Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 33 (2015)

An Ultimate Tradeoff in Propositional Proof Complexity

Alexander Razborov*

April 21, 2015

Abstract

We exhibit an unusually strong tradeoff result in propositional proof complexity that significantly deviates from the established pattern of almost all results of this kind. Namely, restrictions on one resource (width in our case) imply an increase in another resource (tree-like size) that is exponential not only with respect to the complexity of the original problem, but also to the whole class of *all* problems of the same bit size. More specifically, we show that for any parameter $k = k(n)$ there are unsatisfiable k-CNFs that possess refutations of width $O(k)$, but such that any tree-like refutation of width $n^{1-\epsilon}/k$ must necessarily have *double* exponential size $\exp(n^{\Omega(k)})$. This means that there exist contradictions that allow narrow refutations, but in order to keep the size of such a refutation even within a single exponent, it must necessarily use a high degree of parallelism.

Our construction and proof methods combine, in a non-trivial way, two previously known techniques: the hardness escalation method based on substitution formulas and expansion. This combination results in a *hardness compression* approach that strives to preserve hardness of a contradiction while significantly decreasing the number of its variables.

*University of Chicago, razborov@math.uchicago.edu; part of this work was done at Steklov Mathematical Institute, Moscow, and at Toyota Technological Institute, Chicago.
Supported by the Russian Science Foundation, grant # 14-50-00005.

ISSN 1433-8092

1. Introduction

Tradeoff results is quite a popular topic in complexity theory and beyond. Whenever they apply, this serves as a rigorous demonstration of our inherent inability to achieve two conflicting goals at once. These results have the following form that we purposely describe in very generic terms. We are given a task T and a set of protocols $\mathcal{P}_T$ that achieve this task; for simplicity, we confine ourselves to non-uniform models when everything is finite. The set $\mathcal{P}_T$ is equipped with two complexity measures μ and ν, usually of a rather different nature. Then tradeoff results claim that after we restrict $\mathcal{P}_T$ to those protocols P for which $\mu(P)$ is small, the minimum complexity with respect to the second measure, $\min_{P\in\mathcal{P}_T} \nu(P)$ increases drastically, in many cases exponentially. Earlier results along these lines primarily focussed on the pair $\mu =$ "computational space", $\nu =$ "computational time", see [Bor93] for a survey. But, as we already mentioned, at the moment this paradigm is omnipresent and (μ, ν) can be literally anything. Specifically, there are many tradeoff results of all kind in proof complexity [Nor13].

In this paper we demonstrate an example of a tradeoff whose behavior is drastically different from most known cases. After a preliminary version of this paper was disseminated, I was urged by several colleagues to articulate the difference more clearly, and we will now spend some time on it. The reader primarily interested in our technical contribution may safely skip this part and jump to page 4.

It would be instructive to resort to a simple and necessarily schematic picture. On Figure 1, $\mu_{\min}$ is simply $\min_{P\in\mathcal{P}_T} \mu(P)$, that is the complexity of our task T with respect to μ alone. $\mu_{\max}$ is the "saturation level" in the resource μ, that is its maximal amount any "reasonable" protocol can possibly consume. The right end $\mu_{\max}$ of the interval of interest is actually not very important for making our point, whenever it is not clear from the context the reader can assume instead $\mu_{\max} = \infty$.

Tradeoff results of any kind are trying to pinpoint the behavior of the function

$$f(\mu) \stackrel{\text{def}}{=} \min\{\nu(P) \mid P \in \mathcal{P}_T \wedge \mu(P) \leq \mu\}$$

in the interval $\mu \in [\mu_{\min}, \mu_{\max}]$; more specifically, prove that it is (sharply) decreasing. Ordinary tradeoffs normally consist of two parts: a *lower bound* on f, shown in red, along with an *upper bound* stating, as the very minimum, that at the right end $\mu_{\max}$ this lower bound is close to tight. These two facts

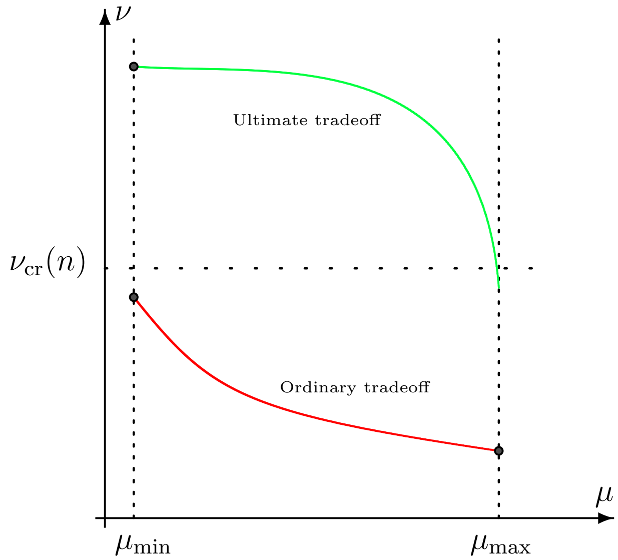

Figure 1: Ordinary and ultimate tradeoffs

together imply that the actual curve $f(\mu)$ does have a slope. We would like to note that we drew the red curve as a nice convex function mostly as a sign of respect to earlier results that almost always had the form of a lower bound on $\mu(P)\nu(P)$. In the plethora of examples found since that, it can be anything non-increasing: we will see below one example (5) that is a lower bound on $\mu(P)\log\nu(P)$; it can be a step function when the lower bound part simply breaks down for some $\mu \in [\mu_{\min}, \mu_{\max}]$ etc. But, repeating our point, two ingredients are paramount to any tradeoff result: a lower bound on $f(\mu)$ at $\mu = \mu_{\min}$ (and, preferably, as much to the right of it as possible), as well as an upper bound at $\mu = \mu_{\max}$ (extended as much as possible to the left).

