comput. complex. **27** (2018), 511 – 559

1016-3328/18/030511-49
published online October 17, 2017
DOI 10.1007/s00037-017-0163-1

computational complexity

ON SPACE AND DEPTH IN RESOLUTION

Alexander Razborov

Abstract. We show that the total space in resolution, as well as in any other reasonable proof system, is equal (up to a polynomial and $(\log n)^{O(1)}$ factors) to the minimum refutation depth. In particular, all these variants of total space are equivalent in this sense. The same conclusion holds for variable space as long as we penalize for excessively (that is, super-exponential) long proofs, which makes the question about equivalence of variable space and depth about the same as the question of (non)-existence of "supercritical" tradeoffs between the variable space and the proof length. We provide a partial negative answer to this question: for all $s(n) \leq n^{1/2}$ there exist CNF contradictions τ_n that possess refutations with variable space $s(n)$ but such that every refutation of τ_n with variable space $o(s^2)$ must have double exponential length $2^{2^{\Omega(s)}}$. We also include a much weaker tradeoff result between variable space and depth in the opposite range $s(n) \ll \log n$ and show that no supercritical tradeoff is possible in this range.

Keywords. Resolution proofs, supercritical tradeoff, variable and total space.

Subject classification. Primary 03F20; Secondary 68Q17.

1. Introduction

The area of propositional proof complexity has seen a rapid development since its inception in the seminal paper Cook & Reckhow (1979). 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 recent years is the interplay

Birkhäuser

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 employing conflict-driven clause learning) are inherently based on the *resolution* proof system dating back to the papers Blake (1937) and Robinson (1965). This somewhat explains the fact that resolution is by far the most studied system in proof complexity, even if recent developments (see, e.g., the survey Barak & Steurer (2014)) seem to be bringing the system of sum-of-squares as a serious rival.[1]

Most of this study concentrated on natural complexity measures of resolution proofs like size, width, depth or space and on their mutual relations; to facilitate further discussion, let us fix some notation (the reader not familiar with some or all of these is referred to Section 2 in which we give all necessary definitions). Namely, we let $S(\tau_n \vdash 0), S_T(\tau_n \vdash 0), w(\tau_n \vdash 0), D(\tau_n \vdash 0), \mathsf{CSpace}(\tau_n \vdash 0), \mathsf{TSpace}(\tau_n \vdash 0)$ and $\mathsf{VSpace}(\tau_n \vdash 0)$ stand for the minimum possible size [tree-like size, width, depth, clause space, total space[2] and variable space, respectively]. $w(\tau_n)$ is the width of the contradiction τ_n itself.

Let us review some prominent relations between these measures. The inequalities $w(\tau_n \vdash 0) \leq D(\tau_n \vdash 0)$ and $\log S_T(\tau_n \vdash 0) \leq D(\tau_n \vdash 0)$ are trivial. Ben-Sasson & Wigderson (2001) conjoined them by proving that

$$w(\tau_n \vdash 0) \leq \log S_T(\tau_n \vdash 0) + w(\tau_n). \tag{1.1}$$

Even more importantly, in the same paper they established the celebrated width-size relation

$$w(\tau_n \vdash 0) \leq O(n \cdot \log S(\tau_n \vdash 0))^{1/2} + w(\tau_n) \tag{1.2}$$

[1] It should be remarked, however, that one of the most prominent SOS lower bound technique dating back to the paper Grigoriev (2001) *is* based on resolution width.

[2] A word of warning about terminology: it is this measure that had been called "variable space" in Alekhnovich *et al.* (2002), and this usage of the term persisted in the literature for a while, see, e.g., Ben-Sasson (2009). But then several good arguments were brought forward as to why it is more natural to reserve the term "variable space" for its connotative meaning, and we follow this revised terminology.

that has steadily grown into a standard method of proving lower bounds on the size of DAG resolution proofs.

In the space world, the obvious relations are $\mathsf{CSpace}(\tau_n \vdash 0) \leq \mathsf{TSpace}(\tau_n \vdash 0)$ and $\mathsf{VSpace}(\tau_n \vdash 0) \leq \mathsf{TSpace}(\tau_n \vdash 0)$. Can $\mathsf{CSpace}(\tau_n \vdash 0)$ and $\mathsf{VSpace}(\tau_n \vdash 0)$ be meaningfully related to each other?

In one direction this was ruled out by (Ben-Sasson 2009, Theorem 3.9): there are 3-CNF contradictions τ_n with $\mathsf{CSpace}(\tau_n \vdash 0) \leq O(1)$ and[3] $\mathsf{VSpace}(\tau_n \vdash 0) \geq \Omega(n/\log n)$.

Whether CSpace can be meaningfully bounded by VSpace is unknown. As will become clear soon, this question is extremely tightly connected to the content of our paper.

Let us mention several prominent and rather non-trivial results connecting "sequential" measures (size, width, depth) and "configurational", space-oriented ones. Atserias & Dalmau (2008) proved that

$$w(\tau_n \vdash 0) \leq \mathsf{CSpace}(\tau_n \vdash 0) + w(\tau_n); \tag{1.3}$$

a simplified version of their proof was presented by Filmus *et al.* (2015) and independently by Razborov (unpublished).

As we already observed, variable space cannot be bounded in terms of clause space, but Urquhart (2011) proved that it can be bounded by depth:

$$\mathsf{VSpace}(\tau_n \vdash 0) \leq D(\tau_n \vdash 0). \tag{1.4}$$

In a recent paper Bonacina (2016), the following connection between width and *total* space was established:

$$w(\tau_n \vdash 0) \leq O(\mathsf{TSpace}(\tau_n \vdash 0))^{1/2} + w(\tau_n) \tag{1.5}$$

that, similarly to (1.2), immediately opens up the possibility of proving super-linear lower bounds on the total space in a systematic way.

[3] Ironically (cf. Footnote 2), although this result was stated in Ben-Sasson (2009) for variable space, it was actually proved there only for what we call here TSpace. However, the extension to VSpace is more or less straightforward, see, e.g., Beck *et al.* (2013).

Finally, it should be mentioned that besides simulations there have been proven quite a great deal of separation and tradeoff results between these measures. They are way too numerous to be meaningfully accounted for here, we refer the interested readers, e.g., to the survey Nordström (2013).

Our contributions. We continue this line of research and prove both simulations and tradeoff results. In the former direction, perhaps the most catchy statement we can make is that $\mathsf{TSpace}(\tau_n \vdash 0)$ and $D(\tau_n \vdash 0)$ are equivalent, up to a polynomial and $\log n$ factors (see Figure 1.1 below for more refined statements). This is arguably the first example when two proof-complexity measures that are quite different in nature and have very different history turn out not only to be tightly related to each other, but actually practically equivalent.

Now, in order to discuss these simulations and their ramifications properly, we need to make up a few definitions.

For a configurational proof[4] π, let $\mathsf{VSpace}^*(\pi) \stackrel{\text{def}}{=} \mathsf{VSpace}(\pi) \cdot \log_2 |\pi|$; a similar definition can be made for the total space and for the clause space although we do not need the latter in our paper. Thus, we penalize refutations in a configurational form for being *excessively* long; let us note that a similar logarithmic normalization naturally pops up in many tradeoff results, see, e.g., Ben-Sasson (2009). Then what we "actually" do is to show that $\mathsf{VSpace}^*(\tau_n \vdash 0)$ is polynomially related to depth; in particular, any small variable space proof can be unfolded into a shallow sequential proof *unless* it is prohibitively long. Given this simulation, the equivalence for the total space is a simple artifact of the observation that proofs with small total space cannot be too long just because there are not that many possible different configurations. More specifically, we have the following picture in which, for the sake of better readability, we have removed $\tau_n \vdash 0$ elsewhere, replaced $f \leq O(g)$ with $f \preceq g$ and taken the liberty to blend new results (that are essentially observations) with previously known ones, like (1.4), and trivial inequalities like $D^2 \leq D^3$.

An immediate corollary is that TSpace, D, TSpace^* and VSpace^*

[4] For definitions see Section 2 below.

$$
\begin{array}{ccccccc}
\mathsf{VSpace} & & \preceq & & \mathsf{TSpace} & \preceq & D^2 \\
\rotatebox{90}{$\geq$} & & & & \rotatebox{90}{$\succeq$} & & \rotatebox{90}{$\geq$} \\
D & \preceq & \mathsf{VSpace}^* & \preceq & \mathsf{TSpace}^* & \preceq & D^3 \\
 & & \rotatebox{90}{$\succeq$} & & \rotatebox{90}{$\succeq$} & & \\
\multicolumn{3}{c}{\mathsf{VSpace} \cdot (\mathsf{VSpace} \cdot \log n + 2^{\mathsf{VSpace}})} & & \mathsf{TSpace}^2 \cdot \log n & &
\end{array}
$$

Figure 1.1: Simulations.

are all equivalent up to a polynomial and $\log n$ factors, and the same applies for semantic versions of TSpace and TSpace^*.

The only difference between TSpace and VSpace is that in the first case we have a decent (that is, singly exponential) bound on the overall number of configurations of small total space. Due to the standard counting argument, this remains true for an *arbitrary* reasonable circuit class, and hence our equivalence uniformly generalizes to the total space based on any one of them: polynomial calculus with resolution, cutting planes etc. (Theorem 3.2). All these measures are essentially depth in disguise, and hence $n^{\Omega(1)}$ depth lower bounds *automatically* imply $n^{\Omega(1)}$ lower bounds on the total space in all those models. For instance, for the total space we have

$$\widetilde{\Omega}(D^{1/2}) \leq \mathsf{TSpace} \leq O(D^2), \tag{1.6}$$

(the question how tight are these bounds will be addressed in the concluding section 7).

In the rest of the paper, we study the relation of variable space itself to these equivalent measures; this question was (apparently) first asked by Urquhart Urquhart (2011). As follows from Figure 1.1, this is equivalent to the following question: can the term 2^{VSpace} in the upper bound on VSpace^* be really dominating or, in other words, can it be the case that the length of a configurational proof must necessarily be *super*-exponential, as long as its variable space is relatively small? Note that this in particular would imply that such a proof must mostly consist of totally non-constructive configurations so this situation may look a bit counterintuitive on the first sight. However, precisely this kind of a behavior dubbed "supercritical" tradeoffs was recently exhibited in Razborov (2016), and several other examples have been found in Berkholz & Nordström (2016a,b) and Razborov (2017).

Our most difficult result (Theorem 3.3) gives a moderate supercritical tradeoff between variable space and proof length: for any $s = s(n) \leq n^{1/2}$, there are $O(1)$-CNF contradictions τ_n with $\mathsf{VSpace}(\tau_n \vdash 0) \leq s$ but such that every refutation π with sub-quadratic variable space $o(s^2)$ must have length $2^{2^{\Omega(s)}}$. Improving the space gap from sub-quadratic to super-polynomial would establish a strong separation between the variable space and the depth, but that would probably require new techniques or at least quite a significant enhancement of ours. As a matter of fact, I am not ready even to *conjecture* that a super-polynomial gap here exists, and perhaps VSpace after all is equivalent to all other measures in Figure 1.1.

The proof of Theorem 3.3 is highly modular and consists of three independent reductions; we review its overall structure at the beginning of Section 6 where the statement is proven. Among previously known ingredients we can mention r-surjective functions Alekhnovich & Razborov (2008), "hardness compression" Razborov (2016) and an extensive usage of the multi-valued logic in space-oriented models Alekhnovich *et al.* (2002). One new idea that we would like to highlight is a "direct product result" Lemma 6.9; results of this sort do not seem to be too frequent in the proof complexity. We use it to amplify our length lower bound for proofs of variable space 1 (that is, consisting of multi-valued literals) to the same lower bound for proofs of larger variable space. This is precisely this step that exponentially blows up the number of multi-valued variables and prevents us from extending this supercritical tradeoff into a super-quadratic space range.

Finally, we look into the opposite range when $\mathsf{VSpace}(\tau_n \vdash 0)$ is very small (say, a constant) and hence the term 2^{VSpace} in Figure 1.1 becomes negligible. In this regime, the syntactic measures $\mathsf{CSpace}, \mathsf{TSpace}$ become constant and, by (1.3), the same applies to width. Razborov (2016) proved a supercritical tradeoff between width and depth, and Berkholz & Nordström (2016b) studied this question for width vs. space, so it seems very natural to ask what kind of tradeoffs might exist between space and depth. We prove both positive and negative results in this direction. First, we observe (Theorem 3.4) that the proof of the relation $D \leq \mathsf{VSpace}^*$ in Figure 1.1 can be generalized to showing that every (semantical) refutation of constant variable space gives rise, for an arbitrary parameter h, to a configurational refutation of variable space $O(h)$ and depth $h^2 \cdot n^{O(1/h)}$; in particular, both space and depth can be made poly-logarithmic, or depth can be brought down to $n^{1/10}$ while space still remains constant. This rules out supercritical tradeoffs in this context, at least as strong as those in Berkholz & Nordström (2016b); Razborov (2016). But we also show that this simulation is essentially the best possible: for the *Induction Principle* $\tau_n = \{x_0,\ x_0 \to x_1, \dots, x_{n-1} \to x_n,\ \bar{x}_n\}$ we show that every refutation π with variable space s must have depth $n^{\Omega(1/s)}$ (Theorem 3.5).

The structure of the paper corresponds to the above overview. In Section 2 we review all the necessary definitions, and in Section 3 we state our main results. The next three sections are devoted to proofs: simulation results in Section 4, small space results in Section 5 and the supercritical tradeoff for large space[5] in Section 6. We conclude with a few remarks and open problems in Section 7.

2. Notation and preliminaries

We let $[n] \stackrel{\text{def}}{=} \{1, 2, \dots, n\}$. $|x|$ is the length of a word x, and xy is the concatenation of two words. u is a *prefix* of v, denoted by $u \leq v$ if $v = uw$ for another (possibly, empty) word w, and Λ is the empty word.

For a Boolean function f, $Vars(f)$ is the set of variables f essentially depends on. $f \models g$ stands for *semantical implication* and means that every assignment α satisfying f satisfies g as well. If τ and τ' are syntactic expressions like CNFs, the semantical implication $\tau \models \tau'$ is understood in terms of the Boolean functions;

[5] We defer our by far most difficult proof to the end.

these expressions represent.

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* is a disjunction (possibly, empty) of literals in which no variable appears along with its negation. A *generalized clause* is either a clause or 1; the set of all generalized clauses makes a lattice in which $\vee$ is the join operator. If C and D are clauses then $C \leq D$ in this lattice if and only if $C \models D$ if and only if every literal appearing in C also appears in D. We will also sometimes say that C is a sub-clause of D in this case. The empty clause will be denoted by 0, and the set of variables occurring in a clause C, either positively or negatively, will be denoted by $Vars(C)$, let also $Vars(1) \stackrel{\text{def}}{=} \emptyset$. This is consistent with the general semantic definition. The *width* of a clause C is defined as $w(C) \stackrel{\text{def}}{=} |Vars(C)|$.

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*.

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

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

Two major topologies used for representing resolution proofs are *sequential* (Hilbert-style) and *configurational* (or *space-oriented*). In order to distinguish between them, we use upper-case letters Π for the former and lower-case π for the latter.

A (sequential) resolution proof Π is a DAG with a unique target node in which all nodes are labeled by clauses, every non-source node v has fan-in 2, and the clause assigned to v can be inferred from clauses sitting at its predecessors via a single application of the resolution rule. A *resolution proof of a clause C from a CNF τ* is a resolution proof Π in which all source nodes are labeled by clauses from τ, and the target node is labeled by a sub-clause[6] of

[6] This is a technicality that is necessary since we did not explicitly include the weakening rule.

C. A *refutation* of a contradiction τ is a proof of 0 from it. The *size* $S(\Pi)$ of a sequential proof is the number of nodes, its *depth* $D(\Pi)$ is the length of the longest path in the underlying DAG, and its *width* $w(\Pi)$ is the maximal possible width $w(C)$ of a clause C appearing in it. For a contradiction τ, we let $S(\tau \vdash 0)$, $D(\tau \vdash 0)$ and $w(\tau \vdash 0)$ denote the minimal possible value of $S(\Pi)$, $D(\Pi)$ and $w(\Pi)$, respectively, taken over all sequential refutations Π of τ.

