

Supercritical Space-Width Trade-Offs for Resolution

Christoph Berkholz*¹ and Jakob Nordström^{†2}

1 Humboldt-Universität zu Berlin, Berlin, Germany

2 KTH Royal Institute of Technology, Stockholm, Sweden

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 resolution proof must have space exceeding by far the linear worst-case upper bound. This significantly strengthens the space-width trade-offs in [Ben-Sasson 2009], and provides one more example of trade-offs in the “supercritical” regime above worst case recently identified by [Razborov 2016]. 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 2008].

1998 ACM Subject Classification F.2.2 Analysis of Algorithms and Problem Complexity: Non-numerical Algorithms and Problems – Complexity of proof procedures, F.1.3 Computation by Abstract Devices: Complexity Measures and Classes, I.2.3 Artificial Intelligence: Deduction and Theorem Proving, F.4.1 Mathematical Logic and Formal Languages: Mathematical Logic – computational logic

Keywords and phrases Proof complexity, resolution, space, width, trade-offs, supercritical

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.57

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 [20] as a way of attacking the problem of showing that $NP \neq coNP$, and hence $P \neq NP$, 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.

Resolution Length, Width, and Space. Arguably the most thoroughly studied proof system in proof complexity is *resolution*, which appeared in [15] and began to be investigated in connection with automated theorem proving in the 1960s [21, 22, 33]. Because of its simplicity 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 [4, 26, 27].

* 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.



© Christoph Berkholz and Jakob Nordström;
licensed under Creative Commons License CC-BY

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).

Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi;
Article No. 57; pp. 57:1–57:14



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



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 a breakthrough result, Haken [24] obtained a length lower bound on the form $\exp(\Omega(n^\delta))$ for general resolution refutations of so-called pigeonhole principle formulas, and this paper was later followed by truly exponential lower bounds $\exp(\Omega(n))$ for other formula families in, e.g., [6, 18, 35].

In a seminal paper [12], 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 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 [3]). 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. 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 [23], 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 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 [23], 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 [1, 9, 23].

These space lower bounds also matched known lower bounds on width, and in a strikingly simple and beautiful result Atserias and Dalmau [2] showed that in fact the resolution width of refuting a k -CNF formula F provides a lower bound for the clause space required.² This allows to recover the space lower bounds mentioned above as immediate consequences of width lower bounds shown in [12]. Furthermore, it follows from [2] 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. In the sequence of papers [28, 30, 10] it was shown, however, that there are formula families that have high space complexity although they have refutations in linear length and constant width.

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, given a formula F , whether there exists a single refutation that can simultaneously optimize these different complexity measures. This question was first raised by Ben-Sasson [8], who gave a strong negative answer for space versus width. He showed 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.

¹ For 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 [1, 13, 16, 17]. In this paper, however, “space” will always mean “clause space” in the sense of [23] unless otherwise stated.

² Note that this is a nontrivial connection since 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.

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 [8] Ben-Sasson also proved a size-space trade-off for tree-like resolution, and building on [8, 10] it was shown in [11] 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 in general resolution. Beame et al. [5] and Beck et al. [7] exhibited formulas over n variables refutable in length polynomial in n where bringing the space down to linear, or even just shaving a constant factor of the polynomial space bound that follows immediately from the length bound, incurs a superpolynomial penalty in proof length.

Regarding length versus width, what was shown in [12] is that a short refutation can be converted to a refutation of small width, but this conversion blows up the length exponentially. Thapen [34] 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 [32] 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 highlight an intriguing property of the trade-off results in [5, 7, 32] that sets them 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 [5, 7] require space even polynomially larger than the worst-case upper bound, and the small-width tree-like proofs in [32] require proofs of length exponential in 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 [31] of [32], although in this paper we will instead use the term *supercritical trade-offs* to indicate that one of the complexity measures is pushed up into the supercritical regime above worst case when the other measure is decreased.

Our Contribution. Answering Razborov’s call in [32] for more examples of the type of trade-offs discussed above, 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 width $\ell(n)$ but for which any refutation in width $o(\ell(n) \log n)$ requires space $n^{\Omega(\ell(n))}$.*

Techniques. In one sentence, we obtain our results by using Razborov’s hardness condensation technique in [32] and combining it with the space lower bounds in [10].