So far we have discussed one particular task, T. But since the complexity theory is about algorithmic problems, T is never in isolation, it is always a member of a large but finite[1] family of tasks $\mathcal{T}_n$, where the subscript n stands for the "size" of T. As it turns out, in virtually all cases of interest there exists a "natural" upper bound $\nu_{\mathrm{cr}}(n)$ on the ν-complexity $\min_{P\in\mathcal{P}_{T^*}} \nu(P)$ of *any* task $T^* \in \mathcal{T}_n$, usually provided by a trivial (aka "straightforward", "brute-force" etc.) protocol. For example, when $\nu =$ "circuit size" in circuit complexity or $\nu =$ "proof length" in proof complexity, $\nu_{\mathrm{cr}}(n) = 2^n$. In

[1]Recall that we are considering non-uniform models.

communication complexity, $\nu_{\mathrm{cr}}(n) = n$, n the bit size of the input. Etc. In some cases, like circuit complexity, this upper bound is known to be tight for the worst or even typical task $T^* \in \mathcal{T}_n$ (so-called *Shannon effect*), and in other cases, like strong propositional proof systems, this is wide open. But, implicitly or explicitly, $\nu_{\mathrm{cr}}(n)$ is always there and determines the range in which interesting developments occur.

As it turns out, in the vast majority of known tradeoff results, the red curve is *below* the critical horizontal line $\nu = \nu_{\mathrm{cr}}(n)$. That is, μ-restricted protocols for T are compared with unrestricted protocols P for the task T itself; they can be managed to be at most as ν-bad as *trivial* protocols for *all* tasks $T^* \in \mathcal{T}_n$.

In this paper we exhibit a tradeoff result that operates entirely *above* the line $\nu = \nu_{\mathrm{cr}}(n)$. In other words, the lower bound in a vicinity of $\mu_{\min}$ (that in our result actually extends almost all the way up to $\mu_{\max}$) is so strong that it exponentially beats the generic upper bound $\nu_{\mathrm{cr}}(n)$ holding for all tasks T^* of comparable size n. The second ingredient in ordinary tradeoff results, which is the upper bound on ν at $\mu = \mu_{\max}$ *need not be proven separately*: it automatically follows from the membership $T \in \mathcal{T}_n$. It is due to this latter property that we propose to call tradeoffs of this kind *ultimate*. This phenomenon seems to be extremely rare (we will review below those few exceptions we have been able to find after an earlier version of this paper was disseminated), and we think it deserves at least a certain amount of attention and contemplation.

Our concrete result belongs to the area of propositional proof complexity that has seen a rapid development since its inception in the seminal paper by Cook and Reckhoff [CR79]. This success is in part due to being well-connected to a number of other disciplines, and one of these connections that has seen a particularly steady growth in very recent years is the interplay between propositional proof complexity and practical SAT solving. As a matter of fact, SAT solvers that seem to completely dominate the landscape at the moment, like those based on conflict-driven clause learning, lead to just one *resolution* proof system dating back to the papers by Blake [Bla37] and [Rob65]. This somewhat explains the fact that resolution is by far the most studied system in proof complexity, and much of this study has concentrated on simple complexity measures for resolution proofs like size, width and space, and on relations existing between them.

Our paper also exclusively deals with the resolution proof system. The first measure of interest to us is *width*. This measure is extremely natural and robust, and in fact it is not very specific to resolution. As is well-known, width w proofs can more instructively be viewed as semantic proofs operating with arbitrary Boolean expressions and equally arbitrary (sound) inference rules with the sole restriction that every line depends on at most w variables.

Ben-Sasson and Wigderson [BW01] showed that short proofs can be transformed into proofs of small width (see (1) and (2) below), while Atserias and Dalmau [AD08] did this for proofs that have small clause space. Thus, despite its deluding simplicity, the class of contradictions possessing small-width refutations is rich.

In this paper we give (yet another) confirmation of this thesis "from the opposite side": there exist contradictions that do have small-width refutations, but the latter are highly complex and non-efficient, and any attempts to simplify them must necessarily lead to a dramatic blow-up in width. Before reviewing some relevant work in this direction and stating our own contribution, it will be convenient to fix some basic notation (we will remind exact definitions in Section 2). For a CNF contradiction τ_n in n variables, let $w(\tau_n \vdash 0)$ $[S(\tau_n \vdash 0)]$ be the minimum possible width [size, respectively] of any resolution refutation of τ_n. $S_T(\tau_n \vdash 0)$ is the minimum size with respect to tree-like refutations, and $w(\tau_n)$ is the maximum width of a clause in τ_n itself.

In this notation, the main results from [BW01] can be stated as follows:

$$w(\tau_n \vdash 0) \leq O(\log S_T(\tau_n \vdash 0)); \quad (1)$$

$$w(\tau_n \vdash 0) \leq O(n \cdot \log S(\tau_n \vdash 0))^{1/2} + w(\tau_n). \quad (2)$$

Bonet and Galesi [BG99] proved that (2) is (almost) tight by exhibiting contradictions τ_n such that $w(\tau_n) \leq O(1)$, $S(\tau_n \vdash 0) \leq n^{O(1)}$ while $w(\tau_n \vdash 0) \geq \Omega(n^{1/2})$. As for (1), it is tight for obvious reasons: the *Complete Tree contradiction* CT_n consisting of all 2^n possible clauses in n variables satisfies $w(CT_n) = w(CT_n \vdash 0) = n$ and $S_T(CT_n \vdash 0) = 2^n$.

In the opposite direction, resolution size can be trivially bounded by width as follows:

$$S(\tau_n \vdash 0) \leq n^{O(w(\tau_n \vdash 0))}; \quad (3)$$

the right-hand side here simply bounds the overall number of all possible clauses of width $\leq w$.

