means that any small-width refutation of the XORified formula must have space complexity roughly $\Omega(N/\log N)$. Choosing parameters so that $N \approx n^w$ yields the bound stated in Theorem 1.1.

We should point out that compared to [Raz16] we get significantly less robust trade-offs, which break down already for a multiplicative logarithmic increase in width. This is mainly due to the fact that we deal not with tree-like resolution as in [Raz16], but with the much stronger general resolution proof system producing DAG-like proofs. We share with [Raz16] the less desirable feature that although our formulas only have n variables they contain on the order of n^w clauses. Thus, measured in terms of formula size our space-width trade-offs do not improve on [Ben09], and the width of our formulas is not constant but scales linearly with w . Still, since the number of variables provides a worst-case upper bound on space (independently of formula size), measured in terms of variables it seems fair to say that the trade-off result in Theorem 1.1 is fairly dramatic.

1.5 Organization of This Paper

The rest of this paper is organized as follows. We start by reviewing some preliminaries in Section 2. In Section 3 we prove our main result assuming a hardness condensation lemma, and this lemma is then established in Section 4. We conclude in Section 5 with a discussion of possible directions for future research. For completeness, proofs of some technical claims are provided in Appendix A.

2 Preliminaries

A *literal* over a Boolean variable x is either the variable x itself (a *positive literal*) or its negation \bar{x} (a *negative literal*). We define $\bar{\bar{x}} = x$. A *clause* $C = a_1 \vee \dots \vee a_k$ is a disjunction of literals over pairwise disjoint variables (without loss of generality we assume that there are no trivial clauses containing both a variable and its negation). A clause C' *subsumes* another clause C if every literal from C' also appears in C . A k -*clause* is a clause that contains at most k literals. A *CNF formula* $F = C_1 \wedge \dots \wedge C_m$ is a conjunction of clauses, and F is a k -*CNF formula* if it consists of k -clauses. We write $\text{Vars}(F)$ to denote the set of variables appearing in a formula F . We think of clauses and CNF formulas as sets: the order of elements is irrelevant and there are no repetitions.

A *resolution refutation* $\pi : F \vdash \perp$ of an unsatisfiable CNF formula F , which can also be referred to as a *resolution proof* for (the unsatisfiability of) F , is an ordered sequence of clauses $\pi = (D_1, \dots, D_\tau)$ such that $D_\tau = \perp$ is the empty clause containing no literals, and each clause D_i , $i \in [\tau] = \{1, \dots, \tau\}$, is either one of the clauses in F (an *axiom*) or is derived from clauses D_j, D_k in π with $j, k < i$ by the *resolution rule*

$$\frac{B \vee x \quad C \vee \bar{x}}{B \vee C} \quad (2.1)$$

For technical reasons, it will also be convenient to permit a *weakening rule*

$$\frac{B}{B \vee C} \quad (2.2)$$

allowing to derive a strictly weaker clause from a clause already derived, although this rule is not essential.

With every resolution proof π we can associate a DAG G_π by having a sequence of vertices v_i on a line in order of increasing i , labelled by the clauses $D_i \in \pi$, and with directed edges (v_j, v_i) and (v_k, v_i) if the clause D_i was derived by resolution from D_j and D_k or an edge (v_j, v_i) if D_i was derived from D_j by weakening. Note that there might be several occurrences of a clause D in the proof π , and if so each occurrence gets its own vertex in G_π .

Now we can formally define the proof complexity measures discussed in Section 1. The *length* $L(\pi)$ of a resolution proof π is the number of clauses in it (counted with repetitions). The *width* $W(C)$ of a clause C is $|C|$, i.e., the number of literals, and the width $W(\pi)$ of a proof π is the size of a largest clause in π . The (*clause*) *space* at step i is the number of clauses C_j , $j < i$, with edges to clauses C_k , $k \geq i$ in G_π , plus 1 for the clause C_i derived at this step. Intuitively, space measures the number of clauses we need to keep in memory at step i , since they were derived before step i but are used to infer new clauses at or after step i . The space $Sp(\pi)$ of a proof π is the maximum space over all steps in π . Taking the minimum over all resolution refutations of a CNF formula F , we define the length, width, and space of refuting F , respectively, as $L(F \vdash \perp) = \min_{\pi: F \vdash \perp} \{L(\pi)\}$, $W(F \vdash \perp) = \min_{\pi: F \vdash \perp} \{W(\pi)\}$, and $Sp(F \vdash \perp) = \min_{\pi: F \vdash \perp} \{Sp(\pi)\}$. We remark that any applications of the weakening rule (2.2) can always be eliminated from a refutation without increasing the length, width, or space.

When reasoning about space, it is sometimes convenient to use a slightly different, but equivalent, description of resolution that makes explicit what clauses are in memory at each point in time. We say that a *configuration-style resolution refutation* is a sequence $(\mathbb{D}_0, \dots, \mathbb{D}_\tau)$ of sets of clauses, or *configurations*, such that $\mathbb{D}_0 = \emptyset$, $\perp \in \mathbb{D}_\tau$, and for all $t \in [\tau]$ the configuration \mathbb{D}_t is obtained from \mathbb{D}_{t-1} by one of the following *derivation steps*:

Axiom download $\mathbb{D}_t = \mathbb{D}_{t-1} \cup \{C\}$, where C is a clause $C \in F$.

Inference $\mathbb{D}_t = \mathbb{D}_{t-1} \cup \{D\}$ for a clause D derived by resolution (2.1) or weakening (2.2) from clauses in \mathbb{D}_{t-1} .

Erasure $\mathbb{D}_t = \mathbb{D}_{t-1} \setminus \mathbb{D}'$ for some $\mathbb{D}' \subseteq \mathbb{D}_{t-1}$.

The length of a configuration-style refutation $\pi = (\mathbb{D}_0, \dots, \mathbb{D}_\tau)$ is the number of axiom downloads and inference steps, the width is the size of a largest clause, as before, and the space is $\max_{t \in [\tau]} \{|\mathbb{D}_t|\}$. Given a refutation as an ordered sequence of clauses $\pi = (D_1, \dots, D_\tau)$, we can construct a configuration-style refutation in the same length, width, and space by deriving each clause D_i via an axiom download or inference step, and interleave with erasures of clauses C_j , $j < i$, as soon as these clauses have no edges to clauses C_k , $k \geq i$, in the associated DAG G_π . In the other direction, taking a configuration-style refutation and listing the sequence of axiom download and inference steps yields a standard resolution refutation in the same length, width, and space (assuming that clauses are erased as soon as possible). Thus, we can switch freely between these two ways of describing resolution refutations.

In this paper, it will be convenient for us to limit our attention to a (slightly non-standard) restricted form of resolution refutations as described next. We define a *homogeneous resolution refutation* to be a refutation where every resolution rule application is of the form

$$\frac{C \vee x \quad C \vee \bar{x}}{C} \quad (2.3)$$

The requirement of homogeneity is essentially without loss of generality, since we need to insert at most two weakening steps before each application of the resolution rule, which increases the width by at most 1, and the weakened clauses can then immediately be forgotten. We state this observation formally for the record.

Observation 2.1. *If a CNF formula F has a standard resolution refutation without weakening steps in length L , width w , and space s , then it has a homogeneous refutation in length at most $3L$, width at most $w + 1$, and space at most $s + 2$.*

As already mentioned, a useful trick to obtain hard CNF formulas for different proof systems and complexity measures, which will play a key role also in this paper, is *XORification*, i.e., substituting variables by exclusive ors of new variables and expanding out in the canonical way to obtain a new CNF formula. For example, the standard way to define binary XOR substitution for a positive literal x is

$$x[\oplus_2] = (x_1 \vee x_2) \wedge (\bar{x}_1 \vee \bar{x}_2) , \quad (2.4)$$

for a negative literal \bar{y} we have

$$\bar{y}[\oplus_2] = (y_1 \vee \bar{y}_2) \wedge (\bar{y}_1 \vee y_2) , \quad (2.5)$$

and applying binary XOR substitution to the clause $x \vee \bar{y}$ we obtain the CNF formula

$$(x \vee \bar{y})[\oplus_2] = x[\oplus_2] \vee \bar{y}[\oplus_2] = (x_1 \vee x_2 \vee y_1 \vee \bar{y}_2) \wedge (x_1 \vee x_2 \vee \bar{y}_1 \vee y_2) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee y_1 \vee \bar{y}_2) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee \bar{y}_1 \vee y_2) . \quad (2.6)$$

The XORification of a CNF formula F is the conjunction of all the formulas corresponding to the XORified clauses of F . We trust that the reader has no problems parsing this slightly informal definition by example or generalising it to substitutions with XOR of arbitrary arity (but see, e.g., Definition 2.12 in [Nor13] for a more rigorous treatment).