The *configurational* (or *space-oriented*) form of propositional proofs was introduced in Alekhnovich *et al.* (2002); Esteban & Torán (2001). A *configuration* $\mathbb{C}$ is a set of generalized clauses that can be viewed as a CNF. A *configurational proof* π from a CNF formula τ is a sequence of configurations $(\mathbb{C}_0, \ldots, \mathbb{C}_T)$ in which $\mathbb{C}_0 = \emptyset$ and every $\mathbb{C}_t$ $(t \in [T])$ is obtained from $\mathbb{C}_{t-1}$ by one of the following rules:

Axiom Download. $\mathbb{C}_t = \mathbb{C}_{t-1} \cup \{A\}$, where $A \in \tau$;

Inference. $\mathbb{C}_t = \mathbb{C}_{t-1} \cup \{C\}$ for some $C \notin \mathbb{C}_{t-1}$ inferrable by a single application of the resolution rule from the clauses in $\mathbb{C}_{t-1}$.

Erasure. $\mathbb{C}_t \subseteq \mathbb{C}_{t-1}$.

π is a (configurational) refutation of τ if $0 \in \mathbb{C}_T$. T is the *length* of π, denoted by $|\pi|$.

The *clause space* of a configuration $\mathbb{C}$ is $|\mathbb{C}|$, its *total space* $\mathsf{TSpace}(\mathbb{C})$ is $\sum_{C \in \mathbb{C}} w(C)$, and its *variable space* $\mathsf{VSpace}(\mathbb{C})$ is $\left|\bigcup_{C \in \mathbb{C}} Vars(C)\right|$. The *clause* [*total*] space $\mathsf{CSpace}(\pi)$ [$\mathsf{TSpace}(\pi)$, respectively] of a configurational proof π is the maximal clause [total, respectively] space of all its configurations, and if τ is a contradiction, then $\mathsf{CSpace}(\tau \vdash 0)$ [$\mathsf{TSpace}(\tau \vdash 0)$] is the minimum value of $\mathsf{CSpace}(\pi)$ [$\mathsf{TSpace}(\pi)$, respectively], where the minimum is taken over all configurational refutations π of τ.

Variable space $\mathsf{VSpace}(\tau \vdash 0)$ can be of course defined analogously, but since this measure is inherently semantical, we prefer to stress this fact by giving a separate, and more robust, definition below.

Definition 2.1. *Let τ be an arbitrary set of Boolean constraints. For a set V of variables, we let*

$$\tau[V] \stackrel{\text{def}}{=} \bigwedge \{ C \mid C \in \tau \wedge Vars(C) \subseteq V \}.$$

A semantical proof *π from τ is a sequence of Boolean functions $(f_0, \ldots, f_1, \ldots, f_T)$ such that $f_0 \equiv 1$ and for every $t \in T$,*

$$f_{t-1} \wedge \tau[Vars(f_{t-1}) \cup Vars(f_t)] \models f_t. \tag{2.2}$$

T is again the length *of π, denoted by $|\pi|$, and π is a* semantical refutation *if $f_T \equiv 0$. $\mathsf{VSpace}(\pi) \stackrel{\text{def}}{=} \max_{0 \le t \le T} |Vars(f_t)|$ and $\mathsf{VSpace}(\tau \vdash 0)$ is the minimum value of $\mathsf{VSpace}(\pi)$ taken over all semantical refutations π of τ.*

In this definition we have combined all three rules (Axiom Download, Inference and Erasure) into one. Every configurational proof turns into a semantical proof of (at most) the same variable space if we replace all configurations in it by the Boolean functions they represent. Hence $\mathsf{VSpace}(\tau \vdash 0)$ never exceeds its syntactical variant, and in the other direction (when τ is actually a CNF), they may differ by at most a factor of 2 simply by expanding all semantical refutations (2.2) into brute-force resolution derivations never leaving the set of variables $Vars(f_{t-1}) \cup Vars(f_t)$.

This purely semantical model also provides a handy uniform way to talk about semantical analogues of more sophisticated space complexity measures. Namely, let $\mathcal{C}$ be a circuit class equipped with a complexity measure $\mu(C)$ $(C \in \mathcal{C})$. Then μ gives rise to the complexity measure on Boolean functions in a standard way: $\mu_{\mathcal{C}}(f)$ is the minimum value of $\mu_{\mathcal{C}}$ taken over all circuits $C \in \mathcal{C}$ computing f. For a semantical refutation π, let us define $\mu_{\mathcal{C}}\text{-}\mathsf{Space}(\pi) \stackrel{\text{def}}{=} \max_{0 \le t \le T} \mu_{\mathcal{C}}(f_t)$ and then $\mu_{\mathcal{C}}\text{-}\mathsf{Space}$ as usual. This definition may seem overly broad at the first glance, but we will see in Theorem 3.2 that under very mild conditions on the circuit class $\mathcal{C}$, this measure will also turn out to be equivalent to depth.

Examples. Semantical analogues of clause and total space studied in the literature before correspond to the case when $\mathcal{C}$ consists of all CNFs, and the measures $\mu_{\mathcal{C}}$ are the number of clauses or overall size, respectively. Semantical analogues of, say, cutting planes space or PCR space are also straightforward in this language.

Finally, we need several mixed, "amortized" measures that penalize configurational proofs for being unreasonably long. We let

$$\begin{aligned} \mathsf{TSpace}^*(\pi) &\stackrel{\text{def}}{=} \mathsf{TSpace}(\pi)\cdot \log_2|\pi|, \\ \mathsf{VSpace}^*(\pi) &\stackrel{\text{def}}{=} \mathsf{VSpace}(\pi)\cdot \log_2|\pi|, \\ \mu_{\mathcal{C}}\text{-}\mathsf{Space}^*(\pi) &\stackrel{\text{def}}{=} \mu_{\mathcal{C}}\text{-}\mathsf{Space}(\pi)\cdot \log_2|\pi|, \end{aligned}$$

and then we define $\mathsf{TSpace}^*(\tau \vdash 0)$ $\mathsf{VSpace}^*(\tau \vdash 0)$ and $\mu_{\mathcal{C}}\text{-}\mathsf{Space}^*$ $(\tau \vdash 0)$ as usual ($\mathsf{CSpace}^*(\tau \vdash 0)$ can be also defined likewise, but we do not need it in this paper).

DEFINITION 2.3. *For a configurational proof* $\pi = (\mathbb{C}_0, \mathbb{C}_1, \ldots, \mathbb{C}_T)$, *define integer valued* depth functions D_t *on* $\mathbb{C}_t$ *by induction on* t. *Since* $\mathbb{C}_0$ *is empty, there is nothing to define. Let* $t > 0$, *assume that* $C \in \mathbb{C}_t$ *and that* D_{t-1} *is already defined. If* $C \in \mathbb{C}_{t-1}$, *we simply let* $D_t(C) \stackrel{\text{def}}{=} D_{t-1}(C)$. *If* $A \in \tau$ *in the* AXIOM DOWNLOAD RULE *then* $D_t(A) \stackrel{\text{def}}{=} 0$. *If* C *is obtained from* $C', C'' \in \mathbb{C}_{t-1}$ *via the resolution rule, we let*

$$D_t(C) \stackrel{\text{def}}{=} \max(D_{t-1}(C'), D_{t-1}(C'')) + 1.$$

Finally, the depth $D(\pi)$ *of a configurational refutation* π *is defined as* $D_T(0)$.

Let us remark that Boolean restrictions naturally act on configurational refutations, and that under this action neither space (any flavor) nor depth may increase.

3. Main results

As all our results were discussed at length in the introduction, here they are listed more or less matter-of-factly.

In order to improve readability, in our first theorem 3.1 we omit the argument $\tau_n \vdash 0$ throughout (τ_n is an arbitrary contradiction in n variables), and also we write $f \preceq g$ for $f \leq O(g)$. Also, we include here both previously known results and trivial simulations; for explicit tags see the proof given in the next section. Arguably, the only non-trivial new part in the whole table is the simulation $D \preceq \mathsf{VSpace}^*$.

Theorem 3.1. *For the proof-complexity measures* $D, \mathsf{TSpace}, \mathsf{VSpace}, \mathsf{TSpace}^*, \mathsf{VSpace}^*$ *introduced in Section 2 we have the following simulations:*

$$\begin{array}{ccccccc}
\mathsf{VSpace} & & \preceq & & \mathsf{TSpace} & \preceq & D^2 \\
\rotatebox{90}{$\succeq$} & & & & \rotatebox{90}{$\succcurlyeq$} & & \rotatebox{90}{$\succeq$} \\
D & \preceq & \mathsf{VSpace}^* & \preceq & \mathsf{TSpace}^* & \preceq & D^3 \\
 & & \rotatebox{90}{$\succcurlyeq$} & & \rotatebox{90}{$\succcurlyeq$} & & \\
\multicolumn{3}{c}{\mathsf{VSpace} \cdot (\mathsf{VSpace} \cdot \log n + 2^{\mathsf{VSpace}})} & & \mathsf{TSpace}^2 \cdot \log n. & &
\end{array}$$

Theorem 3.2. *Let $\mathcal{C}$ be any circuit class that includes CNFs, and let $\mu_{\mathcal{C}}$ be any complexity measure on $\mathcal{C}$ that is intermediate between the number of input variables and the circuit size of $C \in \mathcal{C}$. Then $\mu_{\mathcal{C}}$-$\mathsf{Space}(\tau_n \vdash 0)$ is equivalent, up to a polynomial and $\log n$ factors to $D(\tau_n \vdash 0)$ (and hence to all other measures in Theorem 3.1 except, possibly,* VSpace*).*

Theorem 3.3. *Let $s = s(n) \leq n^{1/2}$ be an arbitrary parameter. Then there exists a CNF τ_n with $\mathsf{VSpace}(\tau_n \vdash 0) \leq s$ but such that for any semantical refutation π of τ_n with $\mathsf{VSpace}(\pi) \leq o(s^2)$ we have $|\pi| \geq \exp(\exp(\Omega(s)))$.*

The next result is a variation on the simulation $D \preceq \mathsf{VSpace}^*$ in Theorem 3.1.

Theorem 3.4. *Assume that a contradiction τ possesses a semantical refutation π with $\mathsf{VSpace}(\pi) = s$ and $|\pi| = S$, and let $h \geq 1$ be an arbitrary parameter. Then τ also has a configurational refutation π' with $\mathsf{VSpace}(\pi') \leq O(sh)$ and $D(\pi') \leq O\left(sh^2 \cdot S^{1/h}\right)$.*

THEOREM 3.5. *Let* $\tau_n = \{x_0, \bar{x}_0 \vee x_1, \bar{x}_1 \vee x_2, \ldots, \bar{x}_{n-1} \vee x_n, \bar{x}_n\}$. *Then for every configurational refutation* π *from* τ_n*, we have the bound*

$$D(\pi) \geq \Omega\left(n^{1/\mathsf{VSpace}(\pi)}\right).$$

4. Proofs of simulations

In this short section, we prove Theorems 3.1 and 3.2.

PROOF OF 3.1. The square

$$\begin{array}{ccc} \mathsf{VSpace} & \leq & \mathsf{TSpace} \\ & & \wedge\!| \\ \mathsf{VSpace}^* & \leq & \mathsf{TSpace}^* \end{array}$$

is obvious.

Variable space is upper bounded by depth.
$\mathsf{VSpace}(\tau \vdash 0) \leq D(\tau \vdash 0)$ is (Urquhart 2011, Theorem 6.1(1)).

Total space is upper bounded by depth squared.
This is a minor variation on (Esteban & Torán 2001, Theorem 2.1). Indeed, let Π be a refutation of a contradiction τ with $D(\Pi) = d$, then w.l.o.g. we can assume that Π is in a tree-like form. Also, $w(\Pi) \leq d$ since every variable in the clause at a node v must be resolved on the path from v to the target (root) node. We now consider the standard pebbling of the underlying tree with $(d+1)$ pebbles and the resulting configuration refutation $\pi = (\mathbb{C}_0, \mathbb{C}_1, \ldots, \mathbb{C}_T)$, as in Esteban & Torán (2001). As a reminder, we can assume w.l.o.g. that Π is a *complete* binary tree of height d, identify its nodes with binary words α of length $\leq d$ in such a way that $\alpha 0$ and $\alpha 1$ are the two children of α, and let C_α be the corresponding clauses so that $\dfrac{C_{\alpha 0} \quad C_{\alpha 1}}{C_\alpha}$ is an instance of the resolution rule. $\mathbb{C}_t$ then consists of all those C_α for which $\alpha 1^{d-|\alpha|} \leq t$, where the left-hand side is interpreted as an integer in the binary notation and that are *maximal* with this property (i.e., either $\alpha = \Lambda$ or $\alpha = \beta 0$ with $\beta 1^{d+1-|\alpha|} > t$). It

can be readily checked that $\mathbb{C}_{t+1}$ is obtained from $\mathbb{C}_t$ by a single AXIOM DOWNLOAD rule, followed by at most d INFLUENCE-ERASURE pairs, and that $|\mathbb{C}_t| \leq d+1$. Since also every clause in π has width $\leq d$ (due to the above remark) and $T \leq 2^{d+1}$, both claims $\mathsf{TSpace}(\tau \vdash 0) \leq D(\tau \vdash 0)(D(\tau \vdash 0)+1)$ and $\mathsf{TSpace}^*(\tau \vdash 0) \leq D(\tau \vdash 0)(D(\tau \vdash 0)+1)^2$ now follow.

"Amortized" space is upper bounded by ordinary space. We need to prove that

$$\mathsf{TSpace}^*(\tau_n \vdash 0) \leq 2\log_2(2n+1)\mathsf{TSpace}(\tau_n \vdash 0)^2$$

and

$$\begin{aligned}&\mathsf{VSpace}^*(\tau_n \vdash 0)\\&\quad\leq \mathsf{VSpace}(\tau_n \vdash 0)\left(\mathsf{VSpace}(\tau_n \vdash 0)\log_2 n + 2^{\mathsf{VSpace}(\tau_n \vdash 0)}\right).\end{aligned}$$

Both bounds follow from the observation that a configurational refutation (be it syntactic or semantic) can w.l.o.g. be assumed not to contain repeated configurations. Now, we estimate the overall number of configurations $\mathbb{C} = \{C_1, \ldots, C_k\}$ with total space $\leq s$ by encoding them as a string $C_1\#C_2\#\ldots\#C_k\#\ldots$ of length $2s$ in which the clauses C_i are written down simply as sequences of literals. We conclude that the overall number of different configurations $\mathbb{C}$ of total space $\leq s$ is bounded by $(2n+1)^{2s}$, which gives us the first statement. Likewise, the overall number of Boolean functions f with $|Vars(f)| \leq s$ is upper bounded by $\binom{n}{s}2^{2^s} \leq n^s 2^{2^s}$, and this gives us the second statement.

Depth is upper bounded by amortized variable space. This is by standard binary search. Let $\pi = (f_0, f_1, \ldots, f_T)$ be a semantical refutation from τ minimizing $\mathsf{VSpace}^*(\pi)$, and let $s \stackrel{\text{def}}{=} \mathsf{VSpace}(\pi)$. We prove by induction on d that for every $0 \leq a < b \leq T$ with $b - a \leq 2^d$ and for any clause C in the straightforward CNF expansion of the implication $f_a \to f_b$ (that is to say, for every clause C with $Vars(C) = Vars(f_a) \cup Vars(f_b)$ and $(f_a \to f_b) \models C$) we have $D(\tau \vdash C) \leq 2s(d+1)$.

Induction base $d = 0$, $b = a + 1$.