Atserias, Lauria and Nordström [ALN14] have recently shown that this bound is also tight: for an arbitrary $w = w(n) \leq n^{1/2-\Omega(1)}$ there exist contradictions τ_n with $w(\tau_n \vdash 0) \leq w$ while $S(\tau_n \vdash 0) \geq n^{\Omega(w)}$. An earlier result by Ben-Sasson, Impagliazzo and Wigderson [BIW04] can be viewed as an ultimate demonstration that no simulation like (3) is possible for *tree-like* resolution. Namely, they gave an example of contradictions τ_n such that

$$w(\tau_n \vdash 0) \leq O(1), \quad S_T(\tau_n \vdash 0) \geq \exp\left(\Omega(n/\log n)\right). \tag{4}$$

Having briefly discussed simulation and separation results, let us review what has been known before in terms of tradeoffs. Ben-Sasson [Ben09] established a tradeoff between width and tree-like resolution size. Namely, he constructed contradictions τ_n that have tree-like refutations of *either* constant width *or* polynomial size but such that

$$w(\Pi) \cdot \log |\Pi| \geq \Omega(n/\log n) \tag{5}$$

for any tree-like refutation Π of τ_n ($w(\Pi)$ and $|\Pi|$ being its width and size respectively). As for the general case, strong tradeoffs are precluded by (3) and the observation that the naive resolution refutation it represents still has the minimum possible width $w(\tau_n \vdash 0)$. Hence, in sharp contrast with (5), every contradiction τ_n has a DAG-like refutation Π such that

$$w(\Pi) \cdot \log |\Pi| \leq O(\log n \cdot w(\tau_n \vdash 0)^2).$$

Even with this severe restriction interesting results along these lines have been reported in [Nor09, Tha14]. They are, however, best viewed in the dual coordinate system $\mu =$ "size", $\nu =$ "width" (in terms of Fugure 1).

Our main result, Theorem 2.4 is a far-reaching generalization of the previous contributions (4), (5). For any parameter $k = k(n)$ we construct a sequence of k-CNF contradictions τ_n such that $w(\tau_n \vdash 0) \leq O(k)$ while

$$|\Pi| \geq \exp\left(n^{\Omega(k)}\right) \tag{6}$$

for any tree-like refutation Π of width $\leq n^{1-\epsilon}/k$. Thus, when k, say, is a sufficiently large constant our bound becomes *super-exponential* in n, and for (say) $k = n^{1/3}$ it becomes *double exponential.* As such, it perfectly fits the paradigm we described in the beginning: in terms of Figure 1, we have $\mu =$ "width", $\nu =$ "tree-like size", $\mu_{\min} = O(k)$, $\mu_{\max} = n$ and $\nu_{\mathrm{cr}}(n) = 2^n$. On less general level, our result is complementary to that of Atserias et al. [ALN14]. Namely, they proved that the obvious brute-search refutation of size $n^{O(w)}$ (cf. (3)) in general can not be shortened. What we prove is that if we additionally want to keep the width reasonably small, and keep the size sane (at most single exponential), we need a high degree of parallelism.

Let us now review a few previous examples that are relevant to this framework.

Resolution proofs, $\mu =$ "width", $\nu =$ "logical depth". Then $\mu_{\max} = n$ and $\nu_{\mathrm{cr}} = 2^n$. Let us choose $\mu_{\min}$ arbitrarily. The proof method of a result by Berkholz [Ber12, Theorem 5] gives an ultimate trade-off similar to ours: every refutation of the minimum width k must have depth at least n^k. This lower bound, however, breaks down already for refutations of width $(k+1)$.

Resolution proofs (and beyond), $\mu =$ "proof length", $\nu =$ "clause space". We have $\mu_{\max} = 2^n$, $\nu_{\mathrm{cr}}(n) = n$. For any $\mu_{\min} \in \left[n^{C\log n}, 2^{n^{\epsilon}}\right]$, the result by Beame, Beck and Impagliazzo [BBI12] gives contradictions τ_n that have refutations of length $\mu_{\min}$ but such that every refutation of length $\leq \mu_{\min}{}^{\frac{\epsilon \log\log n}{\log\log\log n}}$ must have clause space $\mu_{\min}^{\Omega(1)}$. Beck, Nordström and Tang [BNT13, Theorem 4] generalized this result to the polynomial calculus with resolution and also pushed down the lower bound on $\mu_{\min}$ to n^C.

It is worth noting, however, that in the dual regime $\mu =$ "clause space", $\nu =$ "proof length" that is, perhaps, more natural, ultimate tradeoffs are hindered by the observation that every clause space S refutation must necessarily be of length $\exp(O(Sn))$.

Information Complexity. Here $\mu =$ "information complexity", $\nu =$ "communication complexity". This example is different from the previous ones in several important respects. Firstly, the bit size of the problem is deemed to be totally irrelevant, and, in fact, can be thought of as the information complexity itself. Second, the generic upper bound ν_{cr} was proved by Braverman [Bra12]: $CC(T) \leq 2^{IC(T)}$, and it is anything but trivial. Third, and most important, this is an interesting conjecture rather than a result. Namely, separation between information and communication complexities was recently established by Ganor, Kol and Raz in [GKR14]. They conjecture (personal communication) that the same example T provides an ultimate tradeoff between the two complexities, that is every protocol for T of nearly optimal information complexity must be of length that is double exponential in $IC(T)$.

Our construction and the proof combine two very popular techniques in proof complexity: hardness escalation and expansion. The former method converts every contradiction τ_n into another contradiction $\widehat{\tau}_n$ so that relatively mild hardness properties of τ_n transfer to lower bounds for $\widehat{\tau}_n$ in stronger proof systems. So far this technique has been used in two main flavors: *substitution formulas* (see e.g. [Nor13, Section 2.4]) and *lifting formulas* introduced in [BHP10]. One common feature of both approaches is that the price one has to pay for improving hardness is a moderate *increase* in the number of variables.

In our work we change the gears on both these counts and are interested in hardness *preservation*[2] and *variable compression*, that is in (exponentially) decreasing the number of variables. These two conflicting goals are balanced using linear substitutions whose support sets need not necessarily be disjoint as long as they form a good expander. While by now expanders is one of the most common techniques in proof complexity, we are not aware of their previous applications in a similar context.

2. Preliminaries

In this section we give necessary definitions, state some useful facts and formulate, in Section 2.1, our main results.

A *literal* is either a Boolean variable x or its negation $\bar{x}$; we will sometimes use the uniform notation $x^\epsilon \stackrel{\text{def}}{=} \begin{cases} x & \text{if } \epsilon = 1 \\ \bar{x} & \text{if } \epsilon = 0. \end{cases}$. A *clause* C is either a disjunction of literals in which no variable appears along with its negation, or 1. The latter is a convenient technicality (e.g. with this convention the set of all clauses makes a lattice in which $\vee$ is the join operator etc.); 1 should be thought of as a placeholder for all trivially true clauses. C is a *sub-clause of* D, also denoted by $C \leq D$ if either $D = 1$ or $C, D \neq 1$ and every literal appearing in C also appears in D. Two clauses C and D are *consistent* if $C \vee D \neq 1$, that is both C and D are non-trivial and do not

[2] As a matter of fact, our construction also gives the same hardness amplification as ordinary substitution formulas with disjoint sets of variables. This observation, however, plays no role in our conclusions.

contain conflicting literals. The empty clause C will be denoted by 0. The set of variables occurring in a clause C (either positively or negatively) will be denoted by $Vars(C)$ ($Vars(1) \stackrel{\rm def}{=} \emptyset$). The *width* of a clause is defined as $w(C) \stackrel{\rm def}{=} |Vars(C)|$.