In slightly more detail, our starting point are the so-called *pebbling formulas* defined in [12]. These formulas are refutable in constant width, but it was observed in [8] 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 [10] 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 [10] 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 [32] 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 [10] no longer works, since it crucially depends on the fact that all new substitution variables are distinct. What Razborov showed in [32] 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 [8] 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 for general, 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 $\Omega(N/\log N)$, to which we apply w -wise XORification using a much smaller set of n variables. We 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 [32] 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 [32], but with general, DAG-like resolution. We share with [32] 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 [8], 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.

Organization of This Paper. 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. Due to space constraints, we omit some of the proofs in this extended abstract, referring the reader to the upcoming full-length version for the missing details.

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. 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. A k -*CNF formula* is a CNF formula consisting 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 .

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 $W(\pi)$ 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 refutations, 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 or weakening 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 they 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 fact, it will be convenient for us to limit our attention to a (slightly non-standard) restricted form of resolution refutations as described next. Let us say that a *homogeneous resolution refutation* is 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

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

The XORification of a CNF formula F is the conjunction of all the formulas corresponding to the XORified clauses of F . We hope that the reader excuses our slightly informal definition by example and has no problems generalizing it to substitutions with XOR of arbitrary arity (but see, e.g., Definition 2.12 in [29] for a more rigorous treatment).

Usually, the way XORification is done is that any two variables in the original formula are replaced by exclusive ors over disjoint sets of new variables. Razborov [32] 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 [32]). 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), 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 conclude this section with two simple observations that will 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 Theorem 2 in [11] (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 [23]. 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 k -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 $w \leq \frac{\ell}{4k} \log n$ 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** ([10]). *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 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 a resolution refutation 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 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 yields $\delta > 0$ 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 the bipartite graph 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 of $(s - w - 3)2^{-w}$, 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 $w \leq \min(\frac{1}{4k}, \frac{\delta}{2}) \cdot \ell(n) \log n$ (recall that $w = o(\ell(n) \log n)$ by assumption), the sequence of calculations

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

yields the desired space lower bound. ◀

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 the following somewhat intriguing observation. As discussed in the introduction, Theorem 3.2 is shown by using so-called pebbling formulas. 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* $Peb_{\mathcal{D}}$ to consist of the following clauses:

- for all $s \in S$, a 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 weak 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 stronger 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.

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 [8] 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 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 prove 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 k -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:

- $F[\mathcal{G}]$ can be refuted in width ℓ .
- Any refutation $\pi : F[\mathcal{G}] \vdash \perp$ of the XORified formula $F[\mathcal{G}]$ in width $w \leq \frac{\ell}{4k} \log n$ 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 [32], 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 of them that we need. Most of our discussion of boundary expanders can be recovered from [32], but since our setting of parameters is slightly different we give a self-contained presentation below. We refer to the full-length version of this paper for any missing proofs.

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 $\mathcal{N}(U') = \{v \mid \{u, v\} \in E(\mathcal{G}), u \in U'\}$ to denote the set of right neighbours of a left vertex subset $U' \subseteq U$ (and vice versa for right vertex subsets).

► **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 $\partial(U') = \{v \in \mathcal{N}(U') : |\mathcal{N}(v) \cap U'| = 1\}$ 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 $|\mathcal{N}(u)| \leq d$ for all $u \in U$, i.e., the left degree is bounded by d .

For a right vertex subset $V' \subseteq V$ in $\mathcal{G} = (U \dot{\cup} V, E)$ we define the *kernel* $\text{Ker}(V') \subseteq U$ of V' to be the set of all left vertices whose entire neighbourhood is contained in V' , i.e., $\text{Ker}(V') = \{u \in U \mid \mathcal{N}(u) \subseteq V'\}$. We write $\mathcal{G} \setminus V'$ to denote the subgraph of \mathcal{G} induced on $(U \setminus \text{Ker}(V')) \dot{\cup} (V \setminus V')$. That is, we obtain $\mathcal{G} \setminus V'$ from \mathcal{G} by first deleting V' and afterwards all isolated vertices from U .