We have $f_a \wedge \tau[Vars(f_a) \cup Vars(f_{a+1})] \models f_{a+1}$, hence

$$\tau[Vars(f_a) \cup Vars(f_{a+1})] \models (f_a \to f_{a+1}) \models C.$$

Since $|Vars(f_a) \cup Vars(f_{a+1})| \leq 2s$, we can realize the latter semantical refutation by a resolution refutation of depth $\leq 2s$.

Inductive step: $d \geq 1$, $2 \leq b - a \leq 2^d$.

Pick c with $a < c < b$ such that $c - a,\ b - c \leq 2^{d-1}$. Then C has an obvious resolution proof of depth $|Vars(f_c) \setminus (Vars(f_a) \cup Vars(f_b))| \leq s$ from clauses $\widetilde{C}$ appearing in the CNF expansions of $f_a \to f_c$ and $f_c \to f_b$. Since $D(\tau \vdash \widetilde{C}) \leq (d-1)s$ for any such clause by the inductive assumption, the inductive step follows.

In particular, setting $d = \log_2 T$, $a = 0$, $b = T$, $C = 0$, we conclude that $D(\tau \vdash 0) \leq (2s)\log_2 T \leq 2\mathsf{VSpace}^*(\pi)$. □

Proof of 3.2. Take the configurational refutation $(\mathbb{C}_0, \mathbb{C}_1, \ldots, \mathbb{C}_T)$ of τ with total space $\mathsf{TSpace}(\tau \vdash 0)$ and convert it to the semantical form $(f_0, f_1, \ldots, f_T)$. Since $\mathcal{C}$ contains all CNFs and $\mu_{\mathcal{C}}$ does not exceed the circuit size, we conclude that $\mu_{\mathcal{C}}(f_t) \leq O(\mathsf{TSpace}(\mathbb{C}_t))$ and hence $\mu_{\mathcal{C}}\text{-}\mathsf{Space}(\tau \vdash 0) \leq \mathsf{O}(\mathsf{TSpace}(\tau \vdash 0))$. On the other hand, since $\mu_{\mathcal{C}}$ is bounded from below by the number of essential variables, for every semantical proof π we have $\mathsf{VSpace}(\pi) \leq s \stackrel{\text{def}}{=} \mu_{\mathcal{C}}\text{-}\mathsf{Space}(\pi)$. If π in addition is minimal, then the length is bounded by the overall number of circuits C in $\mathcal{C}$ that satisfy $\mu_{\mathcal{C}}(C) \leq s$ and hence, using again the condition on $\mu_{\mathcal{C}}$, have size $\leq s$. Since the number of circuits of size s is bounded (to be on the safe side) by $n^{O(s)}$, the bound $\mathsf{VSpace}^*(\pi) \leq O(s^2 \log n)$ follows. As VSpace^* and TSpace are equivalent up to a polynomial and $\log n$ factors, the same holds for $\mu_{\mathcal{C}}$-Space. □

5. Very small space

In this section we prove Theorems 3.4 and 3.5 (as we noted in the introduction, we defer our most challenging proof to the next section).

Proof of 3.4. As we already remarked, this is a variation on the proof of Theorem 3.1 (the $D(\tau \vdash 0) \leq 2\mathsf{VSpace}^*(\tau \vdash 0)$ part), except that instead of binary search we now do T-ary search for a suitable T. But this time our goal is to come up with a configurational refutation rather than a tree-like one. Hence, an inductive description would be somewhat awkward, and we frame the argument as a direct construction instead.

Let $\pi = (f_0, f_1, \ldots, f_S)$ be a semantical refutation from τ that has variable space $\leq s$. Assume w.l.o.g. that S is of the form $(T+1)^h - 1$ for an integer T, and for $t \in [0..S]$, let $(t_{h-1}, \ldots, t_0)$ be its $(T+1)$-ary representation, that is $t = \sum_{d=0}^{h-1} t_d (T+1)^d$. For $t > 0$, let $\mathrm{ord}(t)$ be the minimal d for which $t_d \neq 0$ (that is, the maximal d for which $(T+1)^d | t$). Let $t^{(k)}$ be the truncation of t by taking k most significant bits: $t^{(k)} \stackrel{\mathrm{def}}{=} \sum_{d=h-k}^{h-1} t_d (T+1)^d$. In particular, $t^{(0)} = 0$ and $t^{(h)} = t$. Let

$$\widehat{f_t} \stackrel{\mathrm{def}}{=} (f_0 \to f_{t^{(1)}}) \wedge (f_{t^{(1)}} \to f_{t^{(2)}}) \wedge \cdots \wedge (f_{t^{(h-1)}} \to f_t).$$

Clearly, $|Vars(\widehat{f_t})| \leq O(hs)$.

Let us now take a look at $\widehat{f}_{t+1}$. Denoting $k \stackrel{\mathrm{def}}{=} h - \mathrm{ord}(t+1)$, let us note that $(t+1)^{(k)} = (t+1)^{(k+1)} = \cdots = (t+1)^{(h-1)} = t+1$ since in all those cases we truncate zeros only. Hence we can remove from $\widehat{f}_{t+1}$ all trivial terms $f_{(t+1)^{(k)}} \to f_{(t+1)^{(k+1)}}, \ldots, f_{(t+1)^{(h-1)}} \to f_{t+1}$ and write it down simply as

$$\begin{aligned} \widehat{f}_{t+1} \equiv (f_0 \to f_{t^{(1)}}) \quad &\wedge \quad (f_{t^{(1)}} \to f_{t^{(2)}}) \wedge \cdots \\ &\wedge \quad (f_{t^{(k-2)}} \to f_{t^{(k-1)}}) \wedge (f_{t^{(k-1)}} \to f_{t+1}). \end{aligned}$$

Hence $\widehat{f_t} \wedge (f_t \to f_{t+1}) \models \widehat{f}_{t+1}$ and $(\widehat{f_0}, \widehat{f_1}, \ldots, \widehat{f_s})$ is also a semantical refutation from τ of the desired variable space $O(hs)$.

We convert it to a configurational resolution refutation as follows. First, for $t \leq t'$ denote by $\mathcal{C}(t, t')$ the straightforward CNF expansion of $f_t \to f_{t'}$. Next, let $\mathcal{C}_t \stackrel{\mathrm{def}}{=} \mathcal{C}(0, t^{(1)}) \cup \mathcal{C}(t^{(1)}, t^{(2)}) \cup \cdots \cup \mathcal{C}(t^{(h-1)}, t)$; this is our chosen CNF representation of the Boolean function $\widehat{f_t}$. Now the conversion is natural: to get from $\mathcal{C}_t$ to $\mathcal{C}_{t+1}$, we first download all axioms in $\tau[Vars(f_t) \cup Vars(f_{t+1})]$, then write down the brute-force inference

$$\begin{aligned} \text{(5.1)} \quad & \mathcal{C}(f_{t^{(k-1)}}, f_{t^{(k)}}), \ldots, \mathcal{C}(f_{t^{(h-1)}}, f_t), \tau[Vars(f_t) \cup Vars(f_{t+1})] \\ & \quad \vdash \mathcal{C}(f_{t^{(k-1)}}, f_{t+1}), \end{aligned}$$

and, finally, erase all clauses in the left-hand side. It remains to bound the depth of this refutation (recall Definition 2.3).

Every individual step (5.1) has depth $O(hs)$ as this is how many variables it involves. To get a bound on the depth of the tree formed by the inferences (5.1), we not that for every $\mathcal{C}(f_a, f_b)$ in the left-hand side either $\mathrm{ord}(b) < \mathrm{ord}(t+1)(= h-k)$: this happens for all configurations but $\mathcal{C}(f_{t^{(k-1)}}, f_{t^{(k)}})$, or $\mathrm{ord}(b) = \mathrm{ord}(t+1)$ and $b < t+1$ $(a = t^{(k-1)},\ b = t^{(k)},\ t_{h-k} \neq 0)$, or it is trivial and can be removed $(a = t^{(k-1)},\ b = t^{(k)},\ t_{h-k} = 0)$. Hence the depth of the proof tree defined by the inferences (5.1) is $O(hT)$, and the required overall bound $O(h^2sT)$ on depth follows. □

Proof of 3.5. Fix a configurational refutation $\pi = (\mathbb{C}_0, \mathbb{C}_1, \dots, \mathbb{C}_T)$ from the Induction Principle τ_n that has variable space s. Let us begin with a few generic remarks.

First, we can assume w.l.o.g. that for every $0 \leq t \leq T-1$, $\mathbb{C}_t$ does not contain the empty clause 0.

Next, let us call a clause *Bi-Horn* if it contains at most one occurrence of a positive literal and at most one occurrence of a negative literal. Since the set of bi-Horn clauses is closed under the Resolution rule, and all axioms in τ_n are bi-Horn, all clauses appearing in our refutation must be also bi-Horn. In other words, for every $t < T$, $\mathbb{C}_t$ must entirely consist of literals and implications of the form $x_i \to x_j$ $(i \neq j)$.

Next, for $t \leq T-1$ we can remove from $\mathbb{C}_t$ all clauses C with $D_t(C) \geq D(\pi)$ and still get a configurational refutation (this reduction corresponds to removing non-essential clauses in Ben-Sasson (2009)). Hence, we can assume that

$$D_t(C) \leq D(\pi) - 1,\ t \leq T-1, C \in \mathbb{C}_t. \tag{5.2}$$

Let us now return to the proof of Theorem 3.5. The configuration $\mathbb{C}_{T-1}$ must contain both literals $x_i, \bar{x}_i$ of some variable i. Let r be the maximal index for which x_r, viewed as a one-variable clause, appears anywhere in π, and let ℓ be the minimal index for which the clause $\bar{x}_\ell$ occurs there. Note that $\ell \leq i \leq r$, and hence $Vars(\mathbb{C}_{T-1})$ has a non-empty intersection with *both* $\{x_0, \dots, x_r\}$ and $\{x_\ell, x_{\ell+1}, \dots, x_n\}$.

Choose a such that

$$Vars(\mathbb{C}_a) \cap \{x_0, x_1, \ldots, x_r\} \neq \emptyset \wedge Vars(\mathbb{C}_a) \cap \{x_\ell, x_{\ell+1}, \ldots, x_n\} \neq \emptyset$$

while for $\mathbb{C}_{a-1}$ one of these properties is violated. By symmetry, we can assume w.lo.g. that $Vars(\mathbb{C}_{a-1}) \cap \{x_0, \ldots, x_r\} = \emptyset$.

Let us now apply to π the restriction $\rho_+ : \ x_0 \to 1, \ x_1 \to 1, \ldots, x_\ell \to 1$. It transforms τ_n to $\tau_{n-\ell-1}$, and since $\bar{x}_\ell$ appears somewhere in the refutation (and is killed by ρ_+), (5.2) implies that $D(\pi|_{\rho_+}) \leq D(\pi) - 1$.

Let us also apply to π the dual restriction $\rho_- : \ x_\ell \to 0, x_{\ell+1} \to 0, \ldots, x_n \to 0$. Then $\tau_n|_{\rho_-} = \tau_{\ell-1}$. Next, every clause C in $\mathbb{C}_{a-1}$ is a bi-Horn clause in the variables $\{x_{r+1}, \ldots, x_n\}$, and, by the definition of r, it may not be a positive literal. Hence C must contain a negative literal which, since $r \geq \ell$, implies $C|_{\rho_-} \equiv 1$. Thus, ρ_- sets to 1 all clauses in $\mathbb{C}_{a-1}$, and since $Vars(\mathbb{C}_b) \cap \{x_\ell, x_{\ell+1}, \ldots, x_n\} \neq \emptyset$ for all $b \geq a$, ρ_- reduces the *space* by at least one: $\mathsf{VSpace}(\pi|_{\rho_-}) \leq \mathsf{VSpace}(\pi) - 1$.

For the purpose of recursion, let $D(n, s)$ be the minimum depth of a configurational refutation of $\tau_{\lfloor n \rfloor}$ that has variable space $\leq s$. We have proved so far that

$$D(n,s) \geq \min_{0 \leq \ell \leq n} \{\max(D(n-\ell-1, s) + 1, D(\ell-1, s-1))\}.$$

Note that $D(n, s)$ is monotone in n. Hence, for any particular ℓ we have

$$\begin{aligned} &\max(D(n-\ell-1, s) + 1, D(\ell-1, s-1)) \\ &\quad \geq \min\left\{D(n - n^{1-1/s} - 2, s) + 1, D(n^{1-1/s}, s-1)\right\}, \end{aligned}$$

and we conclude that

$$D(n,s) \geq \min\left\{D(n - n^{1-1/s} - 2, s) + 1, D(n^{1-1/s}, s-1)\right\}. \tag{5.3}$$

This recurrence resolves to $D(n, s) \geq \Omega\left(n^{1/s}\right)$ since

$$n^{1/s} \leq \left(n - n^{1-1/s} - 2\right)^{1/s} + O(1)$$

(and $n^{1/s} = \left(n^{1-1/s}\right)^{1/(s-1)}$). □

6. A supercritical tradeoff between variable space and length

In this section we prove Theorem 3.3. While it is our most difficult result, its proof naturally splits into three fairly independent parts, and we present it in this modular way, interlaced with necessary definitions.

6.1. Multi-valued logic. Multi-valued logic is the instrument to argue about constraints over larger alphabets just in the same way the propositional logic reasons about Boolean constraints. While the bulk of our proof in this section is carried in the context of multi-valued logic over a large alphabet, we will be content with a purely semantical view and, accordingly, do not attempt to define any syntactic proof system. Our notation more or less follows (Alekhnovich *et al.* 2002, Section 4.3).

DEFINITION 6.1. *(cf. (Alekhnovich* et al. *2002, Definitions 4.5-4.7)). Let D be a finite domain. Instead of Boolean variables, we consider D-*valued variables *X_i ranging over the domain D. A* multi-valued function *$f(X_1,\ldots,X_n)$ is a mapping from D^n to* $\{0,1\}$. *Since the image here is still Boolean, the notions of a (multi-valued) satisfying assignment $\alpha \in D^n$ and the semantical implication $f \models g$ are generalized to the multi-valued logic straightforwardly. So does the definition of the set of essential variables $Vars(f)$.*

*A (D-*valued) literal *is an expression of the form X^P, where X is a (D-valued) variable and $P \subseteq D$ is such that $P \neq \emptyset$ and $P \neq D$. Allowing here also $D = 0$ or $D = P$, we obtain the definition of a* generalized *(D-*valued*)* literal. *A generalized literal X^P is semantically interpreted by the characteristic function of the set P. X^Q is a* weakening *of X^P if $P \subseteq Q$ or, equivalently, $X^P \models X^Q$.*

*A D-*valued clause *[term] is a disjunction [conjunction] of multi-valued literals corresponding to pairwise distinct variables. A* constraint satisfaction problem *(CSP) is simply a set of arbitrary multi-valued functions called in this context "constraints". The* width *of a constraint C is again $|Vars(C)|$, and a CSP is an k-*

CSP *if all constraints in it have width* $\leq k$. *A* semantical D-valued refutation *from a multi-valued CSP* η *and its variable space are defined exactly as in the Boolean case.*

Remark 6.2. *In this section we will be predominantly interested in constraints of width* ≤ 2. *If the reader find this too restrictive, it is perhaps worth reminding that a great deal of celebrated CSPs studied in the combinatorial optimization do have this form.*

6.2. Supercritical tradeoff against variable space 1. Our starting point is the following (quite weak) tradeoff. Before stating it, let us remind that according to our conventions, proofs of variable space 1 make perfect sense and are precisely those in which all configurations are representable by generalized literals.

Lemma 6.3. *For any finite domain* D, *there exists a* D*-valued 2-CSP* η *in four variables such that* η *is refutable in variable space 1, but any such refutation* π *must have length* $\geq \exp\left(D^{\Omega(1)}\right)$.