A *CNF* τ is a conjunction of clauses, often identified with the set of clauses it is comprised of. A CNF is a k-*CNF* if all clauses in it have width at most k. Unsatisfiable CNFs are traditionally called *contradictions*. For CNFs τ, τ', $\tau \models \tau'$ is the *semantical implication* meaning that every truth assignment satisfying τ also satisfies τ'. Thus, τ is a contradiction if and only if $\tau \models 0$. Also, for clauses C and D, $C \leq D$ if and only if $C \models D$. The subscript n in τ_n always stands for the number of variables in the CNF τ_n.

The *resolution proof system* operates with clauses and it consists of the only *resolution rule*

$$\frac{C \vee x \qquad D \vee \bar{x}}{C \vee D}. \tag{7}$$

A *tree-like*[3] *resolution proof* Π is a binary rooted tree in which all nodes all labelled by clauses, and such that the clause assigned to every internal node can be deduced from clauses sitting at its two children via a single application of the resolution rule. A *tree-like resolution proof of a clause C from a CNF τ* is a tree-like resolution proof Π in which all leaves are labelled by clauses from τ, and the root is labelled by a clause $\widetilde{C}$ such that $\widetilde{C} \leq C$ (the latter technicality is necessary since we did not include the weakening rule). A *refutation* of a contradiction is a proof of 0 from it. The *depth* $D(\Pi)$ of a proof Π is the height (the number of edges in the longest path) of its underlying tree, and its *size* $|\Pi|$ is the number of leaves. The *width* $w(\Pi)$ is the maximum width of a clause appearing in Π.

For a CNF τ and a clause C, we let $D(\tau \vdash C)$, $S_T(\tau \vdash C)$ and $w(\tau \vdash C)$ denote the minimum possible value of $D(\Pi)$, $|\Pi|$ and $w(\Pi)$, respectively, taken over all tree-like resolution proofs Π of C from τ (if $\tau \not\models C$, we let all three measures be equal to ∞).

The following result will be one of the starting points for our construction.

Proposition 2.1 ([BIW04]) *There exists an increasing sequence $\{\tau_n\}$ of 4-CNF contradictions such that $w(\tau_n \vdash 0) \leq 6$, but $S_T(\tau_n \vdash 0) \geq \exp(\Omega(n/\log n))$.*

[3]DAG-like proofs are not considered in this paper.

Let A be a $m \times n$ 0-1 matrix. For $i \in [m]$,[4] let $J_i(A) \stackrel{\text{def}}{=} \{ j \in [n] \mid a_{ij} = 1 \}$. For a clause E in the variables $\{y_1, \ldots, y_m\}$, by $E[A]$ we will denote the CNF obtained from E by the $\mathbb{F}_2$-linear substitution $y_i \to \bigoplus_{j \in J_i(A)} x_j$ ($i \in [m]$) followed by expanding the resulted Boolean function as a CNF in the straightforward way. The following easy observation will be important in what follows: for every clause C in $E[A]$,

$$Vars(C) = \bigcup_{y_i \in Vars(E)} \{ x_j \mid j \in J_i(A) \} \tag{8}$$

(we do claim equality here). For a CNF $\tau = E_1 \wedge E_2 \wedge \ldots \wedge E_\ell$, we let $\tau[A] \stackrel{\text{def}}{=} E_1[A] \wedge \ldots \wedge E_\ell[A]$. If τ is a contradiction then evidently $\tau[A]$ is a contradiction, too. The converse need not be true in general, of course.

For $I \subseteq [m]$, the *boundary* of this set of rows is defined as

$$\partial_A(I) \stackrel{\text{def}}{=} \{ j \in [n] \mid |\{ i \in I \mid j \in J_i(A) \}| = 1 \},$$

i.e., it is the set of columns that have *precisely* one 1 at their intersections with I. A is an (r, s, c)-*boundary expander*[5] if $|J_i(A)| \leq s$ for any $i \in [m]$ and $|\partial_A(I)| \geq c|I|$ for every set of rows $I \subseteq [m]$ with $|I| \leq r$. An (r, n, c)-boundary expander (i.e., a $m \times n$ matrix satisfying only the second of these conditions) will be simply called an (r, c)-*boundary expander.*

For a set of columns $J \subseteq [n]$, we let

$$\mathrm{Ker}(J) \stackrel{\text{def}}{=} \{ i \in [m] \mid J_i(A) \subseteq J \}$$

be the set of rows completely contained in J. Let $A \setminus J$ be the sub-matrix of A obtained by removing all columns in J and all rows in $\mathrm{Ker}(J)$.

We need two properties of boundary expanders whose analogues were used, in one or another form, in almost all their applications in proof complexity. The first one, proven by a simple probabilistic argument, says that good expanders exist in the range $m \gg n$ (note that it becomes sub-optimal in the frequently used setting $s, c = O(1),\ m = O(n)$).

---

[4] $[m] \stackrel{\text{def}}{=} \{1, \ldots, m\}$.

[5] In [ABRW04] such matrices were simply called expanders. But since that the research in proof complexity also made good use of ordinary vertex expanders (see e.g. [BG03, AAT11]) and, as a consequence, tended to differentiate between boundary and ordinary expansion. Hence we also adopt this terminological change. What we, however, keep in this paper is the matrix notation as we find it more instructive for many reasons.

Lemma 2.2 *Let $n \to \infty$ and m, s, c be arbitrary integer parameters possibly depending on n such that $c \leq \frac{3}{4}s$ and*

$$r \leq o(n/s) \cdot m^{-\frac{2}{s-c}}. \tag{9}$$

Then for sufficiently large n there exist $m \times n$ (r, s, c)-boundary expanders.

