Short Proofs May Be Spacious: An Optimal Separation of Space and Length in Resolution

Eli Ben-Sasson*
Computer Science Department
Technion — Israel Institute of Technology
Haifa, 32000, Israel
eli@cs.technion.ac.il

Jakob Nordström†
School of Computer Science and Communication
Royal Institute of Technology (KTH)‡
SE-100 44 Stockholm, Sweden
jakobn@kth.se

Abstract

A number of works have looked at the relationship between length and space of resolution proofs. A notorious question has been whether the existence of a short proof implies the existence of a proof that can be verified using limited space.

In this paper we resolve the question by answering it negatively in the strongest possible way. We show that there are families of 6*-CNF formulas of size* n*, for arbitrarily large* n*, that have resolution proofs of length* $\mathrm{O}(n)$ *but for which any proof requires space* $\Omega(n/\log n)$*. This is the strongest asymptotic separation possible since any proof of length* $\mathrm{O}(n)$ *can always be transformed into a proof in space* $\mathrm{O}(n/\log n)$*.*

Our result follows by reducing the space complexity of so called pebbling formulas over a directed acyclic graph to the black-white pebbling price of the graph. The proof is somewhat simpler than previous results (in particular, those reported in [Nordström 2006, Nordström and Håstad 2008]) as it uses a slightly different flavor of pebbling formulas which allows for a rather straightforward reduction of proof space to standard black-white pebbling price.

1. Introduction

Resolution length and space Perhaps the single most studied proof system in propositional proof complexity is *resolution*. This system made its first appearance in 1937 in [9] and began to be investigated in connection with automated theorem proving in the 1960s [13, 14, 29]. Because of the simplicity of resolution—there is only one derivation rule—and because all lines in a proof are clauses, this proof system readily lends itself to proof search algorithms.

Being so simple and fundamental, resolution was also a natural target to attack when developing methods for proving lower bounds in proof complexity. In this context, it is most straightforward to prove bounds on the *length* of refutations, i.e., the number of clauses, rather than on the total size of refutations. The length and size measures are easily seen to be polynomially related. In 1968, Tseitin [36] presented a superpolynomial lower bound on refutation length for a restricted form of resolution, called *regular* resolution, but it was not until almost 20 years later that Haken [21] proved the first superpolynomial lower bound for general resolution. This weakly exponential bound of Haken has later been followed by many other strong results, among others truly exponential lower bounds on resolution refutation length for different formula families in, for instance, [4, 8, 11, 37].

The formal study of *space* in resolution was initiated by Esteban and Torán [16, 34]. Intuitively, the space of a resolution refutation is the maximal number of clauses one needs to keep in memory while verifying the refutation, and the space of refuting the CNF formula F is defined as the minimal space of any resolution refutation of F. A number of upper and lower bounds for refutation space in resolution and other proof systems have subsequently been presented in, for example, [1, 7, 15, 17].

With the definition of space complexity, a natural question to ask is how space relates to other complexity measures of propositional proofs. Esteban and Torán [16] proved that the space is at most logarithmic in the minimal length of a treelike refutation of a formula, which implies that space is bounded by the number of variables appearing in the formula. The question of the relation between space

*Research supported in part by an Alon Fellowship and grants by the Israeli Science Foundation and by the US-Israel Binational Science Foundation.

†Research supported in part by the Ericsson Research Foundation, the Foundation Olle Engkvist Byggmästare, and the Foundation Blanceflor Boncompagni-Ludovisi, née Bildt.

‡Part of this work performed while visiting the Technion.

and length of general resolution proofs, which is the focus of this paper, was raised by the first author in [6] and has also been discussed in, for instance, [17, 33, 35], but there has been no consensus on what the right answer should be. However, these papers identify a plausible formula family for answering the question, namely so-called *pebbling contradictions* defined in terms of pebble games over directed acyclic graphs (DAGs) and these formulas have indeed been used in [25, 27] to make progress and, in this paper, finally resolve the question.

While understanding the relation between space and length seemed stuck, progress was reported on another front — that of space versus *width*. The width measure, first made explicit by Galil in [19], is defined as the maximal number of literals in a clause in the refutation. Atserias and Dalmau showed in [3] that space is always greater than width, raising the possibility of equivalence of these two measures. Notice that width is a different measure of "proof space" as it is the maximal "space" occupied by a single line in the refutation and one may have speculated that the two "space" measures are in fact equivalent.

Progress on the space-length question for general resolution was finally obtained by the second author in [25], which also separated space from width. This was done by exhibiting a k-CNF formula family of size $\mathrm{O}(n)$ refutable in width $\mathrm{O}(1)$ and length $\mathrm{O}(n)$ but requiring space $\Theta(\log n)$. In a recent joint work of the second author with Håstad [27] this separation was improved to width $\mathrm{O}(1)$ and length $\mathrm{O}(n)$ versus space $\Theta(\sqrt{n})$ for a related formula family. We note however that this previous state-of-the-art did not rule out the existence of a space-length tradeoff quantitatively similar to the width-length tradeoff of [8] which says width is at most $\mathrm{O}(\sqrt{n \cdot \text{ length }})$.

Our contribution In this paper, we finally resolve the open question about the relationship between space and length by establishing an optimal separation between the two measures. We do this by studying a somewhat modified variant of pebbling contradictions defined using XORs (see Definition 2.4) and proving lower bounds for such *XOR-pebbling contradictions* in terms of the pebbling price of the underlying DAGs.

Theorem 1.1 (Main). *The space of refuting XOR-pebbling contradictions over any DAG G in resolution is lower-bounded by the black-white pebbling price of G, provided that the number of variables per vertex in the XOR-pebbling contradictions is at least* 2*.*

If we take a constant number of variables per vertex and study DAGs with constant fan-in, it is easy to show that XOR-pebbling contradictions can be refuted in linear length and constant width. Using the result from [20] which exhibits a family of fan-in 2 DAGs $\{G_n\}_{n=1}^{\infty}$ of size $\mathrm{O}(n)$ having pebbling price $\Omega(n/\log n)$, we get the following corollary.

Corollary 1.2 (Main). *There is a family $\{F_n\}_{n=1}^{\infty}$ of* 6*-CNF formulas of size $\mathrm{O}(n)$ that can be refuted in length $\mathrm{O}(n)$ and width $\mathrm{O}(1)$ but require space $\Omega(n/\log n)$.*

Since it can be proven using results from [16, 22] that a refutation of length $\mathrm{O}(n)$ can be carried out in space $\mathrm{O}(n/\log n)$, the separation of space and length in Corollary 1.2 is asymptotically optimal. As an extra bonus, we note that while the constructions in [25, 27] are quite intricate and the proofs very involved, our optimal lower bound proof is relatively clean and straightforward and we discuss it next.

Proof outline For the purposes of analyzing space, a resolution derivation from a CNF formula F can be viewed as a sequence of derivation steps on a blackboard. In each step we may write a clause from F on the blackboard (an *axiom* clause), erase a clause from the blackboard or derive some new clause implied by the clauses currently written on the blackboard. The space of a derivation is then the maximum number of clauses on the blackboard simultaneously.

The black-white pebble game models non-deterministic computation, and the black-white pebbling price of a DAG G is the minimal number of memory registers needed to verify the calculation described by G, where the source vertices contain the input and non-source vertices specify operations on the values of the predecessors. The pebble game on a DAG G can be encoded as an unsatisfiable CNF formula, a so-called *pebbling contradiction* over G.