Usually, XORification is done in such a way that any two variables in the original formula are replaced by exclusive ors over disjoint sets of new variables. Razborov [Raz16] observed that it can sometimes be useful to allow XORification with overlapping sets of variables. Let us define this concept more carefully.

Definition 2.2 (XORification with recycling [Raz16]). Let F be a CNF formula over the set of variables u_1, \dots, u_N and let $\mathcal{G} = (U \dot{\cup} V, E)$ be a bipartite graph with left vertex set $U = \{u_1, \dots, u_N\}$ and right vertex set $V = \{v_1, \dots, v_n\}$. Then for the variables u_i we define the XORified literals $u_i[\mathcal{G}] = \bigoplus_{v \in \mathcal{N}(u_i)} v$ and $\bar{u}_i[\mathcal{G}] = \neg \bigoplus_{v \in \mathcal{N}(u_i)} v$ (where $\mathcal{N}(u_i)$ denotes the neighbours in V of u_i), for clauses $C \in F$ we define $C[\mathcal{G}] = \bigvee_{a \in C} a[\mathcal{G}]$ expanded out to CNF as in (2.6) but with trivial clauses pruned away, and the *XORification of F with respect to \mathcal{G}* is defined to be $F[\mathcal{G}] = \bigwedge_{C \in F} C[\mathcal{G}]$.

Note that if F is an N -variable k -CNF with m clauses and $\mathcal{G} = (\{u_1, \dots, u_N\} \dot{\cup} \{v_1, \dots, v_n\}, E)$ is a bipartite graph of left degree d , then $F[\mathcal{G}]$ is an n -variable kd -CNF formula with most $2^{d-1}m$ clauses. We want to highlight that by definition we have the equality

$$(C \vee a)[\mathcal{G}] = C[\mathcal{G}] \vee a[\mathcal{G}] \quad (2.7)$$

(where we can view the expressions in (2.7) either as the Boolean functions computed by these formulas or as the corresponding clause sets but with trivial clauses removed), and this will be convenient to use in some of our technical arguments.

We conclude this section with two simple observations that will also be useful in what follows.

Observation 2.3. *If F has a (homogeneous) resolution refutation in width w and \mathcal{G} has left degree bounded by d , then $F[\mathcal{G}]$ can be refuted in (homogeneous) resolution in width $2dw$.*

This is not hard to show, and follows, e.g., from the proof of Theorem 2 in [BN11] (strictly speaking, this theorem is for XORification *without* recycling, but recycling can only decrease the width).

Observation 2.4. *If F has a (homogeneous) resolution refutation π such that the associated DAG G_π has depth (i.e., longest path) s , then π can be carried out (in homogeneous resolution) in space $s + 2$ (possibly by repeating and/or reordering clauses in π).*

This second observation is essentially due to [ET01]. To see why this is true, note that the proof DAG G_π can be turned into a binary tree of the same depth by repeating vertices/clauses, and it is then straightforward to show that any tree-like proof DAG in depth s can be realized in space at most $s + 2$.

3 Proof of Main Theorem

In this section we present a proof of Theorem 1.1. The proof makes use of the following hardness condensation lemma, which will be established in the next section and is the main technical contribution of the paper.

Lemma 3.1 (Hardness condensation lemma). *For all $k \in \mathbb{N}^+$ and $\varepsilon > 0$ there exist $n_0 \in \mathbb{N}^+$ and $\delta > 0$ such that the following holds. Let ℓ and n be integers satisfying $n \geq n_0$ and $k \leq \ell \leq n^{\frac{1}{2}-\varepsilon}$, and suppose that F is an unsatisfiable CNF formula over $N = \lfloor n^{\delta\ell} \rfloor$ variables which requires width $W(F \vdash \perp) = k$ and space $Sp(F \vdash \perp) = s$ to be refuted in resolution.*

Then there is a bipartite graph $\mathcal{G} = (U \dot{\cup} V, E)$ with $|U| = N$ and $|V| = n$ such that the n -variable CNF formula $F[\mathcal{G}]$ has the following properties:

- $F[\mathcal{G}]$ can be refuted in width ℓ .
- Any refutation $\pi : F[\mathcal{G}] \vdash \perp$ in width $w \leq \ell \log n$ requires space $Sp(\pi) \geq (s - w - 3)2^{-w}$.

We want to apply this lemma to formulas of low width complexity but high space complexity as stated next.

Theorem 3.2 (BN08). *There is a family $\{F_N\}_{N \in \mathbb{N}}$ of N -variable 6-CNF formulas of size $\Theta(N)$ which can be refuted in width $W(F_N \vdash \perp) = 6$ but require space $Sp(F_N \vdash \perp) = \Omega(N/\log N)$.*

Combining Lemma 3.1 and Theorem 3.2, we can prove our main result.

Proof of Theorem 1.1. Recall that we want to prove that for any constant $\varepsilon > 0$ and any non-decreasing function $\ell(n) \leq n^{\frac{1}{2}-\varepsilon}$ there is a family $\{F_n\}_{n \in \mathbb{N}}$ of n -variable CNF formulas which have resolution refutations of width $\ell(n)$ but for which any refutation of width $o(\ell(n) \log n)$ requires clause space $n^{\Omega(\ell(n))}$.

From Theorem 3.2 we obtain constants $\varepsilon' > 0$ and $N_0 \in \mathbb{N}^+$ and a family of N -variable 6-CNF formulas F_N that require clause space $\varepsilon'N/\log N$ for all $N \geq N_0$. We want to apply hardness condensation as in Lemma 3.1 to these formulas. Let $\varepsilon > 0$ be given in Theorem 1.1 and fix $k = 6$. Plugging this into Lemma 3.1 provides constants $\delta > 0$ and $n_0 \in \mathbb{N}^+$, where in addition we choose n_0 large enough so that $\lfloor n_0^{\delta\ell(n_0)} \rfloor \geq N_0$ (this is always possible since $\delta\ell(n_0) \geq 6\delta > 0$).

For any $n \geq n_0$, set $N = \lfloor n^{\delta\ell(n)} \rfloor \geq N_0$ and let $\mathcal{G} = (U \dot{\cup} V, E)$ with $|U| = N$ and $|V| = n$ be a bipartite graph with properties as guaranteed by Lemma 3.1. Then the lemma says that $F_N[\mathcal{G}]$ is an n -variable formula which can be refuted in width ℓ , but for which every refutation of width $w \leq \frac{\ell}{4k} \log n$ requires clause space $(s - w - 3)2^{-w}$, where $s \geq \varepsilon'N/\log N = \varepsilon' \lfloor n^{\delta\ell(n)} \rfloor / \log \lfloor n^{\delta\ell(n)} \rfloor$ is the space lower bound for F_N . Choosing $w \leq \frac{\delta}{2} \cdot \ell(n) \log n$ (recall that $w = o(\ell(n) \log n)$ by assumption), the sequence of calculations

$$(s - w - 3)2^{-w} \geq (\varepsilon' \lfloor n^{\delta\ell(n)} \rfloor / \log \lfloor n^{\delta\ell(n)} \rfloor) - \frac{\delta}{2} \ell(n) \log n \geq \Omega\left(n^{\frac{\delta}{2}\ell(n)}\right) \quad (3.1)$$

yields the desired space lower bound. \square

If one looks more closely at what is going on inside the proof of Theorem 1.1, where Lemma 3.1 and Theorem 3.2 come together, one can make a somewhat intriguing observation.

As discussed in the introduction, Theorem 3.2 is shown by using so-called pebbling formulas, which we now describe briefly. Given a DAG \mathcal{D} with sources S and a unique sink z , and with all non-sources having fan-in 2, we let every vertex in \mathcal{D} correspond to a variable and define the *pebbling formula* over \mathcal{D} , denoted $Peb_{\mathcal{D}}$, to consist of the following clauses:

- for all $s \in S$, the clause s ;
- For all non-source vertices v with predecessors u_1, u_2 , the clause $\bar{u}_1 \vee \bar{u}_2 \vee v$;
- for the sink z , the clause \bar{z} .

Applying standard binary XOR substitution (without recycling) as in (2.6) to these formulas amplifies lower bounds on the number of variables in memory $VarSp(Peb_{\mathcal{D}} \vdash \perp)$ (which follow from properties of the chosen DAG \mathcal{D}) to lower bounds on the number of clauses $Sp(Peb_{\mathcal{D}}[\oplus_2] \vdash \perp)$. In Lemma 3.1 we then do another round of XOR substitution, this time with recycling, to decrease the number of variables while maintaining the space lower bound for small-width refutations. It is not entirely clear why we would need two separate rounds of XORification to achieve this result. In one sense, it would seem more satisfying to get a clean one-shot argument that just takes pebbling formulas and yields the supercritical trade-offs by only one round of XORification.