Proof. We begin with the observation that was apparently first made in Babai & Seress (1992): the symmetric group $\mathrm{Sym}(D)$ contains elements σ of exponential order. More specifically, let $p_1+\cdots+p_n \leq |D|-2 < p_1+\cdots+p_n+p_{n+1}$, where $p_1 < p_2 < \cdots < p_n < \cdots$ is the list of all prime numbers, take pairwise disjoint $D_i \subseteq D$ with $|D_i| = p_i$ $(1 \leq i \leq n)$, and let σ act cyclically on every D_i and identically on $D\backslash(D_1\cup\cdots\cup D_n)$. Let $\widetilde{P}$ be any transversal of the set $\{D_1,\ldots,D_n\}$, then the orbit of $\widetilde{P}$ in the induced action of σ on $\mathcal{P}(D)$ also has size $\geq \exp(|D|^{\Omega(1)})$. Denote its size by r, then all sets $\widetilde{P}, \sigma(\widetilde{P}), \sigma^2(\widetilde{P})\ldots, \sigma^{r-1}(\widetilde{P})$ are pairwise distinct and $\sigma^r(\widetilde{P}) = \widetilde{P}$. Since all these sets $\left\{\sigma^i(\widetilde{P}) \mid 0 \leq i \leq r-1\right\}$ also have the same size, they are moreover independent w.r.t. inclusion. Let now $P \stackrel{\text{def}}{=} \widetilde{P}\cup\{a\}$, where $a \in D\backslash(D_1\cup\cdots\cup D_n)$ is an arbitrary fixed element. Then (since $|D\backslash(D_1\cup\cdots\cup D_n)| \geq 2$) we additionally have that the $(2r)$ sets

$$\left\{\sigma^i(P) \mid 0 \leq i \leq r-1\right\},\ \left\{D\backslash\sigma^i(P) \mid 0 \leq i \leq r-1\right\} \tag{6.4}$$

are pairwise independent w.r.t. inclusion.

Let now X_0, X_1, X_2, X_3 be D-valued variables. The required 2-CSP η has the following constraints, where $Q \subseteq D$ is an arbitrary subset different from $\emptyset$ and D:

$$\begin{aligned}
&X_0^Q\\
&X_0^Q \to X_1^P\\
&X_2 = X_1\\
&X_3 = X_2\\
&X_1 = \sigma(X_3)\\
&X_3^{\sigma^{r/2}(P)} \to X_0^{D\setminus Q}.
\end{aligned}$$

The extra "buffer" variable X_3 here is needed due to the way a semantical proof is defined (see (2.2)). If we, say, would have included the axioms $X_2 = X_1,\ X_1 = \sigma(X_2)$ instead then a space one proof could have simultaneously downloaded *both* of them, and the subtle argument below would have completely fallen apart.

The refutation π from η with $\mathsf{VSpace}(\pi) = 1$ is straightforward:

$$\begin{gathered}
1, X_0^Q, X_1^P, X_2^P, X_3^P, X_1^{\sigma(P)}, \ldots, X_1^{\sigma^2(P)}, \ldots,\\
X_1^{\sigma^{r/2}(P)}, X_2^{\sigma^{r/2}(P)}, X_3^{\sigma^{r/2}(P)}, X_0^{D\setminus Q}, 0.
\end{gathered}$$

In order to prove the second statement in Lemma 6.3, we show that this refutation, its inverse and its contrapositive are essentially the only non-trivial inferences with variable space 1. More specifically, let

$$\begin{aligned}
\mathcal{L}_t \stackrel{\text{def}}{=} \{X_0^Q\} \quad &\cup \quad \left\{ X_i^{\sigma^h(P)} \mid i \in [3],\ h \in \mathbb{Z},\ |h| \le t-2 \right\}\\
&\cup \quad \left\{ X_i^{D\setminus\sigma^h(P)} \mid i \in [3],\ h \in \mathbb{Z},\ |h - r/2| \le t-2 \right\},
\end{aligned}$$

and let

$$\pi = 1, X_{i_1}^{A_1}, \ldots, X_{i_t}^{A_t}, 0$$

be a semantical refutation of variable space 1. We claim that as long as $t \le r/2$, $X_{i_t}^{A_t}$ is a weakening of a literal in $\mathcal{L}_t$.

Inductive base $t = 1$ is obvious since X_0^Q is the only constraint in η of width 1.

Inductive step.

Let $t \leq r/2$. We have to prove that if $X_i^A \in \mathcal{L}_t$, X_j^B is a generalized literal and

$$X_i^A \wedge \eta[\{X_i, X_j\}] \models X_j^B$$

then X_j^B is a weakening of a literal in $\mathcal{L}_{t+1}$. This is by a routine case analysis; the only case worth mentioning here is $i \in \{1, 3\}$ and $j = 0$, this is where we need the assumption $t \leq r/2$. By symmetry, assume that $i = 1$, then $\eta[\{X_0, X_1\}] \equiv X_0^Q \wedge X_1^P$. But since $X_i^A \in \mathcal{L}_t$ and $t \leq r/2$, we conclude that $A \cap P \neq \emptyset$ since all sets in (6.4) are independent w.r.t. inclusion. Hence $X_1^A \wedge \eta[\{X_0, X_1\}] \models X_0^B$ actually implies that $B \supseteq Q$ and thus X_0^B is a weakening of X_0^Q. $\square$

6.3. Supercritical tradeoffs against logarithmic variable space. In the previous section, we established a numerically strong tradeoff between length and space. On the negative side, it works only against proofs of space 1, but, to compensate for this, the CSP we constructed had only $O(1)$ variables. Our next task is to improve the space range from 1 to h, where h is an arbitrary parameter, but the prize we will have to pay for that is that the number of variables blows up to $\exp(O(h))$. The proof is based on a carefully designed iterative construction of unsatisfiable CSPs that we will call a lexicographic product and will essentially consist in showing a "direct product" theorem for the variable space. We would like to express a cautious hope that this construction may turn out to be of independent interest.

6.3.1. Combinatorial and geometric set-up. For integer parameters $h, \ell \geq 0$ that will be fixed throughout Section 6.3, we let $V \stackrel{\text{def}}{=} \{(i_1, \ldots, i_h) \mid i_\nu \in [\ell]\}$ be the set of all words of length h in the alphabet $[\ell]$. The set V can be alternately viewed as the set of all leaves of a complete ℓ-ary tree of height h, and it is equipped with the natural ultrametric: $\rho(u, v)$ is equal to h minus the length of the longest common prefix of u and v. We have found the geometric view of V as an ultrametric space more instructive for the proof in this section; the reader preferring the language of trees will hopefully have little difficulty with a translation.

We let V^+ be the set of all words u in the same alphabet ℓ such that $1 \leq |u| \leq h$. Its elements correspond to *non-root* vertices of the tree. Conveniently, elements of V^+ can be also naturally identified with *non-trivial* (that is, non-empty and different from the whole space V) balls in the ultrametric ρ. *By a ball we will always mean a non-trivial ball.* Let $r(\mathcal{B})$ be the radius of the ball $\mathcal{B}$, and let $\mathcal{B}(v, r)$ be the ball with center v and radius r.

To help the reader develop some intuition, we compile below simple properties of balls (and give a few definitions on our way) that will be used throughout Section 6.3. All of them immediately follow from the fact that ρ is an ultrametric.

1. Any two balls are either disjoint or one of them contains another.

2. The intersection of any family of balls is either empty or is again a ball.

3. If $u \in \mathcal{B}(v, r)$, then $\mathcal{B}(u, r) = \mathcal{B}(v, r)$. In other words, *every* point in a ball can be taken as its center.

4. If $\mathcal{B}$ and $\mathcal{B}'$ are disjoint balls, then $\rho(u, v)$ takes on the same value for all pairs $u \in \mathcal{B},\ v \in \mathcal{B}'$. We call it *the distance* between $\mathcal{B}$ and $\mathcal{B}'$ and denote by $\rho(\mathcal{B}, \mathcal{B}')$. The distance between two disjoint balls is always *strictly* larger than both $r(\mathcal{B})$ and $r(\mathcal{B}')$.

5. If $\mathcal{B}$ and $\mathcal{B}'$ are disjoint balls, then a ball containing both of them exists if and only if $\rho(\mathcal{B}, \mathcal{B}') < h$. In that case, the minimal ball with these properties is uniquely defined and has radius $\rho(\mathcal{B}, \mathcal{B}')$.

6. Let us call two disjoint balls $\mathcal{B}$ and $\mathcal{B}'$ *adjacent* if they have the same radius r and $\rho(\mathcal{B}, \mathcal{B}') = r + 1$. The relation "two balls are either the same or adjacent" is an equivalence relation.

7. Every partition of V into balls contains at least one equivalence class of this relation.

We conclude with a less trivial combinatorial lemma that will be crucial for our proof.

DEFINITION 6.5. *For a ball $\mathcal{B}$ with $r(\mathcal{B}) \leq h-1$, $\mathcal{B}^+$ is the uniquely defined ball of radius $r(\mathcal{B})+1$ such that $\mathcal{B}^+ \supset \mathcal{B}$.*

LEMMA 6.6. *Any set $V_0 \subseteq V$ with $|V_0| \leq h-2$ can be covered by a collection of balls $\{\mathcal{B}_1, \ldots, \mathcal{B}_w\}$ of radii $\leq h-1$ such that the balls $(\mathcal{B}_1)^+, \ldots, (B_w)^+$ are pairwise disjoint.*

PROOF. Let us call a covering of the set V_0 by pairwise disjoint balls *frugal* if every ball $\mathcal{B}$ in this covering covers at least $r(\mathcal{B})+1$ elements of V_0. Frugal coverings do exist: take, for example, the trivial covering by balls of radius 0. Now pick up a frugal covering with the smallest possible number of balls. We claim that it has all the required properties.

Indeed, the bound $r(\mathcal{B}) \leq h-1$ for a ball $\mathcal{B}$ in our frugal covering simply follows from the definition of frugality and the bound $|V_0| \leq h-2$. Next, if $\mathcal{B}, \mathcal{B}'$ are two different balls in this coloring such that $\mathcal{B}^+ \cap (\mathcal{B}')^+ \neq \emptyset$, then one of these latter balls must contain another, say, $\mathcal{B}^+ \supseteq (\mathcal{B}')^+ \supset \mathcal{B}'$. Replacing $\mathcal{B}$ with $\mathcal{B}^+$ and removing all balls contained in $\mathcal{B}^+$ (including $\mathcal{B}'$!), we will get a frugal covering with a smaller number of balls, a contradiction. Thus, all the balls in the minimal frugal covering are pairwise disjoint. □

6.3.2. Lexicographic products and the main lemma. Let $D_1, \ldots, D_h$ be pairwise disjoint finite domains, and let $\eta_1(X_1, \ldots, X_\ell), \ldots, \eta_h(X_1, \ldots, X_\ell)$ be CSPs, where η_d is D_d-valued. For notational simplicity, we assume that these CSPs have the same number of variables ℓ. In our application of the construction, they will be actually identical (namely, the CSP constructed in the previous section 6.2), and in particular all alphabets D_d's will also be the same. It would have been notationally unwise, however, to identify the domains $D_1, \ldots, D_h$ as well, so we keep them separate and pairwise disjoint.

In this set-up, we are going to define a lexicographic product of $\eta_1, \ldots, \eta_h$, and we begin with an informal description of its intended meaning and the intuition behind it. We start off with a simple naive attempt that does *not* work, and then we will explain what is the problem and how to fix it.

Let us introduce variables $Y_{\mathcal{B}}$ for balls $\mathcal{B}$. The domain of $Y_{\mathcal{B}}$ will be $D_{r(\mathcal{B})+1} \cup \{*\}$, where the intended meaning of $*$ is "undefined". Thus, an assignment to these variables consists of a set $\{\mathcal{B}_1, \ldots, \mathcal{B}_w\}$ of balls and an assignment of the corresponding variables to values in $D_{r(\mathcal{B}_1)+1}, \ldots, D_{r(\mathcal{B}_w)+1}$; all other variables $Y_{\mathcal{B}}$ receive the value $*$. We are interested only in those assignments (let us call them "legal") for which the family $\{\mathcal{B}_1, \ldots, \mathcal{B}_w\}$ makes a *partition* of the space V. Then, as we noted in Section 6.3.1, for every legal assignment there is an equivalence class of the adjacency relation in which all the corresponding variables $Y_{\mathcal{B}}$ are defined, i.e., take values in $D_{r(\mathcal{B})+1}$.

Let us now describe how the CSPs $\eta_1, \ldots, \eta_h$ are used to generate local constraints on legal assignments. Fix a ball $\mathcal{B}$ with $d \stackrel{\text{def}}{=} r(\mathcal{B}) > 0$, and let $\mathcal{B}_1, \ldots, \mathcal{B}_\ell$ be all its (pairwise adjacent) subballs of radius $d-1$. We want to "lift" constraints in $\eta_d(X_1, \ldots, X_\ell)$ to constraints in the variables $Y_{\mathcal{B}_1}, \ldots, Y_{\mathcal{B}_\ell}$ inside this ball. As η_d are all unsatisfiable, it will imply that there is no legal assignment satisfying all these constraints. The only subtle issue is how to treat $*$ values, or, in other words, how to lift a D_h-valued constraint $C(X_{i_1}, \ldots, X_{i_w})$ to a $(D_h \cup \{*\})$-valued constraint $\widehat{C}(Y_1, \ldots, Y_w)$.

It turns out that in order to preserve the bound from Section 6.2, we should do it in the most minimalistic way possible and accept any assignment in which at least one value is $*$. The intuition (that will be made rigorous in the forthcoming subsections) is that, roughly speaking, in this way every variable $Y_{\mathcal{B}_i}$ can "store" information only about itself and of (by the token of legality) larger balls. It does not have any bearing on adjacent balls. This will prevent the refutation from cheating: if, say, it keeps in memory a literal of a variable $Y_{\mathcal{B}_i}$, and wants to apply a semantical inference $Y_{\mathcal{B}_i} \to Y_{\mathcal{B}_j}$, then it is quite possible that instead of following D_d-rules prescribed by this inference, the ball $\mathcal{B}_j$ gets shattered into smaller balls, and we will have to treat them from the scratch, i.e., recursively. As the variable space will be assumed to be small, we will not have enough variables at any given moment to "guard" all levels (Lemma 6.6 will help to make this precise), and the only choice "should be" to proceed lexicographically.

The only problem with this plan lies in the concept of legality. The assumption that balls do not intersect is local and presents no problems. However, the assumption that every point $v \in V$ is covered is represented by the constraint $\bigvee_{\mathcal{B} \ni v}(Y_{\mathcal{B}} \neq *)$ which is prohibitively wide. We circumvent this by replacing $Y_{\mathcal{B}}$ with "extension variables" X_v for any $v \in V$ that will carry the information which of those $Y_{\mathcal{B}}$ is different from $*$, and what is its value. Thus, X_v must be D-valued, where $D = D_1 \dot\cup \cdots \dot\cup D_h$, and then $Y_{\mathcal{B}}$ can be "retrieved" from X_v by

$$Y_{\mathcal{B}} = \begin{cases} X_v & \text{if } X_v \in D_{r(\mathcal{B})+1} \\ * & \text{otherwise.} \end{cases}$$

Naturally, we will need explicit consistency constraints saying that these "intended values" do not depend on the choice of $v \in \mathcal{B}$, and we will sometimes say that a variable X_v is *assigned* in a ball $\mathcal{B} \ni v$ if its assignment is in $D_{r(\mathcal{B})+1}$.

We will now proceed to formal definitions and the statement of the main lemma of Section 6.3.

DEFINITION 6.7. *Let $D_1, \ldots, D_h$ be pairwise disjoint finite sets, $D \stackrel{\text{def}}{=} D_1 \dot\cup \cdots \dot\cup D_h$, and let $\eta_1(X_1, \ldots, X_\ell), \ldots, \eta_h(X_1, \ldots, X_\ell)$ be CSPs, where η_d is D_d-valued. We define their* lexicographic product *$\eta_h \cdot \eta_{h-1} \cdot \ldots \cdot \eta_1$ that will be a D-valued CSP in the variables $(X_v \mid v \in V)$ as follows.*

(i) For $1 \leq d \leq h-1$, let $Con_d(X, Y)$ be the conjunction of the formulas $X^{\{a\}} \equiv Y^{\{a\}}$, where $a \in D_{d+1} \cup \cdots \cup D_h$. We include into $\eta_h \cdot \eta_{h-1} \cdot \ldots \cdot \eta_1$ the constraints $Con_d(X_u, X_v)$ for all $u, v \in V$ with $\rho(u, v) = d$.