Pebble games have been used extensively as a tool to prove time and space lower bounds and tradeoffs for computation. Loosely put, a lower bound for the pebbling price of a graph says that although the computation that the graph describes can be performed quickly, it requires large space. Our hope is that when we encode pebble games in terms of CNF formulas, these formulas should inherit the same properties as the underlying graphs. That is, if we pick a DAG G with high pebbling price, since the corresponding pebbling contradiction encodes a calculation which needs a lot of memory we would like to try to argue that any resolution refutation of this formula should require large space.

Ideally, we would like to give a proof of a lower bound on the resolution refutation space of pebbling contradictions along the following lines:

1. First, find a natural interpretation of sets of clauses currently "on the blackboard" in a refutation of the pebbling contradiction over G in terms of black and white pebbles on the vertices of the DAG G.

2. Then, prove that this interpretation captures the pebble game in the following sense: for any resolution

refutation of a pebbling contradiction over G, looking at consecutive sets of clauses on the blackboard and considering the corresponding sets of pebbles we get a black-white pebbling of G.

3. Finally, show that the interpretation captures clause space in the sense that if the content of the blackboard induces N pebbles on the graph, then there must be at least N clauses on the blackboard.

Combining the above with known lower bounds on the pebbling price of G, this would imply a lower bound on the refutation space of pebbling contradictions. The separation from length and width would then follow since pebbling contradictions are known to be refutable in linear length and constant width.

Unfortunately, this idea does not quite work "off the shelf." Pebblings of DAGs and resolution refutations of CNF formulas are very different objects, and there is no reason a priori that there should be a tight connection between the two. However, relaxing the requirements for the correspondence between resolution and pebbling, the papers [25, 27] made essentially the proof idea above work for two special cases of graphs. In this paper, by using related ideas and studying a slightly modified variant of pebbling contradictions, we can handle any graph, which results in an optimal separation of space and length.

Implications for practical SAT-solvers In recent years, SATISFIABILITY has gone from a problem of mainly theoretical interest to a practical approach for solving applied problems. Although all known Boolean satisfiability solvers (SAT-solvers) have exponential running time in the worst case, enormous progress in performance has led to satisfiability algorithms becoming a standard tool for solving a large number of real-world problems such as hardware and software verification, experiment design, and scheduling.

Perhaps a somewhat surprising aspect of this development is that the most successful SAT-solvers to date are still variants of the resolution-based Davis-Putnam-Logemann-Loveland (DPLL) procedure [13, 14] augmented with *clause learning*. For instance, the great majority of the best algorithms at the 2007 round of the international SAT competitions [32] fit this description. DPLL procedures perform a recursive backtrack search in the space of partial truth value assignments. The idea behind clause learning, or *conflict-driven learning*, is that at each failure (backtrack) point in the search tree, the system derives a reason for the inconsistency in the form of a new clause and then adds this clause to the original CNF formula ("learning" the clause). This can save a lot of work later on in the proof search, when some other partial truth value assignment fails for similar reasons. The main bottleneck for this approach, other than the obvious one of time, is the amount of memory used by the algorithms. Thus, understanding time and memory requirements for clause learning algorithms, and how these requirements are related to one another is a question of great practical importance. We refer to, e.g., [5, 23, 30] for a more detailed discussion of clause learning (and SAT-solving in general) with examples of applications.

In the field of proof complexity, the resources of time and memory correspond to the length and space of resolution proofs. Our work indicates that on certain input formulas, a short proof does not necessarily imply a space-efficient proof exists. Let us give one implication of our result to questions regarding the practical construction of DPLL-based SAT-solvers.

Consider a "frugal" DPLL-based solver augmented with clause learning that tries to save memory by limiting the number of learned clauses as a function of its running time. The reasoning underlying the frugal algorithm is very natural — to save running time, start with the very minimal possible resources and increase them slowly as necessary. Appealing as this strategy may seem, our work shows that on certain inputs it will perform much worse than other, more prodigal, strategies.[1]

Organization of the rest of the paper In Section 2, we state our results formally. Section 3 defines the "resolution-pebbling game" that we use as an intermediate step when translating resolution refutations into black-white pebblings. In Sections 4–6 we provide the proof of our main theorem. Section 7 contains some short concluding remarks.

2. Definitions and Main Results

For the sake of completeness, before presenting our main results we briefly recount (verbatim) from [26] a few basic definitions regarding resolution, pebble games and pebbling contradictions that will be used later on.

Resolution Following the exposition in [16], a resolution proof can be seen as a Turing machine computation, with a special read-only input tape from which the axioms can be downloaded and a working memory where all derivation steps are made. Then the *space* of a resolution proof is the maximum number of clauses that need to be kept in memory simultaneously during a verification of the proof. The formal definitions follow.

[1] This issue is somewhat subtle, however, and out of space considerations we cannot give a full discussion here. Let us just note that there are empirical results like [31] indicating that although pebbling contradictions have very short resolution proofs, these proofs can be very hard to find even for a state-of-the-art SAT-solver.

Definition 2.1 (Resolution ([1])). A *clause configuration* $\mathbb{C}$ is a set of clauses. A sequence of clause configurations $\{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\}$ is a *resolution derivation* from a CNF formula F if $\mathbb{C}_0 = \emptyset$ and for all $t \in [\tau]$, $\mathbb{C}_t$ is obtained from $\mathbb{C}_{t-1}$ by one of the following rules:

Axiom Download $\mathbb{C}_t = \mathbb{C}_{t-1} \cup \{C\}$ for some $C \in F$.

Erasure $\mathbb{C}_t = \mathbb{C}_{t-1} \setminus \{C\}$ for some $C \in \mathbb{C}_{t-1}$.

Inference $\mathbb{C}_t = \mathbb{C}_{t-1} \cup \{C_1 \vee C_2\}$ where $C_1 \vee C_2$ is derived by resolution from the two clauses $C_1 \vee x, C_2 \vee \overline{x} \in \mathbb{C}_{t-1}$.

A resolution derivation $\pi : F \vdash A$ of a clause A from a formula F is a derivation $\{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\}$ such that $\mathbb{C}_\tau = \{A\}$. A *resolution refutation* of F is a derivation of the empty clause 0 from F.

Definition 2.2 (Length, width, space). The *width* $W(C)$ of a clause C is $|C|$, i.e., the number of literals in it. The width of a clause configuration $\mathbb{C}$ is $W(\mathbb{C}) = \max_{C \in \mathbb{C}}\{W(C)\}$. The *space* of a configuration $\mathbb{C}$ is $Sp(\mathbb{C}) = |\mathbb{C}|$, i.e., the number of clauses in $\mathbb{C}$.

Let π be a resolution derivation. Then:

- The *length* $L(\pi)$ of π is the number of axiom download and inference steps in π.
- The *width* of π is $W(\pi) = \max_{\mathbb{C} \in \pi}\{W(\mathbb{C})\}$.
- The *space* of π is $Sp(\pi) = \max_{\mathbb{C} \in \pi}\{Sp(\mathbb{C})\}$.

We define the length of deriving a clause A from F as $L(F \vdash A) = \min_{\pi : F \vdash A}\{L(\pi)\}$, where the minimum is taken over all resolution derivations of A. The width $W(F \vdash A)$ and space $Sp(F \vdash A)$ of deriving A from F are defined completely analogously. The length, width, and space of refuting F is $L(F \vdash 0)$, $W(F \vdash 0)$, and $Sp(F \vdash 0)$, respectively, where as before 0 denotes the contradictory empty clause.