The second property says that in every good expander, the class of small sets of rows whose removal leads to a relatively good expander is in a sense everywhere dense.

Lemma 2.3 *Let A be an $m \times n$ $(r, 2)$-boundary expander. Then for every $J \subseteq [n]$ with $|J| \leq r/4$ there exists $\widehat{J} \supseteq J$ such that $\left|\mathrm{Ker}\left(\widehat{J}\right)\right| \leq 2|J|$ and $A \setminus \widehat{J}$ is an $(r/2, 3/2)$-boundary expander.*

We, however, have not been able to recover these statements from the literature in a referrable form, and for this reason their simple proofs are included in the Appendix.

2.1. Main results

In this brief section we formulate our main results.

Theorem 2.4 *Let $k = k(n) \geq 4$ be any parameter, and let $\epsilon > 0$ be an arbitrary constant. Then there exists a sequence of k-CNF contradictions $\{\tau_n\}$ in n variables such that $w(\tau_n \vdash 0) \leq O(k)$ but for any tree-like refutation Π with $w(\Pi) \leq n^{1-\epsilon}/k$ we have the bound*

$$|\Pi| \geq \exp\left(n^{\Omega(k)}\right).$$

As we noted in Introduction, our main technique is hardness preservation, and since the corresponding statement might be of independent interest, we formulate it here as a separate result.

Theorem 2.5 *Let τ_m be an arbitrary contradiction in the variables $y_1, \ldots, y_m$, and let A be an $m \times n$ $(r, 2)$-boundary expander for some r. Then every tree-like refutation Π of $\tau_m[A]$ with $w(\Pi) \leq r/4$ must satisfy*

$$|\Pi| \geq 2^{2D(\tau_m \vdash 0)/r}.$$

As stated, this is also a hardness escalation result (from depth to tree-like size), but that part alone was known before [Urq11, Theorem 5.4], and one does not need expanders for that.

3. Proofs

In this section we prove Theorems 2.4 and 2.5, and we begin with the latter. We present our proof as a plain inductive argument since, in our view, it is often more instructive than various top-down approaches (cf. the recent simplification of the Atserias-Dalmau bound obtained by Flimus et al. [FLM+14] and independently by Razborov (unpublished)).

Fix an $m \times n$ $(r, 2)$-boundary expander A, where r is an arbitrary parameter. Let us say that a set of columns J is *closed* if $A \setminus J$ is an $(r/2, 3/2)$-boundary expander (cf. Lemma 2.3). Fix now an arbitrary CNF τ_m (that need not necessarily be a contradiction) in the variables $y_1, \ldots, y_m$. We are going to prove the following.

Claim 3.1 *Assume that C is a clause in the variables $x_1, \ldots, x_n$ that possesses a tree-like proof Π from $\tau[A]$ with $w(\Pi) \leq r/4$. Let $J \subseteq [n]$ be an arbitrary closed set with $J \supseteq \{j \mid x_j \in Vars(C)\}$, and let E be any clause in y-variables with*

$$Vars(E) = \{y_i \mid i \in \mathrm{Ker}(J)\}$$

such that

$$E[A] \vee C \not\equiv 1, \tag{10}$$

that is there exists an assignment of x-variables simultaneously falsifying $E[A]$ and C. Then

$$D(\tau \vdash E) \leq \frac{r}{2} \cdot \log_2 |\Pi|.$$

Proof of Claim 3.1. Let C, Π, J and E satisfy the assumptions of our claim. The argument proceeds by induction on $|\Pi|$.

Base $|\Pi| = 1$, i.e. C contains a sub-clause $\widetilde{C}$ that appears in $\widetilde{E}[A]$ for some $\widetilde{E} \in \tau$.

Applying (8) to the clause $\widetilde{E}$, we see that $J \supseteq Vars(C) \supseteq Vars(\widetilde{C})$ implies $\left\{ i \in [m] \,\middle|\, y_i \in Vars(\widetilde{E}) \right\} \subseteq \mathrm{Ker}(J)$, that is $Vars(\widetilde{E}) \subseteq Vars(E)$. Also, E and $\widetilde{E}$ must be consistent since

$$(E \vee \widetilde{E})[A] = E[A] \vee \widetilde{E}[A] \models E[A] \vee \widetilde{C} \models E[A] \vee C,$$

and hence their inconsistency would have implied that $(E \vee \widetilde{E})[A] \equiv E[A] \vee C \equiv 1$ in contradiction with (10). Hence $\widetilde{E} \leq E$ and thus $D(\tau \vdash E) = 0$.

Inductive step $|\Pi| > 1$.
Assume that the last application of the resolution rule has the form

$$\frac{C_0 \vee x_j \qquad C_1 \vee \bar{x}_j}{C_0 \vee C_1}.$$

Fix arbitrarily an assignment α to $\{x_k \mid k \in J\}$ falsifying both $E[A]$ and $C_0 \vee C_1$ that exists by our assumption. Further analysis depends on whether $j \in J$ or not.

Case 1, $j \in J$.
This case is easy. Assume w.l.o.g. that $\alpha(x_j) = 0$. Note that $Vars(C_0 \vee x_j) \subseteq J$ and $\alpha(C_0 \vee x_j) = 0$. Thus we can apply the inductive assumption to the clause $C_0 \vee x_j$, the corresponding sub-proof Π_0 of Π and to the same J and E. We conclude that $D(\tau \vdash E) \leq r \cdot \log_2 |\Pi_0| \leq r \cdot \log_2 |\Pi|$.