And in fact, if we are willing to accept a slightly weaker bound, we could make such a one-shot argument and apply substitution with recycling directly to the pebbling formulas. The reason for this is that one can actually prove a somewhat stronger version of hardness condensation than in Lemma 3.1, as we will see in Section 4. There is no need to require that the original formula should have high space complexity unconditionally, but it suffices that the formula exhibits a strong trade-off between width and clause space. Since the number of clauses times the maximal width of any clause is an upper bound on the total number of distinct variables in memory, for any resolution refutation π we have the inequality $Sp(\pi) \cdot W(\pi) \geq VarSp(\pi)$. In [Ben09] a variable space lower bound $VarSp(Peb_{\mathcal{D}} \vdash \perp) = \Omega(N/\log N)$ was presented (for appropriately chosen DAGs \mathcal{D}), implying that any width- w refutation requires clause space at least $\Omega(N/(w \log N))$. Since our hardness condensation step incurs a loss of a factor $1/2^w$, by starting with standard pebbling formulas and applying XORification with recycling directly we could obtain asymptotically similar bounds to those in Theorem 1.1 in one shot.

However, one can also argue that by combining Lemma 3.1 and Theorem 3.2 in the way done above one obtains a more modular proof, which shows that any formulas satisfying the conditions in Theorem 3.2 can be used for hardness condensation in a black-box fashion. This is why we chose to present the proof in this way.

4 Hardness Condensation

Let us now prove the hardness condensation lemma. We establish a slightly stronger version of the lemma below, which clearly subsumes Lemma 3.1.

Lemma 4.1 (Hardness condensation lemma, strong version). *For all $k \in \mathbb{N}^+$ and $\varepsilon > 0$ there are $n_0 \in \mathbb{N}^+$ and $\delta > 0$ such that the following holds. Let ℓ and n be integers satisfying $n \geq n_0$ and $k \leq \ell \leq n^{\frac{1}{2}-\varepsilon}$ and suppose that F is an unsatisfiable CNF formula over $N = \lfloor n^{\delta\ell} \rfloor$ variables which requires width $W(F \vdash \perp) = k$ to be refuted in resolution.*

Then there is a bipartite graph $\mathcal{G} = (U \dot{\cup} V, E)$ with $|U| = N$ and $|V| = n$ such that the n -variable CNF formula $F[\mathcal{G}]$ has the following properties:

- *The XORified formula $F[\mathcal{G}]$ can be refuted in width ℓ .*
- *Any resolution refutation $\pi : F[\mathcal{G}] \vdash \perp$ of the XORified formula $F[\mathcal{G}]$ in width $w \leq \ell \log n$ requires space $Sp(\pi) \geq (s - w - 3)2^{-w}$, where s is the minimal space of any refutation $\pi' : F \vdash \perp$ of the original formula F in width at most w .*

Clearly, the key to obtain Lemma 4.1 is to choose the right kind of graphs. As in [Raz16], we use boundary expander graphs where the right-hand side is significantly smaller than the left-hand side. Let us start by giving a proper definition of these graphs and reviewing the properties that we need from them. Most of our discussion of boundary expanders can be recovered from [Raz16], but since our setting of parameters is slightly different we give a self-contained presentation and also provide full proofs of all claims in Appendix A for completeness. We remark that there is also a significant overlap with [BN16a] in our treatment of expander graphs below.

In what follows, we will let $\mathcal{G} = (U \dot{\cup} V, E)$ denote a bipartite graph with left vertices U and right vertices V . We write $\mathcal{N}^{\mathcal{G}}(U') = \{v \mid \{u, v\} \in E(\mathcal{G}), u \in U'\}$ to denote the set of right neighbours of a

left vertex subset $U' \subseteq U$ (and vice versa for right vertex subsets), dropping the graph \mathcal{G} from the notation when it is clear from context. For a single vertex v we will use the abbreviation $\mathcal{N}(v) = \mathcal{N}(\{v\})$.

Definition 4.2 (Boundary expander). A bipartite graph $\mathcal{G} = (U \dot{\cup} V, E)$ is an $N \times n$ (r, c) -boundary expander, or *unique neighbour expander*, if $|U| = N$, $|V| = n$, and for every set $U' \subseteq U$, $|U'| \leq r$, it holds that $|\partial(U')| \geq c|U'|$, where $\partial(U') = \{v \in \mathcal{N}^{\mathcal{G}}(U') : |\mathcal{N}^{\mathcal{G}}(v) \cap U'| = 1\}$ is the *boundary* or the set of *unique neighbours* of U' . An (r, d, c) -boundary expander is an (r, c) -boundary expander where additionally $|\mathcal{N}^{\mathcal{G}}(u)| \leq d$ for all $u \in U$, i.e., where the left degree is bounded by d .

An important property of (r, c) -boundary expanders, which holds for arbitrarily small but positive expansion $c > 0$, is that any left vertex subset $U' \subseteq U$ of size $|U'| \leq r$ has a matching into V . In addition, this matching can be chosen in such a way that there is an ordering of the vertices in U' such that every vertex $u_i \in U'$ is matched to a vertex outside of the neighbourhood of the preceding vertices u_1, \dots, u_{i-1} . The proof of this fact uses what is sometimes referred to as a *peeling argument*, which we recapitulate below for the convenience of the reader.

Lemma 4.3 (Peeling lemma). Let $\mathcal{G} = (U \dot{\cup} V, E)$ be an (r, c) -boundary expander with $r \geq 1$ and $c > 0$. Then every left vertex subset $U' \subseteq U$ of size $|U'| = \ell \leq r$ can be ordered $U' = (u_1, \dots, u_\ell)$ in such a way that there is a matching into an ordered right vertex subset $V' = (v_1, \dots, v_\ell) \subseteq V$ for which $v_i \in \mathcal{N}(u_i) \setminus \mathcal{N}(\{u_1, \dots, u_{i-1}\})$.

Proof. The proof is by induction on ℓ . The base case $\ell = 1$ is immediate since $r \geq 1$ and $c > 0$ implies that no left vertex can be isolated. For the induction step, suppose the lemma holds for $\ell - 1$. To define the sequence v_1, \dots, v_ℓ we first fix any $v_\ell \in \partial(U')$, which exists because $|\partial(U')| \geq c|U'| > 0$. Since v_ℓ is in the boundary of U' there exists a unique $u_\ell \in U'$ such that $|\mathcal{N}(v_\ell) \cap U'| = \{u_\ell\}$. Thus, for this pair (u_ℓ, v_ℓ) it holds that $v_\ell \in \mathcal{N}(u_\ell) \setminus \mathcal{N}(U' \setminus \{u_\ell\})$. By the induction hypothesis we can now find sequences $u_1, \dots, u_{\ell-1}$ and $v_1, \dots, v_{\ell-1}$ for $U' \setminus \{u_\ell\}$ such that $v_i \in \mathcal{N}(u_i) \setminus \mathcal{N}(\{u_1, \dots, u_{i-1}\})$, to which we can append u_ℓ and v_ℓ at the end. The lemma follows by the induction principle. \square

For a right vertex subset $V' \subseteq V$ in $\mathcal{G} = (U \dot{\cup} V, E)$ we define the *kernel* $\ker(V') \subseteq U$ of V' to be the set of all left vertices whose entire neighbourhood is contained in V' , i.e.,

$$\ker(V') = \{u \in U \mid \mathcal{N}(u) \subseteq V'\} . \quad (4.1)$$

We write $\mathcal{G} \setminus V'$ to denote the subgraph of \mathcal{G} induced on $(U \setminus \ker(V')) \dot{\cup} (V \setminus V')$. In other words, we can think of $\mathcal{G} \setminus V'$ as being obtained from \mathcal{G} by first deleting V' and afterwards all isolated vertices from U .

Another key property of boundary expanders is that for any small enough right vertex set V' we can always find a *closure* $\gamma(V') \supseteq V'$ with a small kernel on the left such that the subgraph $\mathcal{G} \setminus \gamma(V')$ has good boundary expansion. This is very similar to an analogous lemma in [Raz16], but since our parameters are slightly different we provide a proof of the next lemma in Appendix A.

Lemma 4.4. Let \mathcal{G} be an $(r, 2)$ -boundary expander. Then for every $V' \subseteq V$ with $|V'| \leq r/2$ there exists a set of vertices $\gamma(V') \supseteq V'$ such that $|\ker(\gamma(V'))| \leq |V'|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(r/2, 1)$ -boundary expander.

The next lemma states that there exist $N \times n$ $(r, d, 2)$ -boundary expanders where the size n of the right-hand side is significantly smaller than the size $N = n^{\Theta(d)}$ of the left-hand side. The proof, which closely follows [Raz16, Lemma 2.2], is a standard application of the probabilistic method, but is included in Appendix A for completeness.