Pebble Games The black pebbling price of a DAG G captures the memory space, i.e., the number of registers, required to perform the deterministic computation described by G. The space of a non-deterministic computation is measured by the black-white pebbling price of G. We say that vertices of G with indegree 0 are *sources* and that vertices with outdegree 0 are *sinks* or *targets*. In the following, unless otherwise stated we will assume that all DAGs under discussion have a unique sink, and this sink will always be denoted z. The next definition is adapted from [12], though we use the established pebbling terminology introduced by [22].

Definition 2.3 (Pebble game). Suppose that G is a DAG with sources S and a unique target z. The *black-white pebble game* on G is the following one-player game. At any point in the game, there are black and white pebbles placed on some vertices of G, at most one pebble per vertex. A *pebble configuration* is a pair of subsets $\mathbb{P} = (B, W)$ of $V(G)$, comprising the black-pebbled vertices B and white-pebbled vertices W. The rules of the game are as follows:

1. If all immediate predecessors of an empty vertex v have pebbles on them, a black pebble may be placed on v. In particular, a black pebble can always be placed on any vertex in S.
2. A black pebble may be removed from any vertex at any time.
3. A white pebble may be placed on any empty vertex at any time.
4. If all immediate predecessors of a white-pebbled vertex v have pebbles on them, the white pebble on v may be removed. In particular, a white pebble can always be removed from a source vertex.

A *black-white pebbling* from (B_1, W_1) to (B_2, W_2) in G is a sequence of pebble configurations $\mathcal{P} = \{\mathbb{P}_0, \ldots, \mathbb{P}_\tau\}$ such that $\mathbb{P}_0 = (B_1, W_1)$, $\mathbb{P}_\tau = (B_2, W_2)$, and for all $t \in [\tau]$, $\mathbb{P}_t$ follows from $\mathbb{P}_{t-1}$ by one of the rules above. If $(B_1, W_1) = (\emptyset, \emptyset)$, we say that the pebbling is *unconditional*, otherwise it is *conditional*.

The *cost* of a pebble configuration $\mathbb{P} = (B, W)$ is $\mathit{cost}(\mathbb{P}) = |B \cup W|$ and the cost of a pebbling $\mathcal{P} = \{\mathbb{P}_0, \ldots, \mathbb{P}_\tau\}$ is $\max_{0 \le t \le \tau}\{\mathit{cost}(\mathbb{P}_t)\}$. The *black-white pebbling price* of (B, W), denoted $\mathit{BW\text{-}Peb}(B, W)$, is the minimum cost of any unconditional pebbling reaching (B, W).

A *complete pebbling* of G, also called a *pebbling strategy* for G, is an unconditional pebbling reaching $(\{z\}, \emptyset)$. The *black-white pebbling price* of G, denoted $\mathit{BW\text{-}Peb}(G)$, is the minimum cost of any complete black-white pebbling of G.

Pebbling Formulas Let $\oplus_{i=1}^{d} x_i$ denote the xor of $x_1, \ldots, x_d$ and $\overline{\oplus}_{i=1}^{d} x_i$ denote the negation of this formula. The satisfying assignments of $\oplus_{i=1}^{d} x_i$ ($\overline{\oplus}_{i=1}^{d} x_i$, respectively) are assignments with an odd (even, respectively) number of 1's. In what follows, we associate a Boolean formula with the CNF formula that is logically equivalent to it in the canonical way. For instance, the formula $(x \oplus y) \rightarrow (z \oplus w)$, which is equivalent to $(x \overline{\oplus} y) \vee (z \oplus w)$, is associated with the CNF formula

$$(\overline{x} \vee y \vee z \vee w) \wedge (\overline{x} \vee y \vee \overline{z} \vee \overline{w}) \wedge (x \vee \overline{y} \vee z \vee w) \wedge (x \vee \overline{y} \vee \overline{z} \vee \overline{w}) \ .$$

The next definition is a generalization of formulas previously studied in [8, 10, 28].

Definition 2.4 (XOR-pebbling contradiction). Let G be a DAG with sources S, a unique sink z, and let $d > 0$ be an integer. Associate d distinct variables $v_1, \ldots, v_d$ with every vertex $v \in V(G)$. The dth degree *XOR-pebbling contradiction* over G, denoted $Peb_G^d[\oplus]$, is the CNF obtained from the conjunction of the following formulas over xor-constraints:

- **Source Axioms:** $\bigoplus_{i=1}^d s_i$ for all sources $s \in S$.
- **Pebbling Axioms:** For all vertices $u^{(1)}, \ldots, u^{(\ell)}, v$, such that $u^{(1)}, \ldots, u^{(\ell)}$ are all the immediate predecessors of v, we have $\bigoplus_{i=1}^d u_i^{(1)} \wedge \ldots \wedge \bigoplus_{i=1}^d u_i^{(\ell)} \rightarrow \bigoplus_{i=1}^d v_i$, which is equivalent to the disjunction
$$\overline{\bigoplus}_{i=1}^d u_i^{(1)} \vee \ldots \vee \overline{\bigoplus}_{i=1}^d u_i^{(\ell)} \vee \bigoplus_{i=1}^d v_i .$$
- **Sink Axioms:** $\overline{\bigoplus}_{i=1}^d z_i$ for the sink z.

If G has n vertices and maximal in-degree ℓ, then $Peb_G^d[\oplus]$ is an unsatisfiable $(\ell+1)d$-CNF formula with at most $2^{(\ell+1)(d-1)} \cdot n$ clauses over $d \cdot n$ variables.

We can now give a more precise statement of our lower bound on refutation space for XOR-pebbling contradictions.

Theorem 1.1 (restated). *For every $d > 1$, there is a constant c such that for any DAG G it holds that*
$$Sp(Peb_G^d[\oplus] \vdash 0) \geq \textsf{\textit{BW-Peb}}(G) - c \ .$$

In what follows, a family of formulas $\{F_n\}_{n=1}^\infty$ is said to be *explicitly constructible* if there exists a polynomial time Turing machine that on input 1^n outputs F_n.

Corollary 1.2 (restated). *For every $d > 1$, there is a explicitly constructible family $\{F_n\}_{n=1}^\infty$ of $3d$-CNF formulas of size $\mathrm{O}(n)$ such that $L(F_n \vdash 0) = \mathrm{O}(n)$ and $W(F_n \vdash 0) = \mathrm{O}(1)$ but $Sp(F_n \vdash 0) = \Omega(n/\log n)$.*

Proof. For any DAG G with n vertices, in-degree 2 and a single sink, the CNF formula $Peb_G^2[\oplus]$ is a $3d$-CNF of size $\mathrm{O}(n)$ that can be refuted using proofs of length $\mathrm{O}(n)$ and width $\mathrm{O}(1)$ (for a proof see [6, Theorem 4.3]). The lower bound on space and the explicit constructibility of the formulas follow respectively from Theorem 1.1 and the following lower bound on black-white pebbling price. □

Theorem 2.5 ([20]). *There is a family of explicitly constructible*[2] *DAGs G_n with $\Theta(n)$ vertices and vertex in-degree 2 for all non-sources such that $\textsf{\textit{BW-Peb}}(G) = \Theta(n/\log n)$.*

Proof of Theorem 1.1. There are three main components to our proof of Theorem 1.1. In the next section we define and discuss the *resolution-pebbling price* of a DAG G, denoted $\textsf{\textit{Res-Peb}}(G)$. Then we prove the following pair of statements. The first theorem is proved in Sections 4,5 and the second is proved in Section 6. Taken together, they complete the proof of Theorem 1.1. □

Theorem 2.6. *For every $d > 1$, there is a constant c such that for any DAG G it holds that*
$$Sp(Peb_G^d[\oplus] \vdash 0) \geq \textsf{\textit{Res-Peb}}(G) - c \ .$$

Theorem 2.7. *For any DAG G it holds that*
$$\textsf{\textit{Res-Peb}}(G) \geq \textsf{\textit{BW-Peb}}(G) \ .$$

3. The Resolution-Pebbling Game

In this section we define our modified pebble game that will be used to analyze resolution refutations. The next definition is similar to [27], but somewhat simpler.