Informally, if one of the variables X_u, X_v was assigned in a ball where both u and v belong, then the other variable must be also assigned to the same value. Otherwise, the constraint is vacuous.

(ii) Let $C(X_{i_1}, \ldots, X_{i_w})$ be a (D_d-valued) constraint in η_d. We let the formula $\widehat{C}(Y_1, \ldots, Y_w)$ of the D-valued logic be defined as

$$\bigwedge_{\nu=1}^{w} Y_\nu^{D_d} \Longrightarrow C(Y_1, \dots, Y_w) \tag{6.8}$$

(the right-hand side here makes sense due to the premise $\bigwedge_{\nu=1}^{w} Y_\nu^{D_d}$*). We add to* $\eta_h \cdot \eta_{h-1} \dots \cdot \eta_1$ *all axioms of the form* $\widehat{C}(X_{u_1}, \dots, X_{u_w})$ *as long as* $\rho(u_\nu, u_\mu) = d$ *for all* $\nu \neq \mu$ *(in particular,* $u_1, \dots, u_w$ *share a common prefix of length* $h-d$*) and* $(u_\nu)_{h-d+1} = i_\nu$ $(1 \leq \nu \leq w)$*.*

Informally, if $u_1, \dots, u_w$ *belong to pairwise adjacent balls of radius* $d-1$ *and* all *variables* $X_{u_1}, \dots, X_{u_w}$ *are assigned in these balls, then their assignments must satisfy all applicable constraints in* η_d*. If at least one of these variables is assigned outside of* D_d*, the constraint is vacuous.*

Note that if all $\eta_1, \dots, \eta_h$ are 2-CSP (which is the case we are mostly interested in), then their lexicographic product is also a 2-CSP.

We are now ready to formulate the main result of this section.

LEMMA 6.9. *Assume that* $\eta_1(X_1, \dots, X_\ell), \dots, \eta_h(X_1, \dots, X_\ell)$ *are multi-valued 2-CSPs such that* $\mathsf{VSpace}(\eta_d \vdash 0) = 1$ $(d \in [h])$ *but any refutation* π *of* η_d *with* $\mathsf{VSpace}(\pi) = 1$ *must have length* $> T$*, and let* $\eta \overset{\text{def}}{=} \eta_h \cdot \dots \cdot \eta_1$ *be their lexicographic product. Then* $\mathsf{VSpace}(\eta \vdash 0) = 1$ *(in particular,* η *is a contradiction), but any its refutation* π *with* $\mathsf{VSpace}(\pi) \leq h/2 - 1$ *must also have length* $\geq T$*.*

For the rest of Section 6.3 we fix $\eta_1, \dots, \eta_h, \eta$ *and the length upper bound* T *as in the statement of Lemma 6.9.*

6.3.3. Upper bound. In this section we prove that $\mathsf{VSpace}(\eta \vdash 0) = 1$. We follow the intuition outlined at the beginning of Section 6.3.2 and construct the intended refutation lexicographically. As we do not attempt to store information about different levels, its meagre amount we do carry around can be easily fit into a single literal. This will allow us to stay within space 1.

For every $d \in [h]$ fix a refutation

$$\pi_d = 1, X_{i(d,1)}^{A(d,1)}, \dots, X_{i(d,T-1)}^{A(d,T-1)}, 0$$

of length T, where $i(d,t) \in [\ell]$ and $A(d,t) \subseteq D_d$. For the uniformity of notation, we also let $i(d,0) \stackrel{\text{def}}{=} i(d,1)$, $A(d,0) \stackrel{\text{def}}{=} D_d$ and, likewise, $i(d,T_d) \stackrel{\text{def}}{=} i(d,T_d-1)$, $A(d,T_d) \stackrel{\text{def}}{=} \emptyset$. Denote by $L(d,t) \stackrel{\text{def}}{=} X_{i(d,t)}^{A(d,t)}$ $(t = 0..T)$ the corresponding generalized literal.

For $\vec{t} = (t_h, \ldots, t_{h-1}, \ldots, t_1) \in [0..T]^h$, let $v(\vec{t}) \stackrel{\text{def}}{=} (i(h,t_h), i(h-1,t_{h-1}), \ldots, i(1,t_1)) \in V$ (this is a good place to recall that we enumerate everything from the leaves to the root!) and $L(\vec{t}) \stackrel{\text{def}}{=} X_{v(\vec{t})}^{A(h,t_h)\cup\cdots\cup A(1,t_1)}$ be the corresponding generalized D-valued literal. We claim that the sequence of generalized literals $L(\vec{t})$, taken in the lexicographic order, makes a refutation of η.

Indeed, $L(0,0,\ldots,0) = X_{v(0,\ldots,0)}^{D} \equiv 1$ and $L(T,\ldots,T) = X_{v(T,\ldots,T)}^{\emptyset} \equiv 0$, as required. Given $\vec{t} \neq (T,\ldots,T)$, let $d \in [h]$ be the smallest index such that $t_d \neq T$, say $t_d = s$, so that the next term in the lexicographic order is $\vec{t'} \stackrel{\text{def}}{=} (t_h, \ldots, t_{d+1}, s+1, 0 \ldots, 0)$. We have $L(\vec{t}) = X_{v(\vec{t})}^{B\cup A(d,s)}$ and $L(\vec{t'}) = X_{v(t')}^{B\cup A(d,s+1)\cup D_{d-1}\cup\cdots\cup D_1}$ for the same $B \subseteq D_h \cup \cdots \cup D_{d+1}$.

From the refutation π_d we know that $X_{i(d,s)}^{A(d,s)} \wedge \eta_d [\{X_{i(d,s)}, X_{i(d,s+1)}\}] \models X_{i(d,s+1)}^{A(d,s+1)}$ in the D_d-valued logic. Then $\eta [\{X_{v(\vec{t})}, X_{v(\vec{t'})}\}]$ entails, due to the second group of axioms, that $(X_{v(\vec{t})}^{D_d} \wedge X_{v(\vec{t})'}^{D_d}) \to \eta_d[X_{\vec{t}}, X_{\vec{t'}}]$. Also, as long as $v(\vec{t}) \neq v(\vec{t'})$, $\eta[\{X_{v(\vec{t})}, X_{v(\vec{t'})}\}]$ also contains the first group of axioms $\mathrm{Con}_d(X_{v(\vec{t})}, X_{v(\vec{t'})})$. The required implication

$$X_{v(\vec{t})}^{B\cup A(d,s)} \wedge \eta[\{X_{v(\vec{t})}, X_{v(\vec{t'})}\}] \models X_{v(t')}^{B\cup A_{d,s+1}\cup D_{d-1}\cup\cdots\cup D_1}$$

follows straightforwardly. This completes the proof of $\mathsf{VSpace}(\tau \vdash 0) \leq 1$.

6.3.4. Lower bound. Our overall strategy is quite typical for space complexity: we define a collection of "admissible" configurations $\mathbb{A}$ that is simple enough to be controlled and, on the other hand, everything that we can infer in small space can be majorated by an admissible configuration from $\mathbb{A}$. The only twist is that since we are proving a *length* lower bound, this construction must necessarily be dynamic as well and consist of an increasing sequence

$\mathbb{A}_0 \subseteq \mathbb{A}_1 \subseteq \cdots \subseteq \mathbb{A}_s \subseteq \cdots$, where configurations in $\mathbb{A}_s$ majorate everything that can be inferred in small space *and* length $\leq s$. A relatively simple implementation of this idea was already used in the proof of Lemma 6.3.

Let us now start the formal argument.

DEFINITION 6.10 (normal terms). *Let $\mathcal{B}$ be a ball of radius r, $0 \leq r \leq h-1$, and let $A \subseteq D_{r+1}$ be such that $A \neq \emptyset$ and, moreover, $A \neq D_1$ if $r = 0$. Then we denote by $t_{\mathcal{B},A}$ the following term:*

$$t_{\mathcal{B},A} \stackrel{\mathrm{def}}{=} \bigwedge_{v\in\mathcal{B}} X_v^{D_h\cup\cdots\cup D_{r+2}\cup A}. \tag{6.11}$$

A term t is normal *if it can be represented as*

$$t = t_{\mathcal{B}_1,A_1} \wedge \cdots \wedge t_{\mathcal{B}_w,A_w}, \tag{6.12}$$

where all balls are pairwise disjoint.

REMARK 6.13. *For any D-valued literal X_v^B, the set B* uniquely *determines the term $t_{\mathcal{B},A}$ in which it may possibly appear. Hence the representation (6.12) of a normal term is unique and it what follows we will not distinguish between the two.*

REMARK 6.14. *The reader willing to compare this definition with the informal discussion at the beginning of Section 6.3.2, should be aware that the term $t_{\mathcal{B},A}$ corresponds* not *to the literal $Y_{\mathcal{B}}^A$ but rather to its "monotone closure" asserting that* either *$Y_{\mathcal{B}} \in A$* or *$\mathcal{B}$ is a* proper *sub-ball of a member $\mathcal{B}'$ of the partition; the exact value of $Y_{\mathcal{B}'}$ is irrelevant as long as it is not $*$. This monotonicity reflects the inherently lexicographic nature of all our definitions and proofs.*

DEFINITION 6.15 (sparse terms). *A normal term t as given in (6.12) is* sparse *if no two balls $\mathcal{B}_i, \mathcal{B}_j$ in it are adjacent.*

DEFINITION 6.16 (complexity of normal terms). *Let a ball $\mathcal{B}$ of radius r corresponds to a prefix $(i_h, \ldots, i_{r+1}) \in V^+$, $i_\nu \in [\ell]$. For $A \subseteq D_{r+1}$, let $L(t_{\mathcal{B},A})$ be the minimal length of a space 1 D_{r+1}-*

valued proof of the generalized literal $X^A_{i_{r+1}}$ *from* η_{r+1}. *For a normal term* (6.12), *we let* $L(t) \stackrel{\rm def}{=} \max_{1\leq j\leq w} L(t_{\mathcal{B}_j,A_j})$.

Now we are ready to define the sets of admissible configurations $\mathbb{A}_s$.

Definition 6.17 (admissible configurations). *For a term* t, *we let*

$$t^* \stackrel{\rm def}{=} t \wedge \eta[Vars(t)].$$

We let $\mathbb{A}_s$ *consist of all* t^*, *where* t *is a normal* sparse *term with* $L(t) \leq s$. *Note that obviously* $1 = 1^* \in \mathbb{A}_0 \subseteq \mathbb{A}_1 \subseteq \mathbb{A}_2 \subseteq \cdots$.

Now, to get a better feeling of all these definitions we start with the (simpler) end task. As this kind of reasoning will be recurrent in the proof of much more difficult Lemma 6.19, we outline the argument in perhaps more meticulous way than it deserves.

Lemma 6.18. $0 \not\in \mathbb{A}_{T-1}$.

Proof. Let $t = t_{\mathcal{B}_1,A_1} \wedge \cdots \wedge t_{\mathcal{B}_w,A_w}$ be a normal sparse term with $L(t) \leq T-1$, that is such that $L(t_{\mathcal{B}_j,A_j}) \leq T-1$ for all j; we need to construct an assignment satisfying t^*.

Note first that the property $L(t_{\mathcal{B}_j,A_j}) \leq T-1$ in particular implies that there exists $a_j \in A_j$ satisfying $\eta_d[X^{A_j}_{i(j)}]$, where $d \stackrel{\rm def}{=} r(\mathcal{B}_j)+1$ and $X_{i(j)}$ is the D_d-valued variable corresponding to the ball $\mathcal{B}_j$ as in Definition 6.16. Indeed, otherwise we could have inferred 0 from $X^{A_j}_{i(j)}$ in just one step, in contradiction with the assumption in Lemma 6.9 that η_d does not possess any variable space 1 length T refutations.

Now, assign all variables X_v with $v \in \mathcal{B}_j$ to a_j and assign all other variables (they do not appear in t^*) arbitrarily. Then this assignment clearly satisfies t and it also satisfies all the consistency axioms $\mathrm{Con}_d(X_u, X_v)$ for $X_u, X_v \in Vars(t)$. For the latter, note that if u and v are in the same ball $\mathcal{B}_j$, the axiom is satisfied as X_u, X_v are assigned to the same value a_j, and if X_u, X_v are in different balls, then it is satisfied vacuously as *both* X_u and X_v take values in $D_1 \cup \cdots \cup D_d$ (in fact, even in $D_1 \cup \cdots \cup D_{d-1}$ due to sparsity but we do not need it).

As for the axioms $\widehat{C}(X_{u_1}, \ldots, X_{u_w})$ with $\{X_{u_1}, \ldots, X_{u_w}\} \subseteq Vars(t)$, we invoke the sparsity condition that implies that either all $X_{u_1}, \ldots, X_{u_w}$ are in the same ball $\mathcal{B}_j$ or $w = 1$. In the first case, $\widehat{C}$ is vacuously satisfied due to the left-hand side in (6.8). In the second case, let, say, $u_1 \in \mathcal{B}_1$. The axiom $\widehat{C}(X_{u_1})$ may actually come *from an arbitrary level* d, not necessarily the relevant one $d = r(\mathcal{B}_1) + 1$. But if $d = r(\mathcal{B}_1) + 1$, then C is an axiom of τ_d and hence is satisfied by a_1 due to its choice. On the other hand, if $d \neq r(\mathcal{B}_1) + 1$, then $\widehat{C}(a_1) = 1$ simply because it becomes vacuous according to (6.8). Hence the assignment we have constructed satisfies $\eta[Vars(t)]$ as well. $\square$

The lower bound in Lemma 6.9 now readily follows from the following, which is the heart of our argument.

LEMMA 6.19. *Let* $1 = f_0, \ldots, f_1 \ldots, f_s$ *be a D-valued semantical proof from* $\eta_h \cdot \eta_{h-1} \cdot \ldots \cdot \eta_1$ *of variable space* $\leq h/2 - 1$ *with* $s \leq T - 1$. *Then there exists* $f \in \mathbb{A}_s$ *such that* $f \models f_s$.

PROOF OF 6.19. By induction on s. The base case $s = 0$ is obvious.

For the inductive step, let $t = t_{\mathcal{B}_1, A_1} \wedge \cdots \wedge t_{\mathcal{B}_w, A_w}$ be a normal sparse term such that $L(t) \leq s$, where $s \leq T - 2$, with $t^* \models f_s$. Our goal is to construct a normal sparse term $\widehat{t}$ such that $L(\widehat{t}) \leq s + 1$ and $\widehat{t}^* \models f_{s+1}$. This will complete the proof of Lemma 6.19.

Let[7] $V_0 \stackrel{\text{def}}{=} Vars(f_s) \cup Vars(f_{s+1})$; note that $|V_0| \leq 2(h/2 - 1) = h - 2$. We have $f_s \wedge \eta[V_0] \models f_{s+1}$, and hence it is sufficient to construct a normal sparse term $\widehat{t}$ with $L(\widehat{t}) \leq s + 1$ satisfying

$$\widehat{t}^* \models f_s \wedge \eta[V_0]. \tag{6.20}$$

Fix a collection of balls $\{\mathcal{B}_1^*, \ldots, \mathcal{B}_{w^*}^*\}$ satisfying the conclusion of Lemma 6.6. This collection is not a priori related anyhow to the collection $\{\mathcal{B}_1, \ldots, \mathcal{B}_w\}$ underlying the normal sparse term t, and our task is to merge the two. Let us start with a brief intuitive explanation of what we are going to do and what we hope to

[7] From this point on we freely identify sets of variables and their indices whenever it does not create confusion.

achieve: the rest of the proof of Lemma 6.19 will basically consist of implementing these ideas (and circumventing a few technical difficulties). Before reading the explanation below, the reader is strongly encouraged to refresh elementary properties of balls we stated in Section 6.3.1.