Lemma 4.5. Fix constants $\varepsilon, \delta > 0$ and $d_0 \geq 2$ such that $\delta + \frac{1}{d_0} < \varepsilon/2$. Then there exists an $n_0 \in \mathbb{N}^+$ such that for all n, d , and r satisfying $n \geq n_0$, $d_0 \leq d \leq n^{1/2-\varepsilon}$, and $r \leq n^{1/2}$ there are $\lfloor n^{\delta d} \rfloor \times n$ $(r, d, 2)$ -boundary expanders.

After this review of boundary expanders and their properties we now come to the core argument of the paper, namely that space lower bounds are preserved for small-width resolution refutations when we apply XORification as in Definition 2.2 with respect to an $(r, 2)$ -boundary expander. To get cleaner technical arguments in the proofs we will restrict our attention to homogeneous resolution refutations as in (2.3), which for our purposes is without loss of generality by Observation 2.1.

Lemma 4.6 (Main technical lemma). *Let F be an unsatisfiable CNF-formula and \mathcal{G} an $(r, 2)$ -boundary expander, and suppose that $\pi : F[\mathcal{G}] \vdash \perp$ is a homogeneous resolution refutation in width $w \leq r/2$ of the XORified formula $F[\mathcal{G}]$. Then there is a homogeneous refutation $\pi' : F \vdash \perp$ of the original formula F in width at most w and space $Sp(\pi') \leq 2^w Sp(\pi) + w + 3$.*

Proof. Assume that $\pi = (\mathbb{C}_0, \mathbb{C}_1, \dots, \mathbb{C}_\tau)$ is a configuration-style homogeneous resolution refutation of $F[\mathcal{G}]$ in width $W(\pi) = w \leq r/2$. We will show how to transform π into a refutation π' of the original formula F in width and space as claimed in the lemma. To help the reader navigate the proof, we remark that in what follows we will use the notational conventions that B and C denote clauses over $\text{Vars}(F[\mathcal{G}])$, D denotes a clause over $\text{Vars}(F)$, and A denotes an axiom clause from the original formula F before XORification.

Recall that for clauses $C \in F[\mathcal{G}]$ we have $\text{Vars}(C) \subseteq V$ by construction. For convenience, we will overload notation and write $\ker(C) = \ker(\text{Vars}(C))$, which is a subset of the variables U of the original formula F . Furthermore, for every clause $C \in \pi$ we fix $\gamma(C) := \gamma(\text{Vars}(C)) \subseteq V$ to be a minimal closure with properties as guaranteed by Lemma 4.4, i.e., such that $|\ker(\gamma(V'))| \leq |V'|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(r/2, 1)$ -boundary expander. Note that such closures exist since all clauses $C \in \pi$ have width at most w . It might be worth pointing out, though, that this is a purely existential statement—we have no control over how these closures are constructed, and, in particular, for two clauses B and C such that $B \subseteq C$ it does not necessarily hold that $\gamma(B) \subseteq \gamma(C)$.

An important notion in what follows will be that of *simultaneous falsifiability*, where we say that two CNF formulas F and G are *simultaneously falsifiable* if there is a truth value assignment that at the same time falsifies both F and G . To transform the resolution refutation π of $F[\mathcal{G}]$ into a refutation π' of F we let \mathbb{D}_t be obtained from \mathbb{C}_t by replacing every clause $C \in \mathbb{C}_t$ by the set of clauses

$$\mathcal{G}^{-1}(C) := \{D \mid \text{Vars}(D) = \ker(\gamma(C)); D[\mathcal{G}] \text{ and } C \text{ are simultaneously falsifiable}\} \quad (4.2)$$

and defining

$$\mathbb{D}_t := \bigcup_{C \in \mathbb{C}_t} \mathcal{G}^{-1}(C) \quad (4.3)$$

(where the notation $\mathcal{G}^{-1}(C)$ is chosen to suggest that this is in some intuitive sense the “inverse operation” of XORification with respect to \mathcal{G}). Every clause in $D \in \mathcal{G}^{-1}(C)$ has width at most w , because

$$|\text{Vars}(D)| = |\ker(\gamma(C))| \leq W(C) \leq w, \quad (4.4)$$

where the first inequality is guaranteed by Lemma 4.4 and the second inequality is by assumption. Furthermore, we have $|\mathcal{G}^{-1}(C)| \leq 2^w$, since all clauses in $\mathcal{G}^{-1}(C)$ are over the same set of variables and each variable appears positively or negatively in every clause, and hence

$$|\mathbb{D}_t| \leq 2^w |\mathbb{C}_t| \leq 2^w Sp(\pi). \quad (4.5)$$

We want to argue that the sequence $(\mathbb{D}_0, \mathbb{D}_1, \dots, \mathbb{D}_\tau)$ is the “backbone” of a resolution refutation π' of F , by which we mean that for every t it holds that \mathbb{D}_{t+1} can be derived from \mathbb{D}_t by a sequence of intermediate steps without affecting any proof complexity measure too much.

To make this claim formal, we first observe that for $\mathbb{C}_0 = \emptyset$ we obviously get $\mathbb{D}_0 = \emptyset$ by (4.3). Moreover, it holds that $\mathcal{G}^{-1}(\perp) = \{\perp\}$ and hence $\perp \in \mathbb{D}_\tau$, since the unique minimal closure of the empty set is the empty set itself. We want to show that for every $0 \leq t < \tau$ the configuration \mathbb{D}_{t+1} can be obtained from \mathbb{D}_t by a resolution derivation $(\mathbb{D}_t = \mathbb{D}_{t,0}, \mathbb{D}_{t,1}, \mathbb{D}_{t,2}, \dots, \mathbb{D}_{t,j_t-1}, \mathbb{D}_{t,j_t} = \mathbb{D}_{t+1})$, where the space of every intermediate configuration is bounded by $\max\{Sp(\mathbb{D}_t), Sp(\mathbb{D}_{t+1})\} + w + 3$.

If \mathbb{C}_{t+1} is obtained from \mathbb{C}_t by erasing a clause C , then \mathbb{D}_{t+1} can be obtained from \mathbb{D}_t by erasing all clauses $\mathcal{G}^{-1}(C) \setminus \mathbb{D}_{t+1}$. Suppose that \mathbb{C}_{t+1} is obtained from \mathbb{C}_t by downloading an axiom $C \in F[\mathcal{G}]$. We claim that every clause in $\mathcal{G}^{-1}(C)$ is either an axiom or a weakening of an axiom from F . By the definition of $F[\mathcal{G}]$, every axiom $C \in F[\mathcal{G}]$ is a clause in the CNF formula $A[\mathcal{G}]$ for some original axiom $A \in F$. Fix any axiom $A \in F$ such that $C \in A[\mathcal{G}]$. Then for all $D \in \mathcal{G}^{-1}(C)$ it holds by (4.2) that $\text{Vars}(D) = \ker(\gamma(C)) \supseteq \ker(C) \supseteq \text{Vars}(A)$ and that there is an assignment falsifying both $D[\mathcal{G}]$ and C . To see that this implies that A subsumes D , suppose that there is a variable x appearing positively in A such that $\bar{x} \in D$. Any truth value assignment falsifying $D[\mathcal{G}]$ must falsify $a[\mathcal{G}]$ for all literals $a \in D$, and hence in particular $\bar{x}[\mathcal{G}]$. This means that $x[\mathcal{G}]$ is satisfied by the same assignment, and then so is all of the formula $A[\mathcal{G}]$ including C . But this is a contradiction to the simultaneous falsifiability of $D[\mathcal{G}]$ and C , and so not only does it hold that $\text{Vars}(A) \subseteq \text{Vars}(D)$ but A is in fact a subclause of D as claimed. From this we see that we can add the clauses $\mathcal{G}^{-1}(C)$ to \mathbb{D}_t using axiom download and weakening. After applying a weakening step we immediately delete the old clause. Hence, the additional weakening might increase the space by at most one. It follows that the space of the intermediate configurations need never exceed $Sp(\mathbb{D}_{t+1}) + 1$.

It remains to argue that \mathbb{D}_{t+1} can be derived from \mathbb{D}_t when \mathbb{C}_{t+1} is obtained from \mathbb{C}_t by an inference step. This is stated in the following two claims regarding applications of the resolution and weakening rules.

Claim 4.7. Every clause $D \in \mathcal{G}^{-1}(C)$ can be derived from $\mathcal{G}^{-1}(C \vee x) \cup \mathcal{G}^{-1}(C \vee \bar{x})$ by a homogeneous resolution derivation of width w and depth $w + 1$.

Claim 4.8. For any two clauses B and C with $B \subseteq C$ it holds that every clause $D \in \mathcal{G}^{-1}(C)$ can be derived from $\mathcal{G}^{-1}(B)$ by a homogeneous derivation of width w and depth $w + 1$.