Definition 3.1 (Res-pebbling subconfiguration). If B and W are sets of vertices in a DAG G with $B \neq \emptyset$, $B \cap W = \emptyset$, we say that $[B]\langle W\rangle$ is a *res-pebbling subconfiguration*, or just *subconfiguration*, in G with black pebbles on B and white pebbles on W *supporting* B. A set of subconfigurations $\mathbb{R} = \{[B_i]\langle W_i\rangle \big| i = 1, \ldots, m\}$ is a *res-pebbling configuration* and its *cost* is $\textsf{\textit{cost}}(\mathbb{R}) = |\bigcup_{i=1}^m (B_i \cup W_i)|$.

The game that we play with subconfigurations is also similar to that in [27], although noticeably less complicated.

Definition 3.2 (Resolution-pebbling game). For G a DAG, a *resolution-pebbling*, or *res-pebbling* for short, is a sequence $\mathcal{R} = \{\mathbb{R}_0, \ldots, \mathbb{R}_\tau\}$ of pebbling clause configurations such that for every $t \in [\tau]$, the configuration $\mathbb{R}_t$ is obtained from $\mathbb{R}_{t-1}$ by one of the following rules:

Download $\mathbb{R}_t = \mathbb{R}_{t-1} \cup \{[v]\langle pred(v)\rangle\}$, where $pred(v)$ denotes the set of predecessors of v. (Notice $pred(v) = \emptyset$ for a source node v.)

Resolution $\mathbb{R}_t = \mathbb{R}_{t-1} \cup \{[B_1 \cup B_2]\langle W_1 \cup W_2\rangle\}$ if there exist $[B_1]\langle W_1 \cup \{v\}\rangle$ and $[B_2 \cup \{v\}]\langle W_2\rangle$ in $\mathbb{R}_{t-1}$ such that $B_1 \cap W_2 = \emptyset$.

Weakening $\mathbb{R}_t = \mathbb{R}_{t-1} \cup \{[B \cup B']\langle W \cup W'\rangle\}$ if $[B]\langle W\rangle \in \mathbb{R}_{t-1}$ and $(B \cup B') \cap (W \cup W') = \emptyset$.

Erasure $\mathbb{R}_t = \mathbb{R}_{t-1} \setminus \{[B]\langle W\rangle\}$ for $[B]\langle W\rangle \in \mathbb{R}_{t-1}$.

The cost of a resolution-pebbling is $\textsf{\textit{cost}}(\mathcal{R}) = \max_{t\in[\tau]}\{\textsf{\textit{cost}}(\mathbb{R}_t)\}$. The *resolution-pebbling price* of G is the minimal cost of a resolution pebbling starting with $\mathbb{R}_0 = \emptyset$ and ending with $\mathbb{R}_\tau = \{[z]\langle\emptyset\rangle\}$ where z is the sink of G.

[2] This was not known at the time of the original theorem in [20]. What is needed is an explicit construction of superconcentrators of linear density, and it has since been shown in [18] how to do this with [2] presenting the currently best construction.

Let us try to provide some intuition for the pebbling rules. We interpret a subconfiguration $[B]\langle W\rangle$ as saying "If all vertices in W have a white pebble on them, then a black pebble can be placed somewhere in B via a legal sequence of black-white pebbling moves." A res-pebbling configuration is a set of such statements and the res-pebbling game is a system that allows for deducing new true statements from existing ones. Indeed, going over the four allowed moves in Definition 3.2 one can verify that they give rise to legal statements. For instance, a download step allows us to state: "If all predecessors of v have a white pebble, then a black pebble may be placed on v." The case of the Resolution rule is perhaps the most subtle so we will describe it in detail. The pair *(i)* $[B_1]\langle W_1 \cup \{v\}\rangle$, *(ii)* $[B_2 \cup \{v\}]\langle W_2\rangle$ says: *(i)* "If white pebbles are placed on $W_1 \cup \{v\}$ we may place a black pebble somewhere in B_2", and *(ii)* "If white pebbles are placed on W_2 we may place a black pebble somewhere on $B_2 \cup \{v\}$". The new statement derived by a resolution step says: *(iii)* "If $W_1 \cup W_2$ are covered by white pebbles then a black pebble may be placed somewhere on $B_1 \cup B_2$." Indeed, if all of $W_1 \cup W_2$ have white pebbles, then by statement *(ii)* we know a black pebble may be placed somewhere on $B_2 \cup \{v\}$. If it is placed in B_2 we are done because *(iii)* is true. Otherwise, the black pebble is placed on v. Then by statement *(i)* a black pebble may be placed somewhere on B_1 after which the black pebble can be removed from v. This shows why, intuitively, the resolution step should be valid. The cases of weakening and erasure can be argued in a similar fashion.

4. Resolution Derivations Induce Res-Pebblings

The proof of Theorem 2.6 follows from two main steps. The first step argues that every refutation π of $Peb^d_G[\oplus]$ induces a res-pebbling $\mathcal{R}_\pi$. The second step says that the cost of the induced res-pebbling $\mathcal{R}_\pi$ is a lower bound on the space of π. Together, these two steps imply Theorem 2.6.

In this section, we do the first step by showing how resolution derivations can be interpreted in terms of resolution-pebblings. As in [25, 27], we get a cleaner correspondence between resolution and pebbling if we ignore the sink axioms $\overline{\bigoplus}_{i=1}^{d} z_i$ and instead study resolution derivations of $\bigoplus_{i=1}^{d} z_i$ from the rest of the formula rather than refutations of all of $Peb^d_G[\oplus]$. Let us write ${}^*Peb^d_G[\oplus] = Peb^d_G[\oplus] \setminus \big\{\overline{\bigoplus}_{i=1}^{d} z_i\big\}$ to denote the pebbling formula over G with the sink axioms in the pebbling contradiction removed. The next lemma is the formal statement saying that as long as we keep the pebbling degree d constant, we may just as well study resolution derivations of $\bigoplus_{i=1}^{d} z_i$ from ${}^*Peb^d_G[\oplus]$ instead of refutations of $Peb^d_G[\oplus]$ without losing more than a constant term. The proof, which is similar to [25, 27], is omitted due to space constraints.

Lemma 4.1. *For any DAG G with sink z, it holds that* $Sp\big(Peb^d_G[\oplus] \vdash 0\big) = Sp\big({}^*Peb^d_G[\oplus] \vdash \bigoplus_{i=1}^{d} z_i\big) + \mathrm{O}\big(2^d\big)$.