The terms $t_{\mathcal{B}_1,A_1},\ldots,t_{\mathcal{B}_w,A_w}$ carry all the information our refutation has been able to achieve so far and the balls $\{\mathcal{B}_1^*,\ldots,\mathcal{B}_{w^*}^*\}$ contain all variables that are kept in memory at the moment. Unlike $\mathcal{B}_1,\ldots,\mathcal{B}_w$, these latter balls are new-born and are not equipped with any particular terms yet. Let us, however, postpone this issue and discuss geometry first.

Ideally, we would like to keep in the game all balls

$$\{\mathcal{B}_1,\ldots,\mathcal{B}_w,\mathcal{B}_1^*,\ldots,\mathcal{B}_{w^*}^*\} \tag{6.21}$$

but this is of course impossible because of potential conflicts between them. These conflicts can be of one of the two types: *containment conflicts* and *adjacency conflicts*.

Due to the fact that both collections of balls are individually conflict-free (and $\{\mathcal{B}_1^*,\ldots,\mathcal{B}_{w^*}^*\}$ satisfies even a much stronger property), the picture is actually less chaotic than it may appear on the first glance. No ball in (6.21) may *both* contain another ball and be (properly) contained in one. The adjacency relation, when restricted to (6.21), is a *matching*. Yet another important observation is that no ball in $\{\mathcal{B}_1,\ldots,\mathcal{B}_w\}$ may be involved in *both* containment and adjacency conflicts. Indeed, if $\mathcal{B}_\gamma$ were in an adjacency conflict with $\mathcal{B}_\mu^*$ and in a containment conflict with $\mathcal{B}_{\mu'}^*$ then $(\mathcal{B}_\mu^*)^+ \supset \mathcal{B}_\gamma$ would have a non-empty intersection with $\mathcal{B}_{\mu'}^*$ which is impossible due to the way these balls were chosen (see the statement of Lemma 6.6). This may happen for a new ball $\mathcal{B}_\mu^*$, though; this issue will be addressed when we will talk about designing sets A_μ^* for those balls.

With these remarks in mind, we resolve all containment conflicts in favor of the larger ball and remove the smaller. Adjacency conflicts $(\mathcal{B}_\gamma,\mathcal{B}_\mu^*)$ are always resolved in favor of $\mathcal{B}_\mu^*$: $\mathcal{B}_\gamma$ gets removed. This will define the geometry of the term $\widehat{t}$, it will be normal and sparse as we have resolved all conflicts, and it will contain all variables in V_0 due to the way we have resolved them.

It remains to explain how to transfer information from a term $t_{\mathcal{B}_\gamma, A_\gamma}$ to the new generation of balls if $\mathcal{B}_\gamma$ was slated for extermination, and we still have to define A^*_μ for those new balls $\mathcal{B}^*_\mu$. There are two reasons why $\mathcal{B}_\gamma$ can be removed: since $\mathcal{B}_\gamma \subset \mathcal{B}^*_\mu$ for some μ (containment conflict) or because it is adjacent to some $\mathcal{B}^*_\mu$. In the first case we need not worry, the information will be automatically passed to $\mathcal{B}^*_\mu$ due to monotonicity (cf. Remark 6.14). The second case is more interesting, and the rest of the proof will essentially consist in showing that this is the only way the refutation can do anything meaningful, namely apply an inference $X_i^{A_\gamma} \models X_j^A$ from η_d, for d, i, j chosen in an obvious way. But then among all potential As with this property there is a minimal one (their intersection), and we adorn the ball $\mathcal{B}^*_\mu$ with this A, thus converting it into a normal term. It is crucially important for this step, of course, that $\mathcal{B}_\gamma$ is the only ball adjacent to $\mathcal{B}^*_\mu$; otherwise, the argument would have completely broken apart.

Let us now continue with a formal argument.

Let $\Gamma_0 \subseteq [w]$ consist of those γ for which $\mathcal{B}_\gamma$ is *properly* contained in a ball $\mathcal{B}^*_\mu$ ($\mu \in [w^*]$). Let $M_0 \subseteq [w^*]$ be the set of all those μ for which $\mathcal{B}^*_\mu$ is contained (*not* necessarily properly) in one of the $\mathcal{B}_\gamma$ ($\gamma \in [w]$). By ultrametricity, all balls $\{\mathcal{B}_\gamma \mid \gamma \not\in \Gamma_0\}$, $\left\{\mathcal{B}^*_\mu \mid \mu \not\in M_0\right\}$ are pairwise disjoint. They still may contain adjacent balls, though.

Let $\Gamma_1 \subseteq [w]$ be the set of all balls $\mathcal{B}_\gamma$ such that $\mathcal{B}_\gamma \subseteq (\mathcal{B}^*_\mu)^+ \backslash \mathcal{B}^*_\mu$ for at least one $\mu \in [w^*]$. Note that if $\mathcal{B}_\gamma$ is adjacent to a ball $\mathcal{B}^*_\mu$ then $\gamma \in \Gamma_1$. Also, since all $(\mathcal{B}^*_\mu)^+$ ($\mu \in [w^*]$) are pairwise disjoint, it follows that for any $\gamma \in \Gamma_1$, $\mathcal{B}_\gamma$ is disjoint with *all* balls $\mathcal{B}^*_\mu$ ($\mu \in [w^*]$), *including* $\mu \in M_0$. Hence, in particular, $\Gamma_0 \cap \Gamma_1 = \emptyset$. Moreover, the ball $\mathcal{B}^*_\mu$ with $\mathcal{B}_\gamma \subseteq (\mathcal{B}^*_\mu)^+ \backslash \mathcal{B}^*_\mu$ is uniquely defined, and we let

$$\Gamma_1^\mu \stackrel{\text{def}}{=} \left\{ \gamma \,\middle|\, \mathcal{B}_\gamma \subseteq (\mathcal{B}^*_\mu)^+ \backslash \mathcal{B}^*_\mu \right\};$$

thus, $\Gamma_1 = \dot{\bigcup}_{\mu \in [w^*]} \Gamma_1^\mu$. A word of warning: Γ_1^μ may be non-empty even if $\mu \in M_0$ (more precisely, when $\mathcal{B}^*_\mu = \mathcal{B}_{\gamma'}$ for some $\gamma' \in [w]$, see Claim 6.24 below).

Now, the balls $\{\mathcal{B}_\gamma \mid \gamma \not\in \Gamma_0 \cup \Gamma_1\}$, $\left\{\mathcal{B}^*_\mu \mid \mu \not\in M_0\right\}$ are not only pairwise disjoint but also (due to the definition of Γ_1) pairwise non-adjacent. They will make the support of the sparse term $\widehat{t}$ we are constructing, that is

$$\widehat{t} = \bigwedge_{\gamma \notin \Gamma_0 \cup \Gamma_1} t_{\mathcal{B}_\gamma, A_\gamma} \wedge \bigwedge_{\mu \notin M_0} t_{\mathcal{B}^*_\mu, A^*_\mu}, \tag{6.22}$$

where for $\mu \notin M_0$ the sets A^*_μ are defined as follows. Let $\mu \notin M_0$ and $r \stackrel{\text{def}}{=} r(\mathcal{B}^*_\mu)$.

Case 1. $r(\mathcal{B}_\gamma) < r$ for any $\gamma \in \Gamma_1^\mu$ (which in particular includes the case $\Gamma_1^\mu = \emptyset$).

We simply let $A^*_\mu \stackrel{\text{def}}{=} D_{r+1}$ unless $r = 0$ in which case, due to our convention, we simply remove $t_{\mathcal{B}^*_\mu, A^*_\mu}$ from (6.22).

Case 2. There exists $\gamma \in \Gamma_1^\mu$ with $r(\mathcal{B}_\gamma) = r$.

First note that γ with this property is unique since t is sparse. $\mathcal{B}_\gamma$ and $\mathcal{B}^*_\mu$ are defined by two prefixes of the form $(i_h, i_{h-1}, \ldots, i_{r+2}, i)$ and $(i_h, i_{h-1}, \ldots, i_{r+2}, j)$ with $i \neq j$. We let A^*_μ be the *minimal* subset of D_{r+1} for which

$$X_i^{A_\gamma} \wedge \eta[X_i, X_j] \models X_j^{A^*_\mu} \tag{6.23}$$

in the D_{r+1}-valued logic. We note that $L(t_{\mathcal{B}_\mu, A^*_\mu}) \leq s+1$ and hence (this is quite essential for the upcoming argument!) $A^*_\mu \neq \emptyset$ due to the assumption $s \leq T - 2$.

This completes the construction of the sparse term $\widehat{t}$, and all that remains is to prove (6.20). Let us first state formally a few observations that we already made in our informal explanation above.

Claim 6.24. *If $\Gamma_1^\mu \neq \emptyset$ then $\widehat{t}$ contains a sub-term of the form $t_{\mathcal{B}^*_\mu, A}$ for some A.*

Proof of 6.24. If $\mu \notin M_0$, this is obvious. If $\mu \in M_0$ then $\mathcal{B}^*_\mu \subseteq \mathcal{B}_\gamma$ for some γ and there is another γ' with $\mathcal{B}_{\gamma'} \subseteq (\mathcal{B}^*_\mu)^+ \setminus \mathcal{B}^*_\mu$. We necessarily must have $\mathcal{B}^*_\mu = \mathcal{B}_\gamma$ (otherwise, $\mathcal{B}_{\gamma'} \subseteq \mathcal{B}_\gamma$). Clearly, $\gamma \notin \Gamma_0 \cup \Gamma_1$ and hence $t_{\mathcal{B}_\gamma, A_\gamma} = t_{\mathcal{B}^*_\mu, A_\gamma}$ appears in $\widehat{t}$. $\square$

Claim 6.25. $V_0 \subseteq Vars(\widehat{t})$.

Proof of 6.25. Every $v \in V_0$ is contained in one of the balls $\mathcal{B}^*_\mu$. If $\mu \notin M_0$, we are done; otherwise, there exists $\gamma \in [w]$ with $\mathcal{B}^*_\mu \subseteq \mathcal{B}_\gamma$. Like in the proof of Claim 6.24, $\gamma \notin \Gamma_0 \cup \Gamma_1$, hence $t_{\mathcal{B}_\gamma, A_\gamma}$ appears in $\widehat{t}$ and thus $v \in Vars(\widehat{t})$. □

As an immediate consequence, the second part of the implication in (6.20) is automatic, and we only have to prove that $\widehat{t^*} \models f_s$. Again, let us begin with a simple observation.

Claim 6.26. *Let α be an arbitrary assignment satisfying a term $t_{\mathcal{B},A}$ of the form (6.11). Then α satisfies all axioms $Con_{\rho(u,v)}(x_u, x_v)$ $(u, v \in \mathcal{B})$ if and only if α is constant on $\mathcal{B}$.*

Proof of 6.26. By an easy inspection. □

Our strategy for proving $\widehat{t^*} \models f_s$ given that $Vars(f_s) \subseteq Vars(\widehat{t})$ (by Claim 6.25) and $t^* \models f_s$ is typical for this kind of arguments in proof complexity. Namely, fix an assignment $\alpha \in V^D$ satisfying $\widehat{t^*}$. In order to prove that $f_s(\alpha) = 1$, we only have to show how to modify α to another assignment β such that:

1. α and β agree on V_0;

2. $t^*(\beta) = 1$.

This β will be obtained from α by the "reverse engineering" of the intuition for the construction of $\widehat{t}$ we provided above. The terms $t_{\mathcal{B}_\gamma, A_\gamma}$ with $\gamma \in \Gamma_0$ will be trivially satisfied by monotonicity, we need not worry about them. If $\gamma \in \Gamma_1$, i.e., $\mathcal{B}_\gamma$ is removed in favor of an adjacent ball $\mathcal{B}^*_\mu$, then, due to the minimality in (6.23), every assignment $a \in A^*_\mu$ can be extended to an assignment $b \in A_\gamma$ satisfying the left-hand side, and we use this b to re-assign variables in the ball $\mathcal{B}_\gamma$.

Formally, consider an individual $\mathcal{B}_\gamma$, $\gamma \in \Gamma^\mu_1$ and let $r \stackrel{\text{def}}{=} r(\mathcal{B}^*_\mu)$. By Claim 6.24, $\mathcal{B}^*_\mu \subseteq Vars(\widehat{t})$, and then by Claim 6.26 (since α satisfies $\eta[\mathcal{B}^*_\mu]$), $\alpha|_{\mathcal{B}^*_\mu}$ is a constant a with $a \in D_h \cup \cdots \cup D_{r+2} \cup A^*_\mu$.

Case 1. $a \in D_h \cup \cdots \cup D_{r+2}$**.**
We simply let $\beta|_{\mathcal{B}_\gamma} \equiv a$.

Case 2.1. $a \in A^*_\mu$ **and** $r(\mathcal{B}_\gamma) = r$**, i.e.,** $\mathcal{B}_\gamma$ **and** $\mathcal{B}^*_\mu$ **are adjacent.**
In the notation of (6.23), there exists $b \in A_\gamma$ such that $\eta[\{X_i, X_j\}](b, a) = 1$; otherwise a could have been removed from A^*_μ in violation of the minimality of (6.23). Pick arbitrarily any such b and define $\beta|_{\mathcal{B}_\gamma} \equiv b$.

Case 2.2. $a \in A^*_\mu,\ r(\mathcal{B}_\gamma) < r$**.**
We let $\beta|_{\mathcal{B}_\gamma} \equiv b$, where $b \in D_{r(\mathcal{B}_\gamma)+1}$ is chosen in such a way that $\eta[\{X_i\}](b) = 1$. Here, as before, i is the last entry in the prefix describing the ball $\mathcal{B}_\gamma$.

The construction of β is complete.

Claim 6.27. *α and β agree on all balls $\mathcal{B}^*_\mu$ and on all balls $\mathcal{B}_\gamma$ ($\gamma \notin \Gamma_1$).*

Proof of 6.27. Follows from the above remarks that the balls $\mathcal{B}_\gamma$ ($\gamma \in \Gamma_1$) are disjoint from anything else. □

In particular, α and β agree on V_0, and it remains to show that $t^*(\beta) = 1$.

First we check that $t(\beta) = 1$, that is $t_{\mathcal{B}_\gamma, A_\gamma}(\beta) = 1$ for any $\gamma \in [w]$; this part is relatively easy and was already sufficiently explained above. There are three possibilities.

Case 1. $\gamma \in \Gamma_0$**.**
We have $\mathcal{B}_\gamma \subset \mathcal{B}^*_\mu$ for some $\mu \notin M_0$, and by Claim 6.27, α and β coincide on $\mathcal{B}^*_\mu$. Since $r(\mathcal{B}_\gamma) \leq r(\mathcal{B}^*_\mu) - 1$, we have $t_{\mathcal{B}^*_\mu, A^*_\mu} \models t_{\mathcal{B}_\gamma, A_\gamma}$ *regardless* of the particular value of A^*_μ (over which we do not have any control). But $t_{\mathcal{B}^*_\mu, A^*_\mu}$ appears in $\widehat{t}$ and hence $t_{\mathcal{B}^*_\mu, A^*_\mu}(\alpha) = 1$. $t_{\mathcal{B}_\gamma, A_\gamma}(\beta) = 1$ follows.

Case 2. $\gamma \in \Gamma_1$**.**
In this case $t_{\mathcal{B}_\gamma, A_\gamma}(\beta) = 1$ directly follows from the way the assignment β was constructed.

Case 3. $\gamma \notin \Gamma_0 \cup \Gamma_1$.
Once again, α and β coincide on $\mathcal{B}_\gamma$, and $t_{\mathcal{B}_\gamma, A_\gamma}$ also appears in $\widehat{t}$. Hence $t_{\mathcal{B}_\gamma, A_\gamma}(\beta) = 1$.