Taking these two claims on faith for now, let us see how they allow us to conclude the proof of the lemma. Since the depth of a refutation provides an upper bound on the clause space by Observation 2.4, it follows that in both cases we can derive all clauses in the clause set $\mathcal{G}^{-1}(C)$ one by one by using additional space $w + 3$ to perform the derivations in depth $w + 1$. This shows that F has a homogeneous resolution refutation π' of width w and clause space $Sp(\pi') \leq 2^w Sp(\pi) + w + 3$, which establishes the lemma. \square

We proceed to establish Claims 4.7 and 4.8.

Proof of Claim 4.7. Recall that by Lemma 4.4 the subgraph $\mathcal{G}_C := \mathcal{G} \setminus \gamma(C)$ is an $(r/2, 1)$ -boundary expander and that for $\ker(\gamma(C \vee x)) = \ker(\gamma(C \vee \bar{x}))$ we have $|\ker(\gamma(C \vee x))| \leq W(C \vee x) \leq w \leq r/2$. Therefore, we can apply Lemma 4.3 to the set $K = \ker(\gamma(C \vee x)) \setminus \ker(\gamma(C))$ to obtain an ordering u_1, \dots, u_ℓ of K satisfying $\mathcal{N}^{\mathcal{G}_C}(u_i) \setminus \mathcal{N}^{\mathcal{G}_C}(\{u_1, \dots, u_{i-1}\}) \neq \emptyset$. For $0 \leq i \leq \ell$ we let

$$K^i := (\ker(\gamma(C)) \cap \ker(\gamma(C \vee x))) \cup \{u_j \mid 1 \leq j \leq i\} \quad (4.6)$$

so that $K^\ell = \ker(\gamma(C \vee x))$ and $K^0 \subseteq \ker(\gamma(C))$, and define

$$\mathbb{D}^i := \{D \mid \text{Vars}(D) = K^i; D[\mathcal{G}] \text{ and } C \text{ are simultaneously falsifiable}\} . \quad (4.7)$$

Observe that

$$\begin{aligned} & \mathcal{G}^{-1}(C \vee x) \cup \mathcal{G}^{-1}(C \vee \bar{x}) \\ = & \{D \mid \text{Vars}(D) = \ker(\gamma(C \vee x)); D[\mathcal{G}] \text{ and } C \vee x \text{ are simultaneously falsifiable}\} \\ & \cup \{D \mid \text{Vars}(D) = \ker(\gamma(C \vee \bar{x})); D[\mathcal{G}] \text{ and } C \vee \bar{x} \text{ are simultaneously falsifiable}\} \\ = & \{D \mid \text{Vars}(D) = \ker(\gamma(C \vee x)); D[\mathcal{G}] \text{ and } C \text{ are simultaneously falsifiable}\} \\ = & \mathbb{D}^\ell \end{aligned} \quad (4.8)$$

and that every clause in $\mathcal{G}^{-1}(C)$ is subsumed by a clause in \mathbb{D}^0 since $K^0 \subseteq \ker(\gamma(C))$. Thus, we are done if we can derive all clauses in \mathbb{D}^0 from the clauses in \mathbb{D}^ℓ .

We do so inductively: for $i = \ell, \ell - 1, \dots, 2, 1$ we can obtain any clause $D \in \mathbb{D}^{i-1}$ by an application of the homogeneous resolution rule to the clauses $D \vee u_i$ and $D \vee \bar{u}_i$, which we claim are both available in \mathbb{D}^i . What remains to show is that $D \in \mathbb{D}^{i-1}$ indeed implies that $\{D \vee u_i, D \vee \bar{u}_i\} \subseteq \mathbb{D}^i$. To argue this, note that by the definition of \mathbb{D}^{i-1} in (4.7) there is a (partial) truth value assignment α that simultaneously falsifies $D[\mathcal{G}]$ and C . The peeling lemma guarantees that $\mathcal{N}^{\mathcal{G}C}(u_i) \setminus \text{Vars}(D[\mathcal{G}]) = \mathcal{N}^{\mathcal{G}C}(u_i) \setminus \mathcal{N}^{\mathcal{G}}(K^{i-1})$ has a non-empty intersection with $V \setminus \gamma(C)$, the right-hand side of the expander \mathcal{G}_C . Hence, we can extend α and set the variables in $\mathcal{N}^{\mathcal{G}C}(u_i) \setminus (\text{Vars}(D[\mathcal{G}]) \cup \text{Vars}(C)) \supseteq \mathcal{N}^{\mathcal{G}C}(K^i) \setminus \mathcal{N}^{\mathcal{G}C}(K^{i-1}) \neq \emptyset$ to appropriate values so that the parity $\bigoplus_{v \in \mathcal{N}(u_i)} \alpha(v)$ is even and thus $(D \vee u_i)[\mathcal{G}] = D[\mathcal{G}] \vee u_i[\mathcal{G}]$ is falsified, and we do so without assigning any variables in C , which therefore remains falsified. In an analogous fashion, by instead ensuring that the parity $\bigoplus_{v \in \mathcal{N}(u_i)} \alpha(v)$ is odd we get a falsifying assignment for $(D \vee \bar{u}_i)[\mathcal{G}] \vee C$. Hence, by (4.7) it holds that $D \vee u_i$ and $D \vee \bar{u}_i$ both appear in \mathbb{D}^i .

Finally, to get from \mathbb{D}^0 to $\mathcal{G}^{-1}(C)$ we might need an extra weakening step as observed above. The total depth of the whole derivation is at most $\ell + 1 \leq w + 1$. \square

Proof of Claim 4.8. Note that if $\ker(\gamma(B)) \subseteq \ker(\gamma(C))$ this claim would be easy to establish, but as noted above we have no guarantee that this is the case. Instead, we apply a proof strategy similar to the one for the previous claim. We again have that $\mathcal{G}_C := \mathcal{G} \setminus \gamma(C)$ is an $(r/2, 1)$ -boundary expander, so that we can apply the peeling lemma to the left-hand vertex set $\ker(\gamma(B)) \setminus \ker(\gamma(C))$ to obtain an ordering u_1, \dots, u_ℓ of its vertices satisfying $\mathcal{N}^{\mathcal{G}C}(u_i) \setminus \mathcal{N}^{\mathcal{G}C}(\{u_1, \dots, u_{i-1}\}) \neq \emptyset$. For $0 \leq i \leq \ell$ we let

$$K^i := (\ker(\gamma(C)) \cap \ker(\gamma(B))) \cup \{u_j \mid 1 \leq j \leq i\} \quad (4.9)$$

and as before define

$$\mathbb{D}^i := \{D \mid \text{Vars}(D) = K^i; D[\mathcal{G}] \text{ and } C \text{ are simultaneously falsifiable}\} . \quad (4.10)$$

Note that $\mathbb{D}^\ell \subseteq \mathcal{G}^{-1}(B)$, because if $D[\mathcal{G}]$ and C are simultaneously falsifiable, then $D[\mathcal{G}]$ and $B \subseteq C$ are certainly simultaneously falsifiable. Hence, we can obtain \mathbb{D}^ℓ from $\mathcal{G}^{-1}(B)$ by just erasing clauses. Once more, we apply the peeling argument in an inductive fashion and derive any $D \in \mathbb{D}^{i-1}$ from $D \vee u_i$ and $D \vee \bar{u}_i$ appearing in \mathbb{D}^i . In the end, we can infer any clause in $\mathcal{G}^{-1}(C)$ from \mathbb{D}^0 because every clause in $\mathcal{G}^{-1}(C)$ can be seen to be a weakening of some clause in \mathbb{D}^0 . \square

We can now combine the construction in Lemma 4.6 with the existence of good boundary expanders in Lemma 4.5 to prove the hardness condensation in Lemma 4.1.

Proof of Lemma 4.1. Given $\varepsilon > 0$ and $k \in \mathbb{N}^+$ we choose $\delta := \frac{\varepsilon}{10k}$. Note that we can assume $\varepsilon \leq 1/2$ since otherwise the lemma is vacuous. Suppose ℓ and n are parameters such that $k \leq \ell \leq n^{\frac{1}{2}-\varepsilon}$ and let F be an unsatisfiable CNF formula over $N = \lfloor n^{\delta\ell} \rfloor$ variables that can be refuted in width k . To apply Lemma 4.5 we set $d_0 := \frac{5}{\varepsilon} > 2$ and verify that $\delta + \frac{1}{d_0} = \frac{\varepsilon}{10k} + \frac{\varepsilon}{5} < \frac{\varepsilon}{2}$. We choose the degree of the expander to be $d := \lfloor \frac{\ell}{2k} \rfloor$ and set the size guarantee for expanding left vertex sets to $r := 2\ell \log n$. By the bound on ℓ we have $d \leq \ell \leq n^{\frac{1}{2}-\varepsilon}$. Furthermore, we choose n_0 large enough so that $r \leq 2n^{\frac{1}{2}-\varepsilon} \log n \leq n^{\frac{1}{2}}$ for all $n \geq n_0$.