In view of Lemma 4.1, from now on we will only consider resolution derivations from ${}^*Peb^d_G[\oplus]$ and translate clause configurations in such derivations into sets of black and white pebbles. Note that since ${}^*Peb^d_G[\oplus]$ is non-contradictory and resolution is sound, any clause set $\mathbb{C}$ derived from ${}^*Peb^d_G[\oplus]$ is satisfiable. We next specify how to translate clauses to pebbles.

Definition 4.2 (Induced res-pebbling subconfiguration). Let G be a DAG and $\mathbb{C}$ a set of clauses derived from ${}^*Peb^d_G[\oplus]$. Then $\mathbb{C}$ induces the res-pebbling subconfiguration $[B]\langle W\rangle$ if

$$\mathbb{C} \vDash \Big(\bigvee_{b \in B} \bigoplus\nolimits_{i=1}^{d} b_i\Big) \vee \Big(\bigvee_{w \in W} \overline{\bigoplus}\nolimits_{i=1}^{d} w_i\Big) \tag{1a}$$

but for all strict subsets $B' \subsetneqq B$ and $W' \subsetneqq W$ that

$$\mathbb{C} \nvDash \Big(\bigvee_{b \in B'} \bigoplus\nolimits_{i=1}^{d} b_i\Big) \vee \Big(\bigvee_{w \in W} \overline{\bigoplus}\nolimits_{i=1}^{d} w_i\Big) \text{ , and} \tag{1b}$$

$$\mathbb{C} \nvDash \Big(\bigvee_{b \in B} \bigoplus\nolimits_{i=1}^{d} b_i\Big) \vee \Big(\bigvee_{w \in W'} \overline{\bigoplus}\nolimits_{i=1}^{d} w_i\Big) \ . \tag{1c}$$

To save space, when all conditions (1a)–(1c) hold, we write

$$\mathbb{C} \vartriangleright \Big(\bigvee_{b \in B} \bigoplus\nolimits_{i=1}^{d} b_i\Big) \vee \Big(\bigvee_{w \in W} \overline{\bigoplus}\nolimits_{i=1}^{d} w_i\Big) \tag{2}$$

and refer to this as *precise implication*. We also say that the clause set $\mathbb{C}$ implies $\big(\bigvee_{b \in B} \bigoplus_{i=1}^{d} b_i\big) \vee \big(\bigvee_{w \in W} \overline{\bigoplus}_{i=1}^{d} w_i\big)$ *precisely*. We will also overload the notation and write $\mathbb{C} \vDash [B]\langle W\rangle$, $\mathbb{C} \nvDash [B]\langle W\rangle$, and $\mathbb{C} \vartriangleright [B]\langle W\rangle$ when the corresponding implications or non-implications hold for $\mathbb{C}$ with respect to $\big(\bigvee_{b \in B} \bigoplus_{i=1}^{d} b_i\big) \vee \big(\bigvee_{w \in W'} \overline{\bigoplus}_{i=1}^{d} w_i\big)$. We write

$$\mathbb{R}(\mathbb{C}) = \big\{[B]\langle W\rangle \big| \mathbb{C} \vartriangleright [B]\langle W\rangle\big\} \tag{3}$$

to denote the set of all res-pebbling subconfigurations induced by $\mathbb{C}$.

The following theorem forms the first part of the proof of Theorem 2.6 and says that resolution derivations induce legal res-pebbling sequences.

Theorem 4.3. *Let $\pi = \{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\}$ be a resolution derivation of $\bigoplus_{i=1}^{d} z_i$ from ${}^*Peb^d_G[\oplus]$. Then the induced res-pebbling configurations $\{\mathbb{R}(\mathbb{C}_0), \ldots, \mathbb{R}(\mathbb{C}_\tau)\}$ form the "backbone" of a complete res-pebbling $\mathcal{R}$ of G in the sense that*

1. $\mathbb{R}(\mathbb{C}_0) = \emptyset$,

2. *$\mathbb{R}(\mathbb{C}_\tau) = \{[z]\langle\emptyset\rangle\}$, and*

3. *for every $t \in [\tau]$, the transition $\mathbb{R}(\mathbb{C}_{t-1}) \rightsquigarrow \mathbb{R}(\mathbb{C}_t)$ can be accomplished in accordance with the res-pebbling rules in cost $\max\{\textbf{\textsf{cost}}(\mathbb{R}(\mathbb{C}_{t-1})), \textbf{\textsf{cost}}(\mathbb{R}(\mathbb{C}_t))\} + \mathrm{O}(1)$.*

In particular, to any resolution derivation $\pi : {}^*Peb^d_G[\oplus] \vdash \bigoplus_{i=1}^d z_i$ we can associate a complete res-pebbling $\mathcal{R}_\pi$ of G such that $\textbf{\textsf{cost}}(\mathcal{R}_\pi) \leq \max_{\mathbb{C} \in \pi}\{\textbf{\textsf{cost}}(\mathbb{R}(\mathbb{C}))\} + \mathrm{O}(1)$.

Due to space limitations we omit the proof of Theorem 4.3 but let us try to describe in words what the theorem says. Using the translation of clauses into pebbles in Definition 4.2, clause configurations $\mathbb{C}_0, \mathbb{C}_1, \ldots, \mathbb{C}_\tau$ in a resolution derivation π can be seen to correspond to "snapshots" at different time intervals of a res-pebbling $\mathcal{R}_\pi$ of the DAG G. Furthermore, the cost of this pebbling is essentially upper-bounded by the largest cost we see at any of the snapshots. There may be many pebbling moves needed to go from the pebble configuration corresponding to $\mathbb{C}_{t-1}$ to the one corresponding to $\mathbb{C}_t$, but the maximal cost during this intermediate pebbling moves is at most an additive constant larger than the cost of the pebble configuration corresponding to $\mathbb{C}_{t-1}$ or $\mathbb{C}_t$. Next we use this to show that the cost of the res-pebbling $\mathcal{R}_\pi$ yields a lower bound on the space of the resolution refutation π.

5. Comparing Resolution Space and Res-Pebbling Cost

In this section, we provide the second component in the proof of Theorem 2.6, namely, that the cost of the induced resolution pebbling $\mathcal{R}_\pi$ is a lower bound on the space of π.

We introduce some notation to make the argument more concise. Let us write $\mathit{Vars}^d(u) = \{u_1, \ldots, u_d\}$. We say that a vertex u is *represented* in a clause C derived from ${}^*Peb^d_G[\oplus]$, or that C *mentions* u, if $\mathit{Vars}^d(u) \cap \mathit{Vars}(C) \neq \emptyset$. We write

$$V(C) = \{u \in V(G) \,|\, \mathit{Vars}^d(u) \cap \mathit{Vars}(C) \neq \emptyset\} \qquad (4)$$

to denote all vertices represented in C. We will also refer to $V(C)$ as the set of vertices *mentioned* by C. This notation is extended to sets of clauses by taking unions.

The main component in the proof of Theorem 2.6 is the following theorem. We remark that this is the place in the proof where it is absolutely crucial that we are working with XOR-pebbling contradictions $Peb^d_G[\oplus]$ and not the "standard" pebbling contradictions $Peb^d_G[\vee]$ defined in terms logical or that were used in [6, 25, 27].

Theorem 5.1. *For every clause configuration $\mathbb{C}$ that is derived from ${}^*Peb^d_G[\oplus]$ with $d > 1$, it holds that*

$$|\mathbb{C}| > \textbf{\textsf{cost}}(\mathbb{R}(\mathbb{C})) \ ,$$

where $\textbf{\textsf{cost}}(\mathbb{R}(\mathbb{C})) = \left|\bigcup_{[B]\langle W\rangle \in \mathbb{R}(\mathbb{C})}(W \cup B)\right|$.