A 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 [32]. We omit the proof due to space constraints.

► **Lemma 4.3** ([32]). *Let \mathcal{G} be an $(r, 2)$ -boundary expander. Then for every $V' \subseteq V$ with $|V'| \leq r/2$ there exists a subset $\gamma(V') \supseteq V'$ such that $|\text{Ker}(\gamma(V'))| \leq |V'|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(r/2, 1)$ -boundary expander.*

The next lemma states that there exists $N \times n$ $(r, d, 2)$ -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. This can be proven by a standard application of the probabilistic method.

► **Lemma 4.4.** *Fix constants $\varepsilon, \delta, d_0 > 0$ 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 $n \geq n_0$, $d_0 \leq d \leq n^{\frac{1}{2}-\varepsilon}$, and $r \leq d \log n$ 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 proofs 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 4.5.** *Let F be a 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 $\text{Vars}(C) \subseteq V$. For convenience, we will overload notation and write $\text{Ker}(C) = \text{Ker}(\text{Vars}(C))$, 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 4.3 (such closures exist since all clauses $C \in \pi$ have width at most w). 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

$$\mathcal{G}^{-1}(C) := \{D \mid \text{Vars}(D) = \text{Ker}(\gamma(C)); D[\mathcal{G}] \text{ and } C \text{ are simultaneously falsifiable}\} \quad (4.1)$$

and defining

$$\mathbb{D}_t := \bigcup_{C \in \mathbb{C}_t} \mathcal{G}^{-1}(C) \quad (4.2)$$

(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 $|\text{Vars}(D)| = |\text{Ker}(\gamma(C))| \leq W(C) \leq w$, where the first inequality is guaranteed by Lemma 4.3. 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 $|\mathbb{D}_t| \leq 2^w Sp(\pi)$. 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 show that this is so, we first observe that for $\mathbb{C}_0 = \emptyset$ we obviously get $\mathbb{D}_0 = \emptyset$ by (4.2). 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

$(\mathbb{D}_t = \mathbb{D}_t^0, \mathbb{D}_t^1, \dots, \mathbb{D}_t^j = \mathbb{D}_{t+1})$ where the space of every intermediate configuration is bounded by $\max\{Sp(\mathbb{D}_t), Sp(\mathbb{D}_{t+1}) + w\}$.

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 (4.1) that $\text{Vars}(D) = \text{Ker}(\gamma(C)) \supseteq \text{Ker}(C) \supseteq \text{Vars}(A)$ 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 $\bar{x}[\mathcal{G}]$. But 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, 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 check 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. Here graph expansion comes heavily into play, but due to space constraints we have to defer the proofs to the full-length version of this paper.

► **Claim 4.6.** *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 4.7.** *For any two clauses $B \subseteq C$ 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 $w + 1$.*

Because the depth of a refutation is 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 $w + 1$. It follows that F has a homogeneous resolution refutation π' of width w and clause space $Sp(\pi') \leq 2^w Sp(\pi) + w + 3$.

This concludes the proof of Lemma 4.5. ◀

We can now combine the construction in Lemma 4.5 with the existence of good boundary expanders in Lemma 4.4 to prove the hardness condensation in Lemma 4.1.

Proof of Lemma 4.1. Given $\varepsilon > 0$ and $k \in \mathbb{N}^+$ we choose $\delta := \frac{\varepsilon}{10k}$. Suppose ℓ and n are parameters such that $k \leq \ell \leq n^{\frac{1}{2}-\varepsilon}$ and let F be an unsatisfiable k -CNF formula over $N = \lfloor n^{\delta\ell} \rfloor$ variables that can be refuted in width k . To apply Lemma 4.4 we set $d_0 := \frac{5}{\varepsilon}$ 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 $d := \frac{\ell}{2k}$ and set the expansion guarantee to $r := d \log n$. By the bound on ℓ we have $d \leq \ell \leq n^{\frac{1}{2}-\varepsilon}$.