So far we have proved $t(\beta) = 1$, and what still remains is to show that $\eta[Vars(t)](\beta) = 1$, i.e., we should take care of the "environment" of the term t. This part is technically unpleasant, we have tried several natural possibilities (like modifying the predicate $\models$ in the statement of Lemma 6.19 to pertain to legal assignments only), and none of them has been fully satisfactory. Perhaps, the best intuition toward the remaining part of the proof below is that we have defined both the "forward" construction $t \Longrightarrow \widehat{t}$ and the "reverse" one $\alpha \Longrightarrow \beta$ as natural as possible so that the "supplementary" information contained in $\eta[Vars(t)]$ "floats around" both ways. Let us at least briefly explain a (typical) example of the constraint $\mathrm{Con}_d(X_u, X_v)$, where $u \in \mathcal{B}_\gamma,\ v \in \mathcal{B}_{\gamma'}$ and $\gamma \neq \gamma'$. A problem occurs when at least one of the balls $\mathcal{B}_\gamma, \mathcal{B}_{\gamma'}$ is removed, and, moreover, when it is removed due to an adjacency conflict (containment conflicts only increase $Vars(t)$). If, say, $\mathcal{B}_\gamma$ is removed in favor of an adjacent ball $\mathcal{B}^*_\mu \not\supseteq \mathcal{B}_{\gamma'}$, then we will still have that $\rho(\mathcal{B}^*_\mu, \mathcal{B}_{\gamma'}) = \rho(\mathcal{B}_\gamma, \mathcal{B}_{\gamma'}) = d$ and hence our constraint Con_d still makes perfect sense (and the same content, too) between $\mathcal{B}^*_\mu$ and $\mathcal{B}_{\gamma'}$. We replace X_u in it with X_{u^*}, where $u^* \in \mathcal{B}^*_\mu$ is arbitrary, and show that $\mathrm{Con}_d(\alpha_{u^*}, \beta_v) = 1$ implies $\mathrm{Con}_d(\beta_{u^*}, \beta_v) = 1$. Another kind of care should be given to constraints of width 1, for the same reasons as in the proof of Lemma 6.18, but perhaps at this point it will be simpler just to start the formal argument.

Let us fix $C \in \eta$ with $Vars(C) \subseteq Vars(t)$. We need to prove that

$$C(\beta) = 1. \tag{6.28}$$

Case 1. $Vars(C) \subseteq \mathcal{B}_\gamma$ **for some** γ**.**

Let $r \stackrel{\mathrm{def}}{=} r(\mathcal{B}_\gamma)$.

Case 1.1. $C = \widehat{C}_0(X_u, X_v)$ $(u, v \in \mathcal{B}_\gamma)$**, where** $C_0 \in \eta_d$**,** $d \stackrel{\mathrm{def}}{=} \rho(u, v)$**.**

This case is immediate from the already established fact $t(\beta) = 1$ since it implies $\beta_u, \beta_v \in D_h \cup \cdots \cup D_{r+1}$, while $d \leq r$.

Case 1.2. $C = \mathrm{Con}_{\rho(u,v)}(X_u, X_v)$**.**

For every ball $\mathcal{B}$ occurring in the right-hand side of (6.22), $\alpha|_{\mathcal{B}}$ is constant by Claim 6.26 since α satisfies all consistency axioms

$\mathrm{Con}_{\rho(u,v)}(X_u, X_v)$ with $u, v \in \mathcal{B} \subseteq Vars(\widehat{t})$. Following the same reasoning as in the proof of $t(\beta) = 1$ above, $\beta|_{\mathcal{B}_\gamma}$ is also a constant hence $C(\beta) = 1$ by Claim 6.26.

So far we have treated axioms C of width 2 with $Vars(C) \subseteq \mathcal{B}_\gamma$. We divide the analysis of the case when C is of width 1 into two subcases, according to whether $\gamma \in \Gamma_1$ or not.

Case 1.3. $C = \widehat{C}_0(X_u)$**, where** $C_0 \in \eta_d$ **for some** d **and** $\gamma \notin \Gamma_1$**.**
Since $X_u \in Vars(\widehat{t})$, $C(\alpha_u) = 1$, and since $\gamma \notin \Gamma_1$, $C(\beta_u) = C(\alpha_u)$. This gives (6.28).

Case 1.4. $C = \widehat{C}_0(X_u)$**, where, as before,** $C_0 \in \eta_d$ **for some** d **but** $\gamma \in \Gamma_1$**.**
Let $\gamma \in \Gamma_1^\mu$ and $R \stackrel{\text{def}}{=} r(\mathcal{B}_\mu^*) \geq r$. From our construction, either $\alpha|_{\mathcal{B}_\mu^*} \in D_h \cup \cdots \cup D_{R+2}$ and $\beta|_{\mathcal{B}_\gamma} \equiv \alpha|_{\mathcal{B}_\mu^*}$, or $b \in D_{r+1}$ and $\eta[X_i](b) = 1$, where i is again the last entry in the prefix describing $\mathcal{B}_\gamma$.

Case 1.4.1. $\beta|_{\mathcal{B}_\gamma} \equiv \alpha|_{\mathcal{B}_\mu^*} \in D_h \cup \cdots \cup D_{R+2}$ $(= a)$**.**
We may assume $d \geq R+2$ as otherwise the statement is trivial. Pick arbitrarily $u^* \in \mathcal{B}_\mu^*$, then $\rho(u, u^*) = R + 1$. Hence u and u^* share the prefix of length $h - d \leq h - R - 2$, that is $\widehat{C}_0(X_{u^*})$ is also in η. Now $\widehat{C}_0(a) = 1$ follows from $X_{u^*} \in Vars(\widehat{t})$.

Case 1.4.2. $b \in D_{r+1}$ **and** $\eta[i](b) = 1$**.**
Again, this is obvious if $d \neq r+1$ and follows from $C_0 \in \eta_d[\{i\}]$ otherwise.

We have completed the analysis of the case $Vars(C) \subseteq \mathcal{B}_\gamma$ for a single ball $\mathcal{B}_\gamma$. In particular, we can and will now assume that the width of the constraint C is exactly 2:

Case 2. $Vars(C) = \{u, v\}$**,** $u \in \mathcal{B}_\gamma$ **and** $v \in \mathcal{B}_{\gamma'}$ **with** $\gamma \neq \gamma'$**.**
Let $d \stackrel{\text{def}}{=} \rho(u, v)$, $r \stackrel{\text{def}}{=} r(\mathcal{B}_\gamma)$, $r' \stackrel{\text{def}}{=} r(\mathcal{B}_{\gamma'})$, so that $r, r' \leq d - 1$ and, moreover, at least one of this inequalities is strict (since the balls $\mathcal{B}_\gamma$, $\mathcal{B}_{\gamma'}$ are non-adjacent).

Case 2.1. $\gamma \notin \Gamma_1$ **and** $\gamma' \notin \Gamma_1$**.**
This case is immediate: $\beta_u = \alpha_u$, $\beta_v = \alpha_v$ and (6.28) simply follows from the fact that α satisfies $\eta[Vars(\widehat{t})]$.

Case 2.2. $\gamma \in \Gamma_1$**.**

Let $\gamma \in \Gamma_1^\mu$ and $R \stackrel{\text{def}}{=} r(\mathcal{B}_\mu^*)$ so that $R \geq r$. Let also $a \stackrel{\text{def}}{=} \alpha|_{\mathcal{B}_\mu^*}$; $a \in D_h \cup \cdots \cup D_{R+1}$.

The rest of the analysis splits into two rather different cases according to whether $v \in (\mathcal{B}_\mu^*)^+$ or not.

Case 2.2.1. $v \in (\mathcal{B}_\mu^*)^+$**, that is** $d \leq R+1$**.**

Case 2.2.1.1. $a \in D_h \cup \cdots \cup D_{R+2}$**.**

According to the construction, $\beta_u = \beta_v = a$. $\mathrm{Con}_d(\beta_u, \beta_v) = 1$ follows immediately, and for $\widehat{C}_0(\beta_u, \beta_v)$ ($C_0 \in \tau_d$) we only have to remark that $a \notin D_d$ since $d \leq R+1$.

Case 2.2.1.2. $a \in D_{R+1}$**.**

Case 2.2.1.2.1. $v \in \mathcal{B}_\mu^*$**.**

From our construction, $d = R+1$, $\beta_v = a$, and $\beta_u \in D_{r+1}$. Thus, $\beta_u, \beta_v \in D_d \cup \cdots \cup D_1$, and this proves $\mathrm{Con}_d(\beta_u, \beta_v) = 1$, as well as $\widehat{C}_0(\beta_u, \beta_v) = 1$ unless $r = R$, that is the balls $\mathcal{B}_\gamma$ and $\mathcal{B}_\mu^*$ are adjacent. In this latter case $\widehat{C}_0(\beta_u, a) = 1$ is guaranteed by our choice of β_u.

Case 2.2.1.2.2. $v \in (\mathcal{B}_\mu^*)^+ \setminus \mathcal{B}_\mu^*$**.** In this case $\gamma' \in \Gamma_1^\mu$ as well, and, according to our construction, $\beta_u \in D_{r+1}$ and $\beta_v \in D_{r'+1}$. Recalling that $r+1, r'+1 \leq d$ and, moreover, at least one of the inequalities here is strict, both $\mathrm{Con}_d(\beta_u, \beta_v)$ and $\widehat{C}_0(\beta_u, \beta_v)$ ($C_0 \in \tau_d$) are satisfied for trivial reasons.

At this moment, we are done with the case $v \in (\mathcal{B}_\mu^*)^+$.

Case 2.2.2. $v \notin (\mathcal{B}_\mu^*)^+$ **or, in other words,** $d \geq R+2$**.**

Pick arbitrarily $u^* \in \mathcal{B}_\mu^*$. Since $\rho(u, u^*) = R+1$, by the ultrametric triangle inequality we get $\rho(u^*, v) = d$. In particular, $C(X_{u^*}, X_v)$ is also an axiom of η.

Claim 6.29. $C(\beta_u, \beta_v) = C(\alpha_{u^*}, \beta_v)$.

Proof of 6.29. Readily follows from the dichotomy $\beta_u = \beta_{u^*} = \alpha_{u^*}$ or $\beta_u, \alpha_{u^*} \in D_{R+1} \cup \cdots \cup D_1 \subseteq D_{d-1} \cup \cdots \cup D_1$. $\square$

Thus, if $\gamma' \not\in \Gamma_1$ then $\{u^*, v\} \subseteq Vars(\widehat{t})$, $\beta_v = \alpha_v$ and we are done since α satisfies $\widehat{t}^*$. On the other hand, if $\gamma' \in \Gamma_1^{\mu'}$ for some $\mu' \neq \mu$ then $u^* \not\in (\mathcal{B}^*_{\mu'})^+$, and we simply apply Claim 6.29 once more, this time with $u = v$, $v = u^*, u^* = v^*$, where $v^* \in \mathcal{B}^*_{\mu'}$.

This finally completes our case analysis. To re-cap the overall argument, we have proved (6.28) for any axiom $C \in \eta$ with $Vars(C) \subseteq Vars(t)$. That is, for any assignment $\alpha \in D^V$ with $\widehat{t}^*(\alpha) = 1$ we were able to modify it to some $\beta \in V^D$ so that α and β agree on V_0 and $t^*(\beta) = 1$. This implies (6.20) and completes the inductive step. □

As we observed above, the lower bound in Lemma 6.9 follows immediately.

6.4. From multi-valued logic to the Boolean one. Finally, we need to transfer the tradeoff resulting from Lemmas 6.3, 6.9 to the Boolean setting. This involves two different tasks: the conversion per se and a variable compression in the style of Razborov (2016) that is certainly needed here since the number of variables in the lexicographic product is huge (exponential in h). We will combine both tasks into a single statement, but first we need a few definitions.

Definition 6.30 (Alekhnovich & Razborov 2008). *A function $g : \{0,1\}^s \longrightarrow D$ is* r-surjective *if for any restriction ρ assigning at most r variables, the restricted function $g|_\rho$ is surjective.*

Definition 6.31 (cf. Razborov 2016). *Let A be an $m \times n$ 0-1 matrix in which every row has precisely s ones and $g : \{0,1\}^s \longrightarrow D$ be a function. Let $g[A] : \{0,1\}^n \longrightarrow D^m$ be naturally defined as*

$$g[A](x_1, \ldots, x_n)(i) \stackrel{\mathrm{def}}{=} g(x_{j_1}, \ldots, x_{j_s}),$$

where $j_1 < j_2 < \cdots < j_s$ is the enumeration of ones in the ith row of A. For a D-valued Boolean function $f : D^m \longrightarrow \{0,1\}$, we let the Boolean function $f[g, A] : \{0,1\}^n \longrightarrow \{0,1\}$ be the composition $f \circ g[A]$. Finally, for a D-valued CSP $\eta(Y_1, \ldots, Y_m)$, we let $\eta[g, A] \stackrel{\mathrm{def}}{=} \{C[g, A] \mid C \in \eta\}$.

Definition 6.32. *Let A be a $m \times n$ 0-1 matrix. For $i \in [m]$, let*

$$J_i(A) \stackrel{\text{def}}{=} \{\, j \in [n] \mid a_{ij} = 1 \,\}$$

be the set of all ones in the ith row. For a set of rows $I \subseteq [m]$, the boundary $\partial_A(I)$ *of I 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 in their intersections with I. A is an (r, c)-*boundary expander *if $|\partial_A(I)| \geq c|I|$ for every set of rows $I \subseteq [m]$ with $|I| \leq r$.*

Lemma 6.33. *Let A be an $m \times n$ $\left(2h, \frac{3}{4}s\right)$-boundary expander in which every row has precisely s elements. Let D be a finite domain, $\eta(Y_1, \ldots, Y_m)$ be a D-valued h-CSP and $g : \{0,1\}^s \longrightarrow D$ be an $(3s/4)$-surjective function. Assume that there exists a semantic (Boolean) refutation π from $\eta[g, A]$ with $\mathsf{VSpace}(\pi) \leq (hs)/16$. Then there exists a D-valued refutation $\widehat{\pi}$ of η with $\mathsf{VSpace}(\widehat{\pi}) \leq h$ and $|\widehat{\pi}| = |\pi|$.*

Proof. In the notation of this lemma, fix a semantical (Boolean) refutation $\pi = (f_0, \ldots, f_T)$ from $\eta[g, A]$ with $\mathsf{VSpace}(\pi) \leq (hs)/16$. In order to convert π to a D-valued refutation, we need to recall a few rudimentary facts about expanders.

Definition 6.34. *For a set of columns $J \subseteq [n]$, let*

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

be the set of rows completely contained in J. Let $A \backslash J$ be the submatrix of A obtained by removing all columns in J and all rows in $Ker(J)$.

The following is a part of (Razborov 2016, Lemma 4.4).

Proposition 6.35. *Let A be an $(m \times n)$ (r, c)-boundary expander in which every row has at most s ones, let $c' < c$, and let $J \subseteq [n]$ satisfy $|J| \leq \frac{r}{2}(c - c')$. Then there exists $\widehat{J} \supseteq J$ such that $A \backslash \widehat{J}$ is an $(r/2, c')$-boundary expander and $|\widehat{J}| \leq |J| \left(1 + \frac{s}{c - c'}\right)$.*

We now return to the proof of Lemma 6.33. Let[8] $J_t \stackrel{\mathrm{def}}{=} Vars(f_t)$; $|J_t| \leq (hs)/16$. Apply to this set Proposition 6.35 with $r = 2h$, $c = 3s/4$ and $c' = 5s/8$. We will get $\widehat{J_t} \supseteq J_t$ such that $A \setminus \widehat{J_t}$ is an $(h, 5s/8)$-boundary expander and $|\widehat{J_t}| \leq 9|J_t| \leq 9hs/16$. Let $I_t \stackrel{\mathrm{def}}{=} \mathrm{Ker}(\widehat{J_t})$; we claim that $|I_t| \leq h$. Indeed, assuming the contrary, pick a set $I'_t \subseteq I_t$ with $|I'_t| = h$. Then $|\partial_A(I'_t)| \leq |\widehat{J_t}| \leq 9hs/16$, contrary to the fact that A is an $(h, 3s/4)$-boundary expander.