Proof. Let us write

$$V^* = \bigcup_{[B]\langle W\rangle \in \mathbb{R}(\mathbb{C})}(B \cup W) \qquad (5)$$

to denote all vertices mentioned in the configuration induced by $\mathbb{C}$. At this point, we know nothing about the relationship between V^* and $V(\mathbb{C})$. However, it is intuitively plausible that $V^* \subseteq V(\mathbb{C})$, i.e., that the clause set must mention variables for the vertices on which it induces pebbles, and as we will see later in the proof this is indeed the case.

Consider the bipartite graph with clauses in $\mathbb{C}$ on the left-hand side and vertices in V^* on the right-hand side. We draw an edge between $C \in \mathbb{C}$ and $v \in V^*$ if C mentions v. That is, the set of neighbors of C is $N(C) = V(C) \cap V^*$.

Let $\mathbb{C}_1 \subseteq \mathbb{C}$ be a set of maximal size such that $|\mathbb{C}_1| > |N(\mathbb{C}_1)|$. Let $\mathbb{C}_2 = \mathbb{C} \setminus \mathbb{C}_1$ and define the vertex set $V_1^* = N(\mathbb{C}_1)$. By the maximality of $\mathbb{C}_1$ we have

$$|\mathbb{D}| \leq \left|N(\mathbb{D}) \setminus V_1^*\right| \text{ for all } \mathbb{D} \subseteq \mathbb{C}_2. \qquad (6)$$

This holds trivially in the case $\mathbb{C}_2 = \emptyset$. For the case of nonempty $\mathbb{C}_2$, if, by way of contradiction, $|\mathbb{D}| > \left|N(\mathbb{D}) \setminus V_1^*\right|$, then $\mathbb{C}' = \mathbb{C}_1 \cup \mathbb{D}$ would be a larger set than $\mathbb{C}_1$ with $|\mathbb{C}'| > |N(\mathbb{C}')|$, contradicting the maximality of $\mathbb{C}_1$.

Equation (6) implies, by Hall's marriage theorem, that there is an injective mapping M of $\mathbb{C}_2$ into $V^* \setminus V_1^*$. For $C \in \mathbb{C}_2$ let $v(C) = M(C)$ be the vertex matched to C and let $V_2^* = \{v(C) \,|\, C \in \mathbb{C}_2\}$. We now show $V^* = V_1^* \cup V_2^*$ and this will prove the theorem because $|\mathbb{C}_1| > |V_1^*|$ and $|\mathbb{C}_2| = |V_2^*|$ imply

$$|\mathbb{C}| = |\mathbb{C}_1| + |\mathbb{C}_2| > |V_1^*| + |V_2^*| = |V^*|. \qquad (7)$$

Assume by way of contradiction $V_3^* = V^* \setminus (V_1^* \cup V_2^*) \neq \emptyset$. Fix some $v \in V_3^*$ and $[B]\langle W\rangle \in \mathbb{R}(\mathbb{C})$ such that $v \in (W \cup B)$, which must exist by definition of V^*. By Definition 4.2

$$\mathbb{C} \vartriangleright \Big(\bigvee_{b \in B} \bigoplus_{i=1}^d b_i\Big) \vee \Big(\bigvee_{w \in W} \overline{\bigoplus}_{i=1}^d w_i\Big) \ . \qquad (8)$$

We claim that we can construct a truth value assignment α that makes $\mathbb{C}$ true but $\left(\bigvee_{b \in B} \bigoplus_{i=1}^d b_i\right) \vee \left(\bigvee_{w \in W} \overline{\bigoplus}_{i=1}^d w_i\right)$ false. This clearly contradicts condition (1a) from Definition 4.2 and so the theorem follows.

The desired α will be the union of three partial assignments $\alpha_1 \cup \alpha_2 \cup \alpha_3$ that assign values to distinct variables. For $j = 1, 2$ let $B_j = B \cap V_j^*$ and $W_j = W \cap V_j^*$. By assumption $v \in (B \cup W) \setminus (B_1 \cup W_1)$ so conditions (1b), (1c) in Definition 4.2 imply

$$\mathbb{C} \nvDash \Big(\bigvee_{b \in B_1} \bigoplus_{i=1}^{d} b_i \Big) \vee \Big(\bigvee_{w \in W_1} \overline{\bigoplus}_{i=1}^{d} w_i \Big) \tag{9}$$

so we can find a truth value assignment β that sets $\mathbb{C}$ to true but violates all constraints $\bigoplus_{i=1}^{d} b_i$, $b \in B_1$, and $\overline{\bigoplus}_{i=1}^{d} w_i$, $w \in W_1$. Take α_1 to be the restriction of β to $\mathit{Vars}\big(\mathbb{C}_1\big) \cup \mathit{Vars}^d\big(B_1 \cup W_1\big)$. What is important to notice about α_1 is that it *(i)* does not assign any value to $\mathit{Vars}^d\big(V_2^* \cup V_3^*\big)$, *(ii)* sets $\mathbb{C}_1$ to true, *(iii)* violates all constraints $\bigoplus_{i=1}^{d} b_i$, $b \in B_1$, and $\overline{\bigoplus}_{i=1}^{d} w_i$, $w \in W_1$ and *(iv)* any extension of α_1 will not change *(ii)*, *(iii)*.

To construct α_2 we use the matching M of $\mathbb{C}_2$ into V_2^* to find a distinct vertex $v(C)$ for every $C \in \mathbb{C}_2$ and a literal over some variable $v(C)_i \in \mathit{Vars}^d\big(v(C)\big)$ that fixes C to true. Let γ be this partial assignment. We stress that γ assigns values to at most one variable v_i for any $v \in B_2 \cup W_2$. This means that we can extend γ to an assignment α_2 to $\mathit{Vars}^d\big(V_2^*\big)$ still satisfying $\mathbb{C}_2$ but violating all constraints $\bigoplus_{i=1}^{d} b_i$, $b \in B_2$, and $\overline{\bigoplus}_{i=1}^{d} w_i$, $w \in W_2$. Regarding α_2, notice it *(i)* assigns values only to $\mathit{Vars}^d\big(V_2^*\big)$, *(ii)* sets $\mathbb{C}_2$ to true, *(iii)* violates all constraints $\bigoplus_{i=1}^{d} b_i$, $b \in B_2$, and $\overline{\bigoplus}_{i=1}^{d} w_i$, $w \in W_2$ and *(iv)* any extension of α_2 will not change *(ii)*, *(iii)*.

Finally, to construct α_3 we pick for every $v \in (B \cup W) \cap V_3^*$ an assignment that violates the constraint over v. I.e., if $v \in B$ we set α_3 so that $\bigoplus_{i=1}^{d} v_i$ is false and if $v \in W$ set it so that $\overline{\bigoplus}_{i=1}^{d} v_i$ is false. Notice α_3 assigns values only to variables in $\mathit{Vars}^d\big(V_3^*\big)$. Thus, taking $\alpha = \alpha_1 \cup \alpha_2 \cup \alpha_3$ contradicts (8), which proves the claim. □

Theorem 2.6 now follows from Theorems 4.3 and 5.1 together with Lemma 4.1.

6. From Res-Pebblings to Black-White Pebblings

To complete the proof of Theorem 1.1, we also need to establish lower bounds on res-pebbling price in terms of black-white pebbling price.