Now we have two cases. The first, and interesting, case is when $d \geq d_0$ holds. Then Lemma 4.4 guarantees that there exists an $N \times n$ $(r, d, 2)$ -boundary expander \mathcal{G} . Applying XORification with respect to \mathcal{G} , we obtain an ℓ -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 $w \leq \frac{\ell}{4k} \log n = r/2$ the space lower bound follows from Lemma 4.5.

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 $N = \lfloor n^{\delta \ell} \rfloor = \lfloor n^{\frac{\epsilon}{10k} 2kd} \rfloor \leq \lfloor n^{\frac{\epsilon}{10k} 2kd_0} \rfloor = \lfloor n^{\frac{\epsilon}{10k} 2k \frac{5}{\epsilon}} \rfloor = n$. Thus, 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 $Sp(\pi) \geq s \geq (s - w - 3)2^{-w}$. \blacktriangleleft

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 $o(w \log n)$. Measured in terms of the number of variables n , this is a major improvement over the previous space-width trade-off result in [8], and provides another example of trade-offs in the supercritical regime above worst-case recently identified in [32].

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 $n + 2$. 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 [32], suggesting that the above goal might not be completely out of reach.

Another intriguing open problem from the complexity-theoretic point of view 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 [25] 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 [14]. 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 [1, 19].

Acknowledgements. We wish to thank Alexander Razborov for patiently explaining the hardness condensation technique in [32] during numerous and detailed discussions.

References

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

- 2 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*.
- 3 Albert Atserias, Massimo Lauria, and Jakob Nordström. Narrow proofs may be maximally long. In *Proceedings of the 29th Annual IEEE Conference on Computational Complexity (CCC'14)*, pages 286–297, June 2014.
- 4 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.
- 5 Paul Beame, Chris Beck, and Russell Impagliazzo. Time-space tradeoffs in resolution: Superpolynomial lower bounds for superlinear space. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC'12)*, pages 213–232, May 2012.
- 6 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*.
- 7 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.
- 8 Eli Ben-Sasson. Size-space tradeoffs for resolution. *SIAM Journal on Computing*, 38(6):2511–2525, May 2009. Preliminary version in *STOC'02*.
- 9 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*.
- 10 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.
- 11 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.
- 12 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*.
- 13 Patrick Bennett, Ilario Bonacina, Nicola Galesi, Tony Huynh, Mike Molloy, and Paul Wolan. Space proof complexity for random 3-CNFs. Technical Report 1503.01613, arXiv.org, April 2015.
- 14 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.
- 15 Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.
- 16 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)*, July 2016. To appear.
- 17 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.
- 18 Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988.
- 19 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.

- 20 Stephen A. Cook and Robert Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, March 1979.
- 21 Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem proving. *Communications of the ACM*, 5(7):394–397, July 1962.
- 22 Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.
- 23 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*.
- 24 Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985.
- 25 Alexander Hertel. *Applications of Games to Propositional Proof Complexity*. PhD thesis, University of Toronto, May 2008. Available at <http://www.cs.utoronto.ca/~ahertel/>.
- 26 João P. Marques-Silva and Kareem A. Sakallah. GRASP: A search algorithm for propositional satisfiability. *IEEE Transactions on Computers*, 48(5):506–521, May 1999. Preliminary version in *ICCAD'96*.
- 27 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.
- 28 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*.
- 29 Jakob Nordström. Pebble games, proof complexity and time-space trade-offs. *Logical Methods in Computer Science*, 9:15:1–15:63, September 2013.
- 30 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*.
- 31 Alexander A. Razborov. An ultimate trade-off in propositional proof complexity. Technical Report TR15-033, Electronic Colloquium on Computational Complexity (ECCC), March 2015.
- 32 Alexander A. Razborov. A new kind of tradeoffs in propositional proof complexity. *Journal of the ACM*, 63:16:1–16:14, April 2016.
- 33 John Alan Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, January 1965.
- 34 Neil Thapen. A trade-off between length and width in resolution. Technical Report TR14-137, Electronic Colloquium on Computational Complexity (ECCC), October 2014.
- 35 Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, January 1987.