We now let $\widehat{f_t}$ be the *minimal* D-valued function *in the variables* $\{y_i \mid i \in I_t\}$ such that

$$f_t \models \widehat{f_t}[g, A] \tag{6.36}$$

in the Boolean logic. Then $\left|Vars(\widehat{f_t})\right| \leq h$ and all that remains to show is that $(\widehat{f_0}, \widehat{f_1}, \ldots, \widehat{f_T})$ is indeed a D-valued semantic refutation, that is

$$\widehat{f_t} \wedge \eta[I_t \cup I_{t+1}] \models \widehat{f}_{t+1}, \tag{6.37}$$

for all t.

Let $\alpha \in D^m$ be any assignment satisfying the left-hand side in (6.37). Due to the minimality of $\widehat{f_t}$, if we re-define it to 0 on the input $\alpha|_{I_t}$, this will violate (6.36). In other words, there exists a Boolean assignment $a \in \{0,1\}^n$ such that $f_t(a) = 1$ and

$$g\left(a|_{J_i(A)}\right) = \alpha_i, \tag{6.38}$$

for any $i \in I_t$. We note that these two properties of a depend only on those values a_j for which $j \in \widehat{J_t}$; thus, we can view a as an assignment in $\{0,1\}^{\widehat{J_t}}$, discarding all other values. Our goal is to extend a to an assignment in $\{0,1\}^{\widehat{J_t} \cup \widehat{J}_{t+1}}$ in such a way that (6.38) will be satisfied for all $i \in I_{t+1}$ as well.

This is done by a fairly standard argument. Let $I \stackrel{\mathrm{def}}{=} I_{t+1} \setminus I_t$. Since $A \setminus \widehat{J_t}$ is an $(h, 5s/8)$-boundary expander and $|I| \leq |I_{t+1}| \leq h$, we have $\left|\partial_A(I) \setminus \widehat{J_t}\right| \geq \frac{5s}{8}|I|$. Hence for at least one $i \in I$,

$$\left| J_i(A) \setminus \left(\widehat{J_t} \cup \bigcup_{i' \in I \setminus \{i\}} J_{i'}(A) \right) \right| \geq \frac{5s}{8} \geq \frac{s}{4}.$$

[8] Recall that we often identify sets of variables with sets of their indices.

Removing this i from I and arguing by reverse induction, we can order all rows in I in such a way $I = \{i_1, i_2 \ldots, i_r\}$ that

$$\left| J_{i_\nu}(A) \setminus \left(\widehat{J}_t \cup J_{i_1}(A) \cup \cdots \cup J_{i_{\nu-1}}(A) \right) \right| \geq \frac{s}{4} \tag{6.39}$$

for all $\nu = 1..r$. Using now $(3s/4)$-surjectivity of g, we consecutively extend a to $\widehat{J}_t \cup J_{i_1}(A) \cup \cdots \cup J_{i_\nu}(A)$ enforcing all conditions (6.38).

The partial assignment $a \in \{0,1\}^{\widehat{J}_t \cup \widehat{J}_{t+1}}$ we have constructed still satisfies f_t; we claim that it also satisfies $\eta[g, A] \left[\widehat{J}_t \cup \widehat{J}_{t+1} \right] \supseteq \eta[g, A][J_t \cup J_{t+1}]$.

Indeed, for any $C \in \eta[I_t \cup I_{t+1}]$ this simply follows from the fact that $C(\alpha) = 1$ and the consistency conditions (6.38). One thing we still have to make sure is that $\eta[g, A] \left[\widehat{J}_t \cup \widehat{J}_{t+1} \right]$ does not contain any other, "accidental" constraints.

Claim 6.40. *If C is any constraint of width $\leq h$ and $Vars(C[g, A]) \subseteq \widehat{J}_t \cup \widehat{J}_{t+1}$ then $Vars(C) \subseteq I_t \cup I_{t+1}$.*

Proof of 6.40. By a relatively simple modification of the argument above. Let $I \stackrel{\text{def}}{=} Vars(C)$; $|I| \leq h$, and assume the contrary, that is that there exists $i \in I \setminus (I_t \cup I_{t+1})$. Fix two assignments $\alpha, \beta \in D^I$ differing only in the ith coordinate but such that $C(\alpha) \neq C(\beta)$. We claim that there exist $a, b \in \{0,1\}^n$ such that (cf. (6.38))

$$g\left(a|_{J_i(A)}\right) = \alpha_i, \; g\left(b|_{J_i(A)}\right) = \beta_i, \tag{6.41}$$

for all $i \in I$ while $a|_{\widehat{J}_t \cup \widehat{J}_{t+1}} = b|_{\widehat{J}_t \cup \widehat{J}_{t+1}}$: the first property will imply $C[g, A](a) \neq C[g, A](b)$, and that will contradict $Vars(C[g, A]) \subseteq \widehat{J}_t \cup \widehat{J}_{t+1}$ by the second property.

We construct the promised a, b in two stages. Let $I' \stackrel{\text{def}}{=} I \cap (I_t \cup I_{t+1})$; thus, α and β agree on I'. As before, order the rows in I' in such a way $I' = \{I_1, \ldots, I_r\}$ that

$$\left| J_{i_\nu}(A) \setminus (J_{i_1}(A) \cup \cdots \cup J_{i_{\nu-1}}(A)) \right| \geq s/4$$

holds for all ν (cf. (6.39)), and then satisfy (6.41) with the *same* assignment to $J_{i_1}(A)\cup\cdots\cup J_{i_r}(A)$. Extend it to $\widehat{J}_t\cup\widehat{J}_{t+1}$ arbitrarily; let $c\in\{0,1\}^{\widehat{J}_t\cup\widehat{J}_{t+1}}$ be the resulting assignment.

Now, let $I'' \stackrel{\text{def}}{=} I\setminus(I_t\cup I_{t+1})$, and let A^* be the matrix obtained from A by removing *all* columns J_t, J_{t+1} and all rows I_t, I_{t+1}. Since both $A\setminus J_t$ and $A\setminus J_{t+1}$ are $(h/2, 5s/8)$-expanders, clearly A^* is still an $(h, s/4)$-expander ($\frac{s}{4} = 2\cdot\left(\frac{5}{8}s\right) - s$). This expansion property allows us to extend c, by the same argument as above, to $a, b\in\{0,1\}^n$ that will satisfy (6.41) for $i\in I''$ as well. But, as we remarked above, (6.41) is in contradiction with $Vars(C[g,A])\subseteq \widehat{J}_t\cup\widehat{J}_{t+1}$. $\square$

Since η is an h-CSP by the assumption of Lemma 6.33, we can apply Claim 6.40 to any $C\in\eta$. This gives us that all constraints in $\eta[g,A]\left[\widehat{J}_t\cup\widehat{J}_{t+1}\right]$ are indeed coming from $\eta[I_t\cup I_{t+1}]$. As we already remarked above, this implies that all of them are satisfied by the assignment a, and since $(f_0, f_1,\ldots,f_T)$ is a semantical refutation from $\tau[g,A]$, we conclude that $f_{t+1}(a) = 1$. By (6.36), $\widehat{f}_{t+1}[g,A] = 1$, and since a satisfies (6.38) for all $i\in I_{t+1}$, this means $\widehat{f}_{t+1}(\alpha) = 1$.

We have established (6.37) by showing that any D-valued assignment α satisfying its left-hand side also satisfies the right-hand side. Thus, $(\widehat{f}_0, \widehat{f}_1,\ldots,\widehat{f}_T)$ is indeed a semantical refutation. This completes the proof of Lemma 6.33. $\square$

6.5. Putting it all together.

Proof of 3.3. We are given a function $s = s(n)\le n^{1/2}$. We let $h \stackrel{\text{def}}{=} \lfloor \epsilon s\rfloor \le \epsilon n/s$, where ϵ is a sufficiently small constant. We now need the following standard fact:

Proposition 6.42 (Razborov 2016, Lemma 2.2). *Let $n\to\infty$ and m, s, c be arbitrary integer parameters possibly depending on n such that $c\le \frac{3}{4}s$ and*

$$r \le o(n/s)\cdot m^{-\frac{2}{s-c}}. \qquad (6.43)$$

Then for sufficiently large n there exist $m \times n$ (r, c)-boundary expanders in which every row has $\leq s$ ones.

Note that (6.43) is satisfied with $m := 4^h$, $c := \frac{3}{4}s$ and $r := 2h$ (the term $m^{-\frac{2}{s-c}}$ becomes $\Omega(1)$). This gives us an $(2h, \frac{3}{4}s)$-boundary expander with $m = 4^h$ rows in which every row has $\leq s$ ones. By adding $\leq s$ extra columns if necessary, we can assume w.l.o.g. that every row has exactly s ones.

Next, we claim that if s is sufficiently large (that we can clearly assume w.l.o.g.) a random function $\boldsymbol{g} : \{0,1\}^s \longrightarrow D$ is $(3s/4)$ surjective for $|D| = 2^{s/8}$, with probability $1 - o(1)$. This is straightforward:

$$\begin{aligned}
&\mathbf{P}[\boldsymbol{g} \text{ is not } (3s/4) - \text{surjective}] \\
&\quad \leq 2^{3s/4}\binom{s}{3s/4}\mathbf{P}[\boldsymbol{g}|_\rho \text{ is not surjective}] \\
&\quad \leq \exp(O(s)) \cdot (1 - |D|^{-1})^{2^{s/4}} \leq o(1),
\end{aligned}$$

where ρ is a fixed restriction assigning $(3s/4)$ variables. Pick any such g arbitrarily, and split D into h nearly equal parts, $D = D_1 \dot\cup \cdots \dot\cup D_h$. Let η_d be the D_d-valued 2-CSP in four variables given by Lemma 6.3, that is such that $\mathsf{VSpace}(\eta_d \vdash 0) = 1$ but any its refutation π with $\mathsf{VSpace}(\pi) = 1$ must have length $\geq \exp\left(|D_d|^{\Omega(1)}\right) \geq \exp(\exp(\Omega(s)))$. Let $\eta \stackrel{\text{def}}{=} \eta_h \cdot \ldots \cdot \eta_1$. The D-valued 2-CSP η has $m = 4^h$ variables, say, $Y_1, \ldots, Y_m$ and still satisfies $\mathsf{VSpace}(\eta \vdash 0) = 1$ but now any its refutation $\widehat{\pi}$ with $\mathsf{VSpace}(\widehat{\pi}) \leq h/2 - 1$ has length $\exp(\exp(\Omega(s)))$. The desired contradiction τ_n will be $\eta[g, A]$.

First of all, τ_n has a semantical refutation with variable space $\leq s$. It is obtained simply by taking a D-valued refutation of η with variable space 1 (that is, consisting of generalized literals) and applying the operator $Y_i^P \mapsto Y_i^P[g, A]$ to its configurations. On the other hand, applying Lemma 6.33 in the contrapositive form, every Boolean refutation π from $\eta_g[A]$ with $\mathsf{VSpace}(\pi) \leq (h/2 - 1)s/16$ must have length $\geq \exp(\exp(\Omega(s)))/\exp(O(h))$. As $h = \Theta(s)$, this is $\exp(\exp(\Omega(s)))$. $\square$

7. Conclusion

In this paper we have studied two complexity measures of propositional proofs, variables space and depth, that in our view have been somewhat neglected in the past. We hope that perhaps the nature of the results proved in this paper would help them to find the place in the overall hierarchy that, in our opinion, they fully deserve by the token of being very clean, robust and natural.

That said, the most interesting question about them remains open: whether variable space and depth are polynomially related or, equivalently, whether there exists a supercritical tradeoff between them. In a slightly less precise form this was asked in (Urquhart 2011, Problem 7.2); we have proved a quadratic gap, but the general problem looks quite challenging.

A positive answer to this question would immediately imply that clause space is polynomially bounded by variable space. Even if these two problems seem to be extremely tightly related, we still would like to ask this separately: is it correct that

$$\mathsf{CSpace}(\tau_n \vdash 0) \leq (\mathsf{VSpace}(\tau_n \vdash 0) \log n)^{O(1)}?$$

In the opposite range, of (barely) constant variable space, all refutations a priori have small length, and we have shown that the depth can be reduced to, say, n while keeping the variable space constant and length polynomial. We would like to take this opportunity and re-iterate an interesting question of (somewhat) similar flavor asked in (Nordström 2013, Open Problem 16). Assume that we have a configurational refutation of constant *clause* space. Is it always possible to reduce *length* to polynomial while keeping the clause space constant? As with our first question, this one also looks quite challenging.

Finally, there still remains a considerable amount of work to be done on refining simulations in Theorem 3.1. For example, let us take a closer look at (1.6). By Bonacina's result (1.5), every $O(1)$-CNF τ_n with $w(\tau_n \vdash 0) = \Theta(n)$ automatically provides an example with $\mathsf{TSpace}(\tau_n \vdash 0) = \Theta(D^2)$ $(= \Theta(n^2))$. But what about the lower bound in (1.6)? Can, say, TSpace be sub-linear in depth or the bound can be improved to $\widetilde{\Omega}(D)$? This does not seem to easily follow from any known results.

Acknowledgements

This work was supported by the Russian Science Foundation, grant # 14-50-00005.

The author is deeply indebted to anonymous referees of this paper for constructive criticism and a great deal of very useful remarks.

References

M. Alekhnovich, E. Ben-Sasson, A. Razborov & A. Wigderson (2002). Space complexity in propositional calculus. *SIAM Journal on Computing* **31**(4), 1184–1211.

M. Alekhnovich & A. Razborov (2008). Resolution is not automatizable unless $W[P]$ is tractable. *SIAM Journal on Computing* **38**(4), 1347–1363.

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

L. Babai & A. Seress (1992). On the Diameter of Permutation Groups. *European Journal of Combinatorics* **13**, 231–243.

B. Barak & D. Steurer (2014). Sum-of-squares proofs and the quest toward optimal algorithms. In *Proceedings of International Congress of Mathematicians (ICM)*, volume IV, 509–533.

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

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

E. Ben-Sasson & A. Wigderson (2001). Short Proofs are Narrow - Resolution made Simple. *Journal of the ACM* **48**(2), 149–169.

C. Berkholz & J. Nordström (2016a). Near-Optimal Lower Bounds on Quantifier Depth and Weisfeiler-Leman Refinement Steps. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*, 267–276.

C. Berkholz & J. Nordström (2016b). Supercritical Space-Width Trade-offs for Resolution. In *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming*, 57:1–57:14.

A. Blake (1937). *Canonical expressions in Boolean algebra*. Ph.D. thesis, University of Chicago.

I. Bonacina (2016). Total space in Resolution is at least width squared. In *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming*, 56:1–56:13.

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

J. L. Esteban & J. Torán (2001). Space bounds for resolution. *Information and Computation* **171**(1), 84–97.

Y. Filmus, M. Lauria, M. Mikša, J. Nordström & M. Vinyals (2015). From Small Space to Small Width in Resolution. *ACM Transactions on Computational Logic* **16**(4), article 28.

D. Grigoriev (2001). Linear lower bounds on degrees of Postivestellensatz calculus proofs for the parity. *Theoretical Computer Science* **259**, 613–622.

J. Nordström (2013). Pebble Games, Proof Complexity and Time-Space Trade-Offs. *Logical Methods in Computer Science* **9**, 1–63.

A. Razborov (2016). A New Kind of Tradeoffs in Propositional Proof Complexity. *Journal of the ACM* **62**(3), article 16.

A. Razborov (2017). On the Width of Semi-Algebraic Proofs and Algorithms. *Mathematics of Operation Research* http://pubsonline.informs.org/doi/abs/10.1287/moor.2016.0840.

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

A. Urquhart (2011). The Depth of Resolution Proofs. *Studia Logica* **99**, 349–364.

Manuscript received 7 December 2016

Alexander Razborov
University of Chicago
Chicago IL 60637, USA
razborov@math.uchicago.edu

Steklov Mathematical Institute
Moscow 117418 Russia
razborov@mi.ras.ru.