Theorem 2.7 (restated). *For any DAG G it holds that $\textsf{Res-Peb}(G) \geq \textsf{BW-Peb}(G)$.*

On the face of it, the resolution-pebbling game might seem quite different from the standard black-white pebble game. The lower bounds on black-white pebbling depend critically on the fact that the rules for black pebble placement and white pebble removal are very strict. In the resolution-pebbling game, however, we can always remove any white pebbles by doing an erasure, and by weakening we can always black-pebble any vertex although no white pebbles are even near this vertex. However, the fact that we collect black pebbles B and white pebbles W in subconfigurations $[B]\langle W\rangle$, and only allow operations on these subconfigurations, makes it relatively straightforward to show Theorem 2.7. The proof follows immediately from the following pair of lemmas, proved next.

Lemma 6.1. *Given any complete res-pebbling $\mathcal{R}$ of G using weakening, there is a complete res-pebbling $\mathcal{R}'$ which never makes any weakening moves and has $\textsf{cost}(\mathcal{R}') \leq \textsf{cost}(\mathcal{R})$.*

Lemma 6.2. *Given any complete res-pebbling $\mathcal{R}'$ of G that does not make any weakening moves, there is a complete standard black-white pebbling $\mathcal{P}$ of G such that $\textsf{cost}(\mathcal{P}) \leq \textsf{cost}(\mathcal{R}')$.*

Proof of Lemma 6.1. This is true since we can always construct a shadow pebbling that matches download, resolution, and erasure moves but ignores weakening moves. Such a pebbling can have at most the same cost as the pebbling that it is shadowing.

Formally, given any complete res-pebbling $\mathcal{R} = \big\{\mathbb{R}_0, \ldots, \mathbb{R}_\tau\big\}$ of G, we construct our pebbling $\mathcal{R}' = \big\{\mathbb{R}'_0, \ldots, \mathbb{R}'_\tau\big\}$ inductively by maintaining the following invariant: For every $\mathbb{R}_t \in \mathcal{R}$ there is a surjective function $g_t : \mathbb{R}_t \mapsto \mathbb{R}'_t$ such that whenever $g_t([B]\langle W\rangle) = [b]\langle W_b\rangle$ it holds that $b \in B$ and $W_b \subseteq W$. If we can construct such a function g_t for every t we are clearly done, since $\textsf{cost}(\mathbb{R}'_t) = \textsf{cost}(g_t(\mathbb{R}_t)) \leq \textsf{cost}(\mathbb{R}_t)$ and we must have $g_\tau([z]\langle\emptyset\rangle) = \{[z]\langle\emptyset\rangle\}$. The base case $\mathbb{R}_0 = \emptyset$ is trivial. We make a case analysis over the pebbling move made at time t.

Download $\mathbb{R}_t = \mathbb{R}_{t-1} \cup \{[v]\langle \mathit{pred}(v)\rangle\}$: Make the same download move in $\mathcal{R}'$, set $g_t([v]\langle \mathit{pred}(v)\rangle) = [v]\langle \mathit{pred}(v)\rangle$ and let $g_t = g_{t-1}$ for all other subconfigurations in $\mathbb{R}_{t-1}$.

Erasure $\mathbb{R}_t = \mathbb{R}_{t-1} \setminus \big\{[B]\langle W\rangle\big\}$: Set $\mathbb{R}'_t = g_{t-1}(\mathbb{R}_t)$ (which might result in an erasure or leave $\mathbb{R}'_t = \mathbb{R}'_{t-1}$ unchanged).

Weakening $\mathbb{R}_t = \mathbb{R}_{t-1} \cup \big\{[B \cup B']\langle W \cup W'\rangle\big\}$ for some subconfiguration $[B]\langle W\rangle \in \mathbb{R}_{t-1}$: set $g_t([B \cup B']\langle W \cup W'\rangle) = g_{t-1}([B]\langle W\rangle)$ and let $g_t = g_{t-1}$ for all other subconfigurations (leaving $\mathbb{R}'_t = \mathbb{R}'_{t-1}$ unchanged).

Resolution $\mathbb{R}_t = \mathbb{R}_{t-1} \cup \big\{[B_1 \cup B_2]\langle W_1 \cup W_2\rangle\big\}$ derived from $[B_1]\langle W_1 \cup \{v\}\rangle, [B_2 \cup \{v\}]\langle W_2\rangle \in \mathbb{R}_{t-1}$: This is the only nontrivial case. Let $g_{t-1}([B_1]\langle W_1 \cup \{v\}\rangle) = [b_1]\langle W'_1\rangle$ and $g_{t-1}([B_2 \cup \{v\}]\langle W_2\rangle) = [b_2]\langle W'_2\rangle$. Note that by the induction hypothesis we have $b_1 \in B_1 \subseteq B_1 \cup B_2$ and $W'_2 \subseteq W_2 \subseteq W_1 \cup W_2$. We get three subcases:

1. $v \notin W'_1$: Then $W'_1 \subseteq W_1 \subseteq W_1 \cup W_2$, so we can set $g_t([B_1 \cup B_2]\langle W_1 \cup W_2\rangle) = [b_1]\langle W'_1\rangle$.
2. $v \neq b_2$: Then $b_2 \in B_2 \subseteq B_1 \cup B_2$, so we can set $g_t([B_1 \cup B_2]\langle W_1 \cup W_2\rangle) = [b_2]\langle W'_2\rangle$.
3. Otherwise, we have $v = b_2$ and $v \in W'_1$, so we can resolve $[b_1]\langle W'_1\rangle$ and $[b_2]\langle W'_2\rangle$ to get $[b_1]\langle (W'_1 \cup W'_2) \setminus \{b_2\}\rangle$ and set $g_t([B_1 \cup B_2]\langle W_1 \cup W_2\rangle) = [b_1]\langle (W'_1 \cup W'_2) \setminus \{b_2\}\rangle$.

Let $g_t = g_{t-1}$ for all other subconfigurations in $\mathbb{R}_{t-1}$.

Since in all cases we can construct a surjective function $g_t : \mathbb{R}_t \mapsto \mathbb{R}'_t$ satisfying the invariant conditions, the lemma follows. □

Proof of Lemma 6.2. We assume without loss of generality that $\mathcal{R}'$ terminates at time τ once it contains a subconfiguration $[z]\langle\emptyset\rangle$ where z is the sink of G. Next, we define the *essential* subconfigurations of $\mathcal{R}'$ by backwards induction as follows. The only essential subconfiguration of $\mathbb{R}_\tau$ is $[z]\langle\emptyset\rangle$. For $t < \tau$, we say a subconfiguration is essential in $\mathbb{R}_t$ iff it is either *(i)* essential at time $t+1$, or *(ii)* one of the two subconfigurations used in a resolution step resulting in an essential subconfiguration. To prove the lemma it is sufficient to show that the set of pebbles mentioned in essential subconfigurations forms a legal black-white pebbling of G. Formally, let

$$\mathbb{B}_t = \{\cup B \mid [B]\langle W\rangle \text{ is essential in } \mathbb{R}_t\}$$

and

$$\mathbb{W}_t = \{\cup W \mid [B]\langle W\rangle \text{ is essential in } \mathbb{R}_t\} \setminus \mathbb{B}_t .$$

We claim the sequence $\{(\mathbb{B}_0, \mathbb{W}_0), \ldots, (\mathbb{B}_\tau, \mathbb{W}_\tau)\}$ is a legal black-white pebbling of G and this proves our lemma.