Case 2, $j \notin J$.
One of the two sub-trees Π_0, Π_1 (say, Π_0) determined by the children of the root has size $\leq |\Pi|/2$, and we assume w.l.o.g. that it corresponds to the child labeled by $C_0 \vee x_j$. Since $w(C_0 \vee x_j) \leq r/4$ by our assumption, we can apply Lemma 2.3 to the set $J' \stackrel{\text{def}}{=} \{j' \mid x_{j'} \in Vars(C_0 \vee x_j)\}$. This will give us a closed $\widehat{J} \supseteq \{j' \mid x_{j'} \in Vars(C_0 \vee x_j)\}$ with $\left|\text{Ker}\left(\widehat{J}\right)\right| \leq r/2$, and our goal is to prove that *every* clause $\widehat{E}$ with $Vars\left(\widehat{E}\right) = \left\{y_i \,\middle|\, i \in \text{Ker}\left(\widehat{J}\right)\right\}$ *and consistent with* E satisfies the assumptions of Claim 3.1 with $C := C_0 \vee x_j,\ J := \widehat{J}$ and $E := \widehat{E}$ (the rest will be easy). For that we only have to extend our original assignment α to the variables $\left\{x_j \,\middle|\, j \in J \cup \widehat{J}\right\}$ in such a way that it will falsify both $\widehat{E}[A]$ and $C_0 \vee x_j$.

Since $C_0 \leq C_0 \vee C_1 \leq C$ is already falsified by α, the latter task can be achieved simply by setting additionally $\alpha(x_j) \stackrel{\text{def}}{=} 0$ (recall that $j \notin J$). Also, every literal y_i^ϵ of a variable $y_i \in Vars(E) \cap Vars\left(\widehat{E}\right)$ maps to $y_i^\epsilon[A] = \bigoplus_{j \in J_i(A)} x_j \oplus \bar{\epsilon}$ and, since $J_i(A) \subseteq J$, it has been already decided by α. As E and $\widehat{E}$ are consistent by our assumption (and α falsifies E), $\alpha(y_i[A])$ is actually $\bar{\epsilon}$. It only remains to show that α' can be extended in such a way that it sets all $y_i[A]$ for $i \in \text{Ker}\left(\widehat{J}\right) \setminus \text{Ker}(J)$ to fixed values predetermined to falsify the formula $\widehat{E}[A]$.

Let A' be the matrix obtained from $A \setminus J$ by additionally removing the column j from it. Since $A \setminus J$ is an $(r/2, 3/2)$-boundary expander, A' is an $(r/2, 1/2)$-boundary expander. Also, $\text{Ker}\left(\widehat{J}\right) \setminus \text{Ker}(J)$ is a set of rows of cardinality $\leq r/2$, therefore $\partial_{A'}(I) \neq \emptyset$ for every non-empty subset $I \subseteq \mathrm{Ker}\left(\widehat{J}\right) \setminus \mathrm{Ker}(J)$. This allows us to order, by reverse induction, the rows in $\mathrm{Ker}\left(\widehat{J}\right) \setminus \mathrm{Ker}(J)$ in such a way $\mathrm{Ker}\left(\widehat{J}\right) \setminus \mathrm{Ker}(J) = \{i_1, \ldots, i_\ell\}$ that for every $\nu \in [\ell]$ the set of columns $J_{i_\nu}(A') \setminus \bigcup_{\mu=1}^{\nu-1} J_{i_\mu}(A')$ is not empty; fix arbitrarily $j_\nu \in J_{i_\nu}(A') \setminus \bigcup_{\mu=1}^{\nu-1} J_{i_\mu}(A')$. Now, we first extend α' to $\{x_j \mid j \in (J \cup J') \setminus \{j_1, j_2, \ldots, j_\ell\}\}$ arbitrarily (say, by zeros) and then consecutively extend it to $x_{j_1}, \ldots, x_{j_\ell}$ so that the linear forms $\bigoplus_{j \in J_1(A)} x_j, \ldots, \bigoplus_{j \in J_\ell(A)} x_j$ are set to the right values.

In conclusion, $\widehat{E}$ satisfies assumptions of Claim 3.1 with $C := C_0 \vee x_j$. Since $C_0 \vee x_j$ has a proof from $\tau[A]$ of width $\leq r/4$ and size $\leq |\Pi|/2$, $D(\tau \vdash E) \leq \frac{r}{2}(\log_2 |\Pi| - 1)$. This conclusion holds for an arbitrary clause $\widehat{E}$ in the variables $\left\{ y_i \,\middle|\, i \in \mathrm{Ker}\left(\widehat{J}\right) \right\}$ consistent with E. Now we resolve all these clauses in the brute-force way along all the variables $\left\{ y_i \,\middle|\, i \in \mathrm{Ker}\left(\widehat{J}\right) \setminus \mathrm{Ker}(J) \right\}$. Since the depth of this original proof is at most $r/2$, we get a proof of E in depth $\frac{r}{2} \log_2 |\Pi|$.

This completes the analysis in case 2 of the inductive step. Claim 3.1 is proved.$\blacksquare$

Theorem 2.5 is now immediate. If τ is a contradiction and Π is a refutation of $\tau[A]$ with $w(\Pi) \leq r/4$, we simply apply Claim 3.1 with $C := 0,\ J := \emptyset$ and $E := 0$.

For Theorem 2.4, we are going to apply Theorem 2.5 to $\tau[A]$, where τ is the contradiction from Proposition 2.1 and A is the (random) matrix guaranteed by Lemma 2.2.

Proof of Theorem 2.4.
First of all we can assume that $k \geq 12$ since otherwise already the contradictions from Proposition 2.1 will do. Set $w := n^{1-\epsilon}/k$, $r := 4w$, $s := \lfloor k/4 \rfloor \geq 3$, $c := 2$ and choose the parameter m as the smallest value for which (9) is satisfied. Clearly, $m \geq (n/kw)^{\Omega(k)} \geq n^{\Omega(k)}$. If it turns out that $m \leq n^2$ then, as before, we simply take the 4-CNF contradiction τ_n provided by Proposition 2.1. Otherwise we take the formula τ_m provided by that proposition and compose it with an $m \times n$ $(r, s, 2)$-expander A guaranteed by Lemma 2.2.

Recall that $D(\tau_m \vdash 0) \geq \Omega(m/\log m)$. Hence Theorem 2.5 implies that every tree-like refutation Π of the k-CNF contradiction $\tau_m[A]$ with $w(\Pi) \leq w$ must have size at least

$$|\Pi| \geq \exp\left(\Omega\left(\frac{m}{r \log m}\right)\right) \geq \exp\left(\Omega\left(\frac{m}{n \log m}\right)\right) \overset{\text{since } m \geq n^2}{\geq} \exp(m^{\Omega(1)}) \geq \exp(n^{\Omega(k)}).$$

It only remains to remark that the width 6 refutation of τ_m stipulated by Proposition 2.1 can be converted into a width $O(k)$ refutation of $\tau_m[A]$ simply by applying the operator $E \mapsto E[A]$ to its lines.