Now we have two cases. The first, and interesting, case is when $d \geq d_0$ holds. Then Lemma 4.5 guarantees that there exists an $N \times n$ $(r, d, 2)$ -boundary expander \mathcal{G} . Applying XORification with respect to \mathcal{G} , we obtain a CNF formula $F[\mathcal{G}]$ with n variables. By Observation 2.3 it holds that $F[\mathcal{G}]$ has a resolution refutation of width $2dk \leq \ell$. Now suppose that $\pi : F[\mathcal{G}] \vdash \perp$ is a refutation of width w . Because $w \leq \ell \log n = r/2$ the space lower bound follows from Lemma 4.6.

The second case is when $d < d_0$. Then we do not actually need any XORification but can use the original formula. Formally, let $\mathcal{G} = (U \dot{\cup} (V \cup V'), E)$ be a matching between two sets U and V of size $|U| = |V| = N$ plus some isolated vertices V' on the right-hand side such that $|V \cup V'| = n$. To check that this is well defined we have to verify that $N \leq n$, which follows from the calculations $N = \lfloor n^{\delta\ell} \rfloor = \lfloor n^{\frac{\varepsilon}{10k} 2kd} \rfloor \leq \lfloor n^{\frac{\varepsilon}{10k} 2kd_0} \rfloor = \lfloor n^{\frac{\varepsilon}{10k} 2k \frac{5}{\varepsilon}} \rfloor = n$. In this somewhat convoluted way we obtain $F[\mathcal{G}] = F$ (plus some left-over dummy variables) and we have $W(F[\mathcal{G}] \vdash \perp) = W(F \vdash \perp) = k \leq \ell$ as well as $Sp(\pi) \geq s \geq (s - w - 3)2^{-w}$. The lemma follows. \square

5 Concluding Remarks

In this paper we prove that there are CNF formulas over n variables exhibiting an $n^{\Omega(w)}$ clause space lower bound for resolution refutations in width w . This lower bound is optimal (up to constants in the exponent) as every refutation in width w has length, and hence space, at most $n^{O(w)}$. Our lower bounds do not only hold for the minimal refutation width w but remain valid for any refutations in width asymptotically smaller than $w \log n$. Measured in terms of the number of variables n , this is a major improvement over the previous space-width trade-off result in [Ben09], and provides another example of trade-offs in the supercritical regime above worst-case recently identified in [Raz16].

Regarding possible future research directions, a first open problem is whether the range of applicability can be extended even further so that the space lower bound holds true up to width $o(n)$. It is clear that the lower bound has to break down at some point, since if one is allowed maximal width n any formula can be refuted in clause space $n + 2$ [ET01]. A supercritical trade-off on resolution proof depth over width ranging from w all the way up to $n^{1-\epsilon}/w$ was shown in [Raz16], suggesting that the above goal might not be completely out of reach.

Another intriguing open problem is to prove space trade-offs that are superlinear not only in terms of the number of variables but measured also in formula size. Such lower bounds cannot be obtained by the techniques used in this paper, but they are likely to exist as the following argument shows (see [Her08] for a more detailed discussion). Suppose that every refutation in width $w(n)$ can be transformed into a refutation that has width $w(n)$ and clause space polynomial in the size of the formula. Then we can find such a refutation non-deterministically in polynomial space by keeping the current configuration in memory and guessing the inference steps. Thus, by Savitch's theorem, finding refutations of width $w(n)$ would be in deterministic PSPACE. On the other hand, it has been shown by the first author that the problem of finding resolution refutations of bounded width is EXPTIME-complete [Ber12]. Hence, unless EXPTIME = PSPACE there are formulas where every refutation of minimal width needs clause space that is superpolynomial in the size of the formula.

Finally, it would be interesting to study if the supercritical trade-offs between clause space and width in resolution shown in this paper could be extended to similar trade-offs between monomial space and degree for polynomial calculus or polynomial calculus resolution as defined in [ABRW02, CEI96].

Acknowledgements

We wish to thank Alexander Razborov for patiently explaining the hardness condensation technique in [Raz16] during numerous and detailed discussions.

Part of the work of the first author was performed while at KTH Royal Institute of Technology supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD). The research of the second author was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement no. 279611 and by Swedish Research Council grants 621-2010-4797 and 621-2012-5645.

References

- [ABRW02] Michael Alekhovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *SIAM Journal on Computing*, 31(4):1184–1211, 2002. Preliminary version in *STOC '00*.
- [AD08] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, May 2008. Preliminary version in *CCC '03*.

- [ALN16] Albert Atserias, Massimo Lauria, and Jakob Nordström. Narrow proofs may be maximally long. *ACM Transactions on Computational Logic*, 17:19:1–19:30, May 2016. Preliminary version in *CCC '14*.
- [BBG⁺15] Patrick Bennett, Ilario Bonacina, Nicola Galesi, Tony Huynh, Mike Molloy, and Paul Wollan. Space proof complexity for random 3-CNFs. Technical Report 1503.01613, arXiv.org, April 2015.
- [BBI16] Paul Beame, Chris Beck, and Russell Impagliazzo. Time-space tradeoffs in resolution: Superpolynomial lower bounds for superlinear space. *SIAM Journal on Computing*, 45(4):1612–1645, August 2016. Preliminary version in *STOC '12*.
- [Ben09] Eli Ben-Sasson. Size-space tradeoffs for resolution. *SIAM Journal on Computing*, 38(6):2511–2525, May 2009. Preliminary version in *STOC '02*.
- [Ber12] Christoph Berkholz. On the complexity of finding narrow proofs. In *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS '12)*, pages 351–360, October 2012.
- [BG03] Eli Ben-Sasson and Nicola Galesi. Space complexity of random formulae in resolution. *Random Structures and Algorithms*, 23(1):92–109, August 2003. Preliminary version in *CCC '01*.
- [BGT14] Ilario Bonacina, Nicola Galesi, and Neil Thapen. Total space in resolution. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS '14)*, pages 641–650, October 2014.
- [BKPS02] Paul Beame, Richard Karp, Toniann Pitassi, and Michael Saks. The efficiency of resolution and Davis-Putnam procedures. *SIAM Journal on Computing*, 31(4):1048–1075, 2002. Preliminary versions of these results appeared in *FOCS '96* and *STOC '98*.
- [Bla37] Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.
- [BN08] Eli Ben-Sasson and Jakob Nordström. Short proofs may be spacious: An optimal separation of space and length in resolution. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 709–718, October 2008.
- [BN11] Eli Ben-Sasson and Jakob Nordström. Understanding space in proof complexity: Separations and trade-offs via substitutions. In *Proceedings of the 2nd Symposium on Innovations in Computer Science (ICS '11)*, pages 401–416, January 2011.
- [BN16a] Christoph Berkholz and Jakob Nordström. Near-optimal lower bounds on quantifier depth and Weisfeiler-Leman refinement steps. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '16)*, pages 267–276, July 2016.
- [BN16b] Christoph Berkholz and Jakob Nordström. Near-optimal lower bounds on quantifier depth and Weisfeiler-Leman refinement steps. Technical Report TR16-135, Electronic Colloquium on Computational Complexity (ECCC), August 2016. Preliminary version in *LICS '16*.
- [BNT13] Chris Beck, Jakob Nordström, and Bangsheng Tang. Some trade-off results for polynomial calculus. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC '13)*, pages 813–822, May 2013.

References

- [Bon16] Ilario Bonacina. Total space in resolution is at least width squared. In *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP '16)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 56:1–56:13, July 2016.
- [BS97] Roberto J. Bayardo Jr. and Robert Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI '97)*, pages 203–208, July 1997.
- [BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version in *STOC '99*.
- [CEI96] Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 174–183, May 1996.
- [CR79] Stephen A. Cook and Robert Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, March 1979.
- [CS88] Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988.
- [DLL62] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem proving. *Communications of the ACM*, 5(7):394–397, July 1962.
- [DP60] Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.
- [ET01] Juan Luis Esteban and Jacobo Torán. Space bounds for resolution. *Information and Computation*, 171(1):84–97, 2001. Preliminary versions of these results appeared in *STACS '99* and *CSL '99*.
- [Hak85] Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985.
- [Her08] Alexander Hertel. *Applications of Games to Propositional Proof Complexity*. PhD thesis, University of Toronto, May 2008. Available at <http://www.cs.utoronto.ca/~ahertel/>.
- [MMZ⁺01] Matthew W. Moskewicz, Conor F. Madigan, Ying Zhao, Lintao Zhang, and Sharad Malik. Chaff: Engineering an efficient SAT solver. In *Proceedings of the 38th Design Automation Conference (DAC '01)*, pages 530–535, June 2001.
- [MS99] João P. Marques-Silva and Karem A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 48(5):506–521, May 1999. Preliminary version in *ICCAD '96*.
- [NH13] Jakob Nordström and Johan Håstad. Towards an optimal separation of space and length in resolution. *Theory of Computing*, 9:471–557, May 2013. Preliminary version in *STOC '08*.
- [Nor09] Jakob Nordström. Narrow proofs may be spacious: Separating space and width in resolution. *SIAM Journal on Computing*, 39(1):59–121, May 2009. Preliminary version in *STOC '06*.
- [Nor13] Jakob Nordström. Pebble games, proof complexity and time-space trade-offs. *Logical Methods in Computer Science*, 9:15:1–15:63, September 2013.