By construction $\mathbb{B}_0 = \mathbb{W}_0 = \emptyset$ and $\mathbb{B}_\tau = \{z\}, \mathbb{W}_\tau = \emptyset$ so we only need to argue that intermediate steps are legal black-white moves. By definition of essentiality we do not need to worry about erasure moves because only unessential clauses can be erased. Thus, if the t^{th} step is an erasure then $(\mathbb{B}_{t-1}, \mathbb{W}_{t-1}) = (\mathbb{B}_t, \mathbb{W}_t)$. By assumption, there are no weakening moves so we only need to handle downloads and resolution steps which is what we do next.

Download Suppose the t^{th} step is a download of an essential subconfiguration corresponding to vertex v. Then $\mathbb{B}_t = \mathbb{B}_{t-1} \cup \{v\}$ and $\mathbb{W}_t = (\mathbb{W}_t \cup \mathit{pred}(v)) \setminus (\mathbb{B}_{t-1} \cup) \{v\})$ and this transition corresponds to a sequence of legal pebbling moves involving *(i)* placing white pebbles on all predecessors of v that are not covered by $\mathbb{W}_t \cup \mathbb{B}_t$, *(ii)* removing a white pebble from v, if $v \in \mathbb{W}_t$, which is legal because all of v's predecessors are pebbled, and *(iii)* placing a black pebble on v. Notice the overall number of pebbles throughout this sequence is at most $|\mathbb{B}_t \cup \mathbb{W}_t|$.

Resolution Suppose the t^{th} move is a resolution step deriving an essential subconfiguration. By definition, the two subconfigurations used in the resolution step are essential at time $t-1$. Furthermore, if v is the vertex that is removed in this step we have $v \in \mathbb{B}_{t-1}$. Inspection reveals $\mathbb{B}_{t-1} \supseteq \mathbb{B}_t \supseteq \mathbb{B}_{t-1} \setminus \{v\}$ and $\mathbb{W}_t \supseteq \mathbb{W}_{t-1}$ implying we can reach $(\mathbb{B}_t, \mathbb{W}_t)$ by a legal sequence of pebbling moves because we need only remove the black pebble from v and perhaps place a white one on it. This completes the proof of the lemma and with it the proof of Theorem 2.7 is complete. □

7. Concluding Remarks

We have proven an asymptotically optimal separation of space and length in resolution. This answers an open question discussed in, for instance, [17, 33, 35].

It would be interesting to see if the proof technique used in this paper can be extended to yield length-space tradeoffs in the sense that there are CNF formulas that can be refuted in short length and small space, but where any short refutation must have large space.[3]

Another natural question is whether our lower bounds can be extended to stronger proof systems than resolution. One obvious candidate would be the k-DNF resolution proof systems $\mathfrak{R}(k)$ introduced by Krajíček [24], where the lines in the proofs are k-DNF formulas instead of clauses and one can "resolve" over up to k variables simultaneously. We believe that XOR-pebbling contradictions $Peb_G^{k+1}[\oplus]$ should separate k-DNF resolution and $(k+1)$-DNF resolution with respect to space. If so, this would establish that the k-DNF resolution proof systems form a strict hierarchy with respect to space. Currently, all that is known is the separation result in [15] for the restricted case of tree-like k-DNF resolution.

Acknowledgements

The second author wants to acknowledge that Albert Atserias mentioned already back in 2006, albeit in a slightly different context, that working on XOR-pebbling contradictions might be a possible way forward.

[3] At the time of submission of this paper, the answer seems to be a firm yes, but the manuscript is still in a very early stage of preparation.

References

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

[2] N. Alon and M. Capalbo. Smaller explicit superconcentrators. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '03)*, pages 340–346, 2003.

[3] A. Atserias and V. Dalmau. A combinatorical characterization of resolution width. In *Proceedings of the 18th IEEE Annual Conference on Computational Complexity (CCC '03)*, pages 239–247, July 2003. Journal version to appear in *Journal of Computer and System Sciences*.

[4] P. Beame, R. Karp, T. Pitassi, and M. Saks. The efficiency of resolution and Davis-Putnam procedures. *SIAM Journal on Computing*, 31(4):1048–1075, 2002.

[5] P. Beame, H. Kautz, and A. Sabharwal. Understanding the power of clause learning. In *Proceedings of the 18th International Joint Conference in Artificial Intelligence (IJCAI '03)*, pages 94–99, 2003.

[6] E. Ben-Sasson. Size space tradeoffs for resolution. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 457–464, May 2002.

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

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

[9] A. Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.

[10] M. L. Bonet, J. L. Esteban, N. Galesi, and J. Johannsen. On the relative complexity of resolution refinements and cutting planes proof systems. *SIAM Journal on Computing*, 30(5):1462–1484, 2000.

[11] V. Chvátal and E. Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, Oct. 1988.

[12] S. A. Cook and R. Sethi. Storage requirements for deterministic polynomial time recognizable languages. *Journal of Computer and System Sciences*, 13(1):25–37, 1976.

[13] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem proving. *Communications of the ACM*, 5(7):394–397, July 1962.

[14] M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.

[15] J. L. Esteban, N. Galesi, and J. Messner. On the complexity of resolution with bounded conjunctions. *Theoretical Computer Science*, 321(2-3):347–370, Aug. 2004.

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

[17] J. L. Esteban and J. Torán. A combinatorial characterization of treelike resolution space. *Information Processing Letters*, 87(6):295–300, 2003.

[18] O. Gabber and Z. Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, June 1981.

[19] Z. Galil. On resolution with clauses of bounded size. *SIAM Journal on Computing*, 6(3):444–459, 1977.

[20] J. R. Gilbert and R. E. Tarjan. Variations of a pebble game on graphs. Technical Report STAN-CS-78-661, Stanford University, 1978.

[21] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, Aug. 1985.

[22] J. Hopcroft, W. Paul, and L. Valiant. On time versus space. *Journal of the ACM*, 24(2):332–337, Apr. 1977.

[23] H. Kautz and B. Selman. The state of SAT. *Discrete Applied Mathematics*, 155(12):1514–1524, June 2007.

[24] J. Krajíček. On the weak pigeonhole principle. *Fundamenta Mathematicae*, 170(1-3):123–140, 2001.

[25] J. Nordström. Narrow proofs may be spacious: Separating space and width in resolution (Extended abstract). In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC '06)*, pages 507–516, May 2006.

[26] J. Nordström. *Short Proofs May Be Spacious: Understanding Space in Resolution*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, May 2008. Available at `http://www.csc.kth.se/~jakobn/research/`.

[27] J. Nordström and J. Håstad. Towards an optimal separation of space and length in resolution. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*. ACM, May 2008.

[28] R. Raz and P. McKenzie. Separation of the monotone NC hierarchy. *Combinatorica*, 19(3):403–435, Mar. 1999.

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

[30] A. Sabharwal. *Algorithmic Applications of Propositional Proof Complexity*. PhD thesis, University of Washington, Seattle, 2005.

[31] A. Sabharwal, P. Beame, and H. Kautz. Using problem structure for efficient clause learning. In *6th International Conference on Theory and Applications of Satisfiability Testing (SAT '03), Selected Revised Papers*, volume 2919 of *Lecture Notes in Computer Science*, pages 242–256. Springer, 2004.

[32] The international SAT Competitions web page. `http://www.satcompetition.org`.

[33] N. Segerlind. The complexity of propositional proofs. *Bulletin of