4. Open problems

Our first (vaguely defined) open question is clear from the context: identify or prove more ultimate tradeoffs. At this point it is already clear that this peculiar phenomenon is rare, but is it *extremely* rare? It looks like one natural place to look for ultimate tradeoffs (given the abundance of traditional ones) is propositional space complexity.

Our result is a bit incomplete since the lower bound is double exponential only in the number of variables, not in the size of the contradiction. We remark that by contrast, in [Ber12] the size of the contradiction stays polynomial even when the minimum refutation width is unbounded. Is there any way to combine the small size of contradictions provided by Berkholz's method with a larger interval $[\mu_{\min}, \mu]$ in which the lower bound holds, as in our paper?

Attempting a rigorous formulation (there are many other ways to pinpoint this question), do there exist contradictions τ_n that possess constant width refutations, but such that any such refutation must necessarily have tree-like size $\exp(n^2)$?

Acknowledgement

I am greatly indebted to Jakob Nordström for extensive and very useful comments on an earlier version of this paper.

References

[AAT11] M. Alekhnovich, S. Arora, and I. Tourlakis. Toward strong nonapproximability results in the Lovász-Schrijver hierarchy. *Computational Complexity*, (4):615–648, 2011.

[ABRW04] M. Alekhnovich, E. Ben-Sasson, A. Razborov, and A. Wigderson. Pseudorandom generators in propositional proof complexity. *SIAM Journal on Computing*, 34(1):67–88, 2004.

[AD08] A. Atserias and V. Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, 2008.

[ALN14] A. Atserias, M. Lauria, and J. Nordström. Narrow proofs may be maximally long. In *Proceedings of the 29th IEEE Conference on Computational Complexity*, pages 286–297, 2014.

[BBI12] P. Beame, C. Beck, and R. Impagliazzo. Time-space tradeoffs in resolution: superpolynomial lower bounds for superlinear space. In *Proceedings of the 44th ACM Symposium on the Theory of Computing*, pages 213–232, 2012.

[Ben09] E. Ben-Sasson. Size-space tradeoffs for resolution. *SIAM Journal on Computing*, 38(6):2511–2525, 2009.

[Ber12] C. Berkholz. On the complexity of finding narrow proofs. Technical Report 1204.0775 [cs.LO], arXiv, 2012.

[BG99] M. Bonet and N. Galesi. A study of proof search algorithms for Resolution and Polynomial Calculus. In *Proceedings of the 40th IEEE FOCS*, pages 422–431, 1999.

[BG03] E. Ben-Sasson and N. Galesi. Space complexity of random formulae in resolution. *Random Structures and Algorithms*, 23(1):92–109, 2003.

[BHP10] P. Beame, T. Huynh, and T. Pitassi. Hardness amplification in proof complexity. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing,*, pages 87–96, 2010.

[BIW04] E. Ben-Sasson, R. Impagliazzo, and A. Wigderson. Near optimal separation of tree-like and general resolution. *Combinatorica*, 24(4):585–603, 2004.

[Bla37] A. Blake. *Canonical expressions in Boolean algebra*. PhD thesis, University of Chicago, 1937.

[BNT13] C. Beck, J. Nordström, and B. Tang. Some trade-off results for polynomial calculus: extended abstract. In *Proceedings of the 45th ACM Symposium on the Theory of Computing*, pages 813–822, 2013.

[Bor93] A Borodin. Time space tradeoffs (getting closer to the barrier?). In *Proceedings of the* 4th *International Symposium on Algorithms and Computation*, pages 209–220, 1993.

[Bra12] M. Braverman. Interactive information complexity. In *Proceedings of the 44th ACM Symposium on the Theory of Computing*, pages 505–524, 2012.

[BW01] E. Ben-Sasson and A. Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM*, 48(2):149–169, 2001.

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

[FLM+14] Y. Filmus, M. Lauria, M. Mikša, J. Nordström, and M. Vinyals. From small space to small width in resolution. Technical Report TR14-081, Electronic Colloquium on Computational Complexity, 2014.

[GKR14] A. Ganor, G. Kol, and R. Raz. Exponential separation of information and communication. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science*, pages 176–185, 2014.

[Nor09] J. Nordström. A simplified way of proving trade-off results for resolution. *Information Processing Letters*, 109:1030–1035, 2009.

[Nor13] J. Nordström. Pebble games, proof complexity and time-space trade-offs. *Logical Methods in Computer Science*, 9:1–63, 2013.

[Rob65] J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965.

[Tha14] N. Thapen. A trade-off between length and width in resolution. Technical Report TR14-137, Electronic Colloquium on Computational Complexity, 2014.

[Urq11] A. Urquhart. The depth of resolution proofs. *Studia Logica*, 99:349–364, 2011.

A. Appendix

Here we give self-contained proofs of Lemmas 2.2 and 2.3.

Lemma 2.2. *Let $n \to \infty$ and m, s, c be arbitrary integer parameters possibly depending on n such that $c \leq \frac{3}{4}s$ and*

$$r \leq o(n/s) \cdot m^{-\frac{2}{s-c}}.$$

Then for sufficiently large n there exist $m \times n$ (r, s, c)-boundary expanders.

Proof. This lemma and its proof is identical to [ABRW04, Theorem 5.1], except that we relax some restrictions on the parameters. We construct a random $m \times n$ matrix $\boldsymbol{A}$ by picking independently in each row s random entries with repetitions (the latter feature is not crucial, but it does make calculations neater). That is, we let $J_i(\boldsymbol{A}) \stackrel{\text{def}}{=} \{\boldsymbol{j_{i1}}, \ldots, \boldsymbol{j_{is}}\}$, where $\{\boldsymbol{j_{i\nu}}\}$ $(i \in [m], \nu \in [s])$ is a collection of ms independent random $[n]$-valued variables.

Recall that a matrix A is an (ordinary) (r, s, c)-*expander* if, again, $|J_i(A)| \leq s$ for all $i \in [m]$, and for every $I \subseteq [m]$ with $|I| \leq r$ we have $|\bigcup_{i \in I} J_i(A)| \geq c \cdot |I|$. Thus, the only difference from boundary expanders consists in replacing $\partial_A(I)$ with $\bigcup_{i \in I} J_i(A)$.