- [Raz15] Alexander A. Razborov. An ultimate trade-off in propositional proof complexity. Technical Report TR15-033, Electronic Colloquium on Computational Complexity (ECCC), March 2015.
- [Raz16] Alexander A. Razborov. A new kind of tradeoffs in propositional proof complexity. *Journal of the ACM*, 63:16:1–16:14, April 2016.
- [Rob65] John Alan Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, January 1965.
- [Tha14] Neil Thapen. A trade-off between length and width in resolution. Technical Report TR14-137, Electronic Colloquium on Computational Complexity (ECCC), October 2014.
- [Urq87] Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, January 1987.

A Appendix

In this appendix we give proofs for Lemmas 4.4 and 4.5. As already mentioned, most of this material appears in a similar form in [Raz16] (although the exact parameters are slightly different), and there is also a substantial overlap with analogous technical lemmas in [BN16b]. In fact, Lemma 4.4 is exactly as stated in in [BN16b], but we present a proof below to give a self-contained exposition of our version of Razborov’s hardness condensation technique adapted to general resolution.

Lemma 4.4 (restated). *Let \mathcal{G} be a bipartite $(r, 2)$ -boundary expander. Then for every right vertex set $V' \subseteq V$ of size $|V'| \leq r/2$ there exists a superset $\gamma(V') \supseteq V'$ such that $|\ker(\gamma(V'))| \leq |V'|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(r/2, 1)$ -boundary expander.*

Proof. With assumptions as in the lemma, let $\mathcal{G} = (U \dot{\cup} V, E)$ be an $(r, 2)$ -boundary expander and let $V' \subseteq V$ be a right vertex set of size $|V'| \leq r/2$. We will construct an increasing sequence of right vertex sets $V' = V_0 \subset V_1 \subset \dots \subset V_\tau$ such that for $\gamma(V') = V_\tau$ it holds that $\mathcal{G} \setminus V_\tau$ is an $(r/2, 1)$ -boundary expander.

If $\mathcal{G} \setminus V_0$ is an $(r/2, 1)$ -boundary expander, then we can stop right away, but otherwise there must exist a left vertex set U_1 of size at most $r/2$ such that $|\partial^{\mathcal{G} \setminus V_0}(U_1)| \leq |U_1|$. Delete U_1 and all its neighbours from $\mathcal{G} \setminus V_0$. If now the resulting graph is an $(r/2, 1)$ -boundary expander, then we are done, but otherwise we repeat this process and iteratively delete vertex sets that violate the expansion requirements. Formally, for $i \geq 1$ we let U_i be any left vertex set of size at most $r/2$ such that $|\partial^{\mathcal{G} \setminus V_{i-1}}(U_i)| \leq |U_i|$, where we set $V_i := V_0 \cup \bigcup_{j=1}^i \mathcal{N}^{\mathcal{G}}(U_j)$ (and where we note that what is deleted at the i th step is $\mathcal{N}^{\mathcal{G} \setminus V_{i-1}}(U_i)$ together with the kernel $\ker(\mathcal{N}^{\mathcal{G} \setminus V_{i-1}}(U_i))$ of this right vertex set, so that after the i th step all of $\mathcal{N}^{\mathcal{G}}(U_i)$ and $\ker(\mathcal{N}^{\mathcal{G}}(U_i))$ has been removed from the graph).

Since all sets U_i constructed above are non-empty, this process must terminate for some $i = \tau$ and the resulting graph $\mathcal{G} \setminus V_\tau$ is then an $(r/2, 1)$ -boundary expander (if nothing else, an empty graph without vertices vacuously satisfies the expansion condition). However, we need to check that the condition $|\ker(V_\tau)| \leq |V_0|$ holds. This follows from the next claim.

Claim A.1. Let $V_{-1} = U_0 = \emptyset$ and suppose that $i \geq 0$. Then for U_i and V_i as constructed above we have the following properties:

1. For all U' such that $\ker(V_{i-1}) \cup U_i \subseteq U' \subseteq \ker(V_i)$ it holds that $|\partial^{\mathcal{G}}(U') \setminus V_0| \leq |\ker(V_i)|$.
2. The kernel of V_i has size $|\ker(V_i)| \leq |V_0|$.

We establish Claim A.1 by induction. For the base case $i = 0$, Property 1 holds since $U' \subseteq \ker(V_0)$ implies that $\partial^{\mathcal{G}}(U') \subseteq V_0$. For Property 2, suppose that $|\ker(V_0)| \leq r$. Then by the expansion of \mathcal{G} we have that $2|\ker(V_0)| \leq |\partial^{\mathcal{G}}(\ker(V_0))|$, and in combination with $\partial^{\mathcal{G}}(\ker(V_0)) \subseteq V_0$ this implies

$|\ker(V_0)| \leq \frac{1}{2}|V_0|$. If instead $|\ker(V_0)| > r$, then we can find a subset $U' \subseteq \ker(V_0)$ of size $|U'| = r$ for which it holds by expansion that $|\partial^{\mathcal{G}}(U')| \geq 2r$. But this is a contradiction since as argued above we should have $|\partial^{\mathcal{G}}(U')| \leq |V_0| \leq r/2$.

For the induction step, suppose that Property 1 and Property 2 both hold for $i - 1$. Let us write $U^* = \ker(V_{i-1}) \cup U_i$ and consider any U' such that $U^* \subseteq U' \subseteq \ker(V_i)$. We claim that every vertex in $\partial^{\mathcal{G}}(U')$ is either in the boundary $\partial^{\mathcal{G}}(U^*)$ or is a member of V_0 . To see why this is so, observe that since $U' \subseteq \ker(V_i)$ we have $\partial^{\mathcal{G}}(U') \subseteq V_i = V_0 \cup \bigcup_{j=1}^i \mathcal{N}^{\mathcal{G}}(U_j)$. Furthermore, note that $\bigcup_{j=1}^i U_j \subseteq U^* \subseteq U'$ holds (which is due to the fact that $\mathcal{N}(\ker(V')) \subseteq V'$ for any V'). Hence, for any $v \in \partial^{\mathcal{G}}(U') \setminus V_0$ it must be the case that $v \in \bigcup_{j=1}^i \mathcal{N}^{\mathcal{G}}(U_j)$, and so the unique neighbour of v on the left is contained in $\bigcup_{j=1}^i U_j$ and therefore also in U^* , implying that $v \in \partial(U^*)$. From this we can conclude that

$$\partial^{\mathcal{G}}(U') \setminus V_0 \subseteq \partial^{\mathcal{G}}(U^*) \setminus V_0 , \quad (\text{A.1})$$

and we will use this to show that

$$|\partial^{\mathcal{G}}(U^*) \setminus V_0| = |\partial^{\mathcal{G}}(\ker(V_{i-1}) \cup U_i) \setminus V_0| \leq |\ker(V_i)| \quad (\text{A.2})$$

in order to prove Property 1.

By definition, it holds that every vertex in $V_{i-1} \setminus V_0$ has at least one neighbour in $\ker(V_{i-1})$. It follows that for $U^* = \ker(V_{i-1}) \cup U_i$ all new boundary vertices in $\partial^{\mathcal{G}}(U^*) \setminus \partial^{\mathcal{G}}(\ker(V_{i-1}))$ are either from V_0 or from the boundary $\partial^{\mathcal{G}\setminus V_{i-1}}(U_i)$ of U_i that lies outside of V_{i-1} . Therefore we have

$$\partial^{\mathcal{G}}(U^*) \setminus V_0 = \partial^{\mathcal{G}}(\ker(V_{i-1}) \cup U_i) \setminus V_0 \subseteq (\partial^{\mathcal{G}}(\ker(V_{i-1})) \setminus V_0) \cup \partial^{\mathcal{G}\setminus V_{i-1}}(U_i) . \quad (\text{A.3})$$

Since we have chosen U_i so that it does not satisfy the expansion condition we know that

$$|\partial^{\mathcal{G}\setminus V_{i-1}}(U_i)| \leq |U_i| \quad (\text{A.4})$$

and by the inductive hypothesis for Property 1 it holds that

$$|\partial^{\mathcal{G}}(\ker(V_{i-1})) \setminus V_0| \leq |\ker(V_{i-1})| . \quad (\text{A.5})$$

Combining (A.1) with (A.3)–(A.5) we conclude that

$$\begin{aligned} |\partial^{\mathcal{G}}(U') \setminus V_0| &\leq |\partial^{\mathcal{G}}(\ker(V_{i-1}) \cup U_i) \setminus V_0| \leq \\ &\leq (|\partial^{\mathcal{G}}(\ker(V_{i-1})) \setminus V_0| + |\partial^{\mathcal{G}\setminus V_{i-1}}(U_i)|) \leq |\ker(V_{i-1})| + |U_i| \leq |\ker(V_i)| , \end{aligned} \quad (\text{A.6})$$

where the last inequality holds since $\ker(V_{i-1})$ and U_i are disjoint subsets of $\ker(V_i)$. This completes the inductive step for Property 1.

To show Property 2, let us first assume that $|\ker(V_i)| \leq r$. Then by the expansion properties of \mathcal{G} together with Property 1 applied to the set $U' = \ker(V_i)$ we have

$$2|\ker(V_i)| \leq |\partial^{\mathcal{G}}(\ker(V_i))| \leq |V_0| + |\ker(V_i)| , \quad (\text{A.7})$$

from which it follows that

$$|\ker(V_i)| \leq |V_0| . \quad (\text{A.8})$$

If instead $|\ker(V_i)| > r$, then by the inductive hypothesis we know that $|\ker(V_{i-1})| \leq |V_0| \leq r/2$ and by construction we have $|U_i| \leq r/2$. Therefore, there must exist a vertex set U' of size r satisfying the condition $\ker(V_{i-1}) \cup U_i \subseteq U' \subseteq \ker(V_i)$ in Property 1. From the expansion properties of \mathcal{G} we conclude that $|\partial(U')| \geq 2r$, which is a contradiction because for sets U' satisfying the conditions in Property 1 we derived (A.6), which implies that $|\partial(U')| \leq |V_0| + |\ker(V_{i-1})| + |U_i| \leq 3r/2$. The claim follows by the induction principle. \square

Lemma 4.5 (restated). Fix constants $\varepsilon, \delta > 0$ and $d_0 \geq 2$ such that $\delta + \frac{1}{d_0} < \varepsilon/2$. Then there exists an $n_0 \in \mathbb{N}^+$ such that for all n, d , and r satisfying $n \geq n_0$, $d_0 \leq d \leq n^{1/2-\varepsilon}$, and $r \leq n^{1/2}$ there are $\lfloor n^{\delta d} \rfloor \times n$ $(r, d, 2)$ -boundary expanders.

Proof. Let U and V be two disjoint sets of vertices of size $|U| = N = \lfloor n^{\delta d} \rfloor$ and $|V| = n$. For every $u \in U$ we choose d times a neighbour $v \in V$ uniformly at random with repetitions. This gives us a bipartite graph $\mathcal{G} = (U \cup V, E)$ of left-degree at most d . In the sequel we show that \mathcal{G} is almost surely an $(r, d, 2)$ -boundary expander as $n \rightarrow \infty$.

First note that for every set $U' \subseteq U$ all neighbours $v \in \mathcal{N}(U') \setminus \partial(U')$ that are not in the boundary of U' have at least two neighbours in U' . Since there are at most $d|U'| - |\partial(U')|$ edges between U' and $\mathcal{N}(U') \setminus \partial(U')$, it follows that $|\mathcal{N}(U') \setminus \partial(U')| \leq (d|U'| - |\partial(U')|)/2$ and hence

$$\begin{aligned} |\mathcal{N}(U')| &= |\mathcal{N}(U') \cap \partial(U')| + |\mathcal{N}(U') \setminus \partial(U')| \leq \\ &\leq |\partial(U')| + \frac{d|U'| - |\partial(U')|}{2} = \frac{d|U'| + |\partial(U')|}{2}. \end{aligned} \quad (\text{A.9})$$

If \mathcal{G} is not an $(r, d, 2)$ -boundary expander, then there is a set U' of size $\ell \leq r$ that has a boundary $\partial(U')$ of size at most 2ℓ and from (A.9) it follows that $|\mathcal{N}(U')| \leq (1 + d/2)\ell$. By a union bound argument we obtain

$$\Pr[\mathcal{G} \text{ is not an } (r, d, 2)\text{-boundary expander}] \quad (\text{A.10a})$$

$$\leq \sum_{\ell=1}^r \sum_{U' \subseteq [N]; |U'|=\ell} \Pr[|\partial(U')| \leq 2\ell] \quad (\text{A.10b})$$

$$\leq \sum_{\ell=1}^r \binom{N}{\ell} \Pr[|\mathcal{N}(U')| \leq (1 + d/2)\ell \text{ for some fixed } |U'| = \ell] \quad (\text{A.10c})$$

$$\leq \sum_{\ell=1}^r \binom{N}{\ell} \binom{n}{(1 + d/2)\ell} \left(\frac{(1 + d/2)\ell}{n} \right)^{d\ell} \quad (\text{A.10d})$$

$$\leq \sum_{\ell=1}^r N^\ell \left(\frac{en}{(1 + d/2)\ell} \right)^{(1+d/2)\ell} ((1 + d/2)\ell)^{d\ell} n^{-d\ell} \quad (\text{A.10e})$$

$$= \sum_{\ell=1}^r N^\ell (en)^{(1+d/2)\ell} ((1 + d/2)\ell)^{(d/2-1)\ell} n^{-d\ell} \quad (\text{A.10f})$$

$$\leq \sum_{\ell=1}^r n^{\delta d \ell} (en)^{(1+d/2)\ell} ((1 + d/2)\ell)^{(d/2-1)\ell} n^{-d\ell} \quad (\text{A.10g})$$

$$= \sum_{\ell=1}^r n^{\delta d \ell} n^{\frac{\log e}{\log n} (1+d/2)\ell} n^{\frac{1}{\log n} \log((d/2+1)\ell)^{(d/2-1)\ell}} n^{(-d/2+1)\ell} \quad (\text{A.10h})$$

$$\leq \sum_{\ell=1}^r n^{\left(\frac{\log e}{\log n} d + \frac{1}{\log n} \log(dr)(d/2-1) - d/2 + 1 + \delta d \right) \ell} \quad (\text{A.10i})$$

$$= \sum_{\ell=1}^r n^{\left(\frac{\log e}{\log n} + \frac{1}{\log n} \log(dr)(1/2-1/d) - 1/2 + 1/d + \delta \right) d \ell}, \quad (\text{A.10j})$$

where to get from line (A.10d) to (A.10e) we used that $\binom{n}{k} \leq \left(\frac{en}{k}\right)^k$ for the Euler number e , from (A.10g) to (A.10h) we used that $n^{\log a / \log n} = a$, and from (A.10h) to (A.10i) that $d \geq d_0 \geq 2$ and $\ell \leq r$. In order to show that (A.10j) is bounded away from 1, it suffices to demonstrate that the expression

$$\frac{\log e}{\log n} + \frac{1}{\log n} \log(dr)(1/2 - 1/d) - 1/2 + 1/d + \delta \quad (\text{A.11})$$

A Appendix

is negative and bounded away from zero. Set $\lambda = \varepsilon/2 - 1/d_0 - \delta > 0$ and choose $n_0 = 3^{2/\lambda}$. By the upper bounds on r and d it follows that

$$\log e / \log n + \log(dr)(1/2 - 1/d) / \log n - 1/2 + 1/d + \delta \quad (\text{A.12a})$$

$$\leq \log e / \log n + (1/2) \log(n^{\frac{1}{2}-\varepsilon} n^{\frac{1}{2}}) / \log n - 1/2 + 1/d + \delta \quad (\text{A.12b})$$

$$= \log e / \log n - \varepsilon/2 + 1/d + \delta \quad (\text{A.12c})$$

$$\leq \log e / \log n - \varepsilon/2 + 1/d_0 + \delta \quad (\text{A.12d})$$

$$= \log e / \log n - \lambda \quad (\text{A.12e})$$

$$\leq -\lambda/2, \quad (\text{A.12f})$$

where the last inequality holds since $n \geq n_0 > e^{2/\lambda}$. It follows that the probability that \mathcal{G} is not an $(r, d, 2)$ -boundary expander is bounded by

$$\sum_{\ell=1}^r n^{(-\lambda/2)d\ell} \leq \sum_{\ell=1}^r n_0^{(-\lambda/2)d\ell} \leq \sum_{\ell=1}^{\infty} \left(\frac{1}{3}\right)^{d\ell} \leq \frac{1}{2}, \quad (\text{A.13})$$

which establishes the lemma. □