Claim A.1 *Every $\left(r, s, \frac{s+c}{2}\right)$-expander is an (r, s, c)-boundary expander.*

Proof of Claim A.1. Since every column $j \in \bigcup_{i \in I} J_i(A) \setminus \partial_A(I)$ belongs to at least two sets $J_i(A)$ $(i \in I)$, we have the bound

$$\left|\bigcup_{i \in I} J_i(A)\right| \leq |\partial_A(I)| + \frac{1}{2}\left(\sum_{i \in I} |J_i(A)| - \partial_A(I)\right) \leq \frac{1}{2}(s|I| + |\partial_A(I)|).$$

On the other hand, $|\bigcup_{i \in I} J_i(A)| \geq \frac{s+c}{2}|I|$ since A is an $\left(r, s, \frac{s+c}{2}\right)$-expander. The required inequality $|\partial_A(I)| \geq c \cdot |I|$ follows.■

Thus, it remains to prove that $\boldsymbol{A}$ is an (r, s, c')-expander a.s., where $c' \stackrel{\text{def}}{=} \frac{c+s}{2}$. Let p_ℓ be the probability that any given ℓ rows of the matrix $\boldsymbol{A}$ violate the expansion property. Then

$$\mathbf{P}[\boldsymbol{A} \text{ is not a } (r, s, c')\text{-expander}] \le \sum_{\ell=1}^{r} p_\ell \cdot m^\ell.$$

On the other hand,

$$\begin{aligned} p_\ell = \mathbf{P}[|\{\boldsymbol{j_{i\nu}} \mid i \in I,\ \nu \in [s]\}| \le c'\ell] &\le \binom{n}{c'\ell} \cdot \left(\frac{c'\ell}{n}\right)^{s\ell} \\ \le O(1)^{c'\ell} \cdot \left(\frac{c'\ell}{n}\right)^{(s-c')\ell} &\le \{O((sl)/n)\}^{(s-c')\ell}, \end{aligned}$$

where for the last inequality we used that $c' \le \frac{7}{8}c \le \frac{7}{8}s$ and hence $c' \le O(s-c)$. Thus,

$$\mathbf{P}[\boldsymbol{A} \text{ is not a } (r, s, c')\text{-expander}] \le \sum_{\ell=1}^{r} \{O((sl)/n)\}^{(s-c')\ell} m^\ell \le \sum_{\ell=1}^{r} \left(\{O((sr)/n)\}^{(s-c')\ell} m\right)^\ell,$$

and since $m(sr/n)^{s-c'} = m(sr/n)^{(s-c)/2} \le o(1)$ by our assumption, Lemma 2.2 follows.$\blacksquare$

Lemma 2.3. *Let A be an $m \times n$ $(r, 2)$-boundary expander. Then for every $J \subseteq [n]$ with $|J| \le r/4$ there exists $\widehat{J} \supseteq J$ sich that $\left|\mathrm{Ker}\left(\widehat{J}\right)\right| \le 2|J|$ and $A \setminus \widehat{J}$ is an $(r/2, 3/2)$-boundary expander.*

Proof. We define a strictly increasing sequence of sets of columns $J_0 \subset J_1 \subset \ldots \subset J_t \subset \ldots$ as follows. Let $J_0 \stackrel{\text{def}}{=} J$. For $t > 0$, we first let S_t be an arbitrary set of rows violating the $(r/2, 3/2)$-boundary expansion condition in $A \setminus J_{t-1}$ if such a set exists; otherwise, the construction terminates. Then we let

$$J_t \stackrel{\text{def}}{=} J_{t-1} \cup \bigcup_{i \in S_t} J_i(A).$$

Note that since the chain $J_0 \subset J_1 \subset \ldots \subset J_t \ldots$ is strictly increasing, the process does terminate at some point; let J_T be the final set in this chain. We claim that $\widehat{J} := J_T$ has the required properties, and the only thing that has to be checked is that $|\mathrm{Ker}(J_T)| \le 2|J|$. For that we prove by induction on $t = 0, \ldots, T$ that $|\mathrm{Ker}(J_t)| \le 2|J|$.

Base case $|\mathrm{Ker}(J)| \leq 2|J|$ immediately follows from the fact that A is an $(r, 2)$-boundary expander and $|J| \leq r/4$.

Inductive step. Assume that $|\mathrm{Ker}(J_{t-1})| \leq 2|J|$ for some $1 \leq t \leq T$, and let us prove that $|\mathrm{Ker}(J_t)| \leq 2|J|$.

Since $|S_t| \leq r/2$, $|\mathrm{Ker}(J_{t-1})| \leq 2|J| \leq r/2$ and $\mathrm{Ker}(J_{t-1}) \cup S_t \subseteq \mathrm{Ker}(J_t)$, we can choose a set of rows I such that $\mathrm{Ker}(J_{t-1}) \cup S_t \subseteq I \subseteq \mathrm{Ker}(J_t)$ and

$$|I| = \min(r, |\mathrm{Ker}(J_t)|). \tag{11}$$

Applying to I the expansion condition, we get

$$|\partial_A(I)| \geq 2|I|.$$

On the other hand, $I \subseteq \mathrm{Ker}(J_t)$ implies that

$$\partial_A(I) \subseteq J \cup \bigcup_{s=1}^{t} \partial_{A \setminus J_{s-1}}(S_s).$$

Since S_s's violate the $(r/2, 3/2)$-boundary expansion conditions in respective matrices, we conclude that

$$|\partial_A(I)| \leq |J| + \frac{3}{2} \sum_{s=1}^{t} |S_s| \leq |J| + \frac{3}{2}|I|,$$

where for the latter inequality we used the fact $I \supseteq S_1 \dot\cup S_2 \dot\cup \ldots \dot\cup S_t$. Comparing these two inequalities, we find that $|I| \leq 2|J| \leq r/2$. Now (11) implies that in fact $|I| = |\mathrm{Ker}(J_t)|$ that completes the inductive step in the proof of Lemma 2.3.■

ECCC ISSN 1433-8092

http://eccc.hpi-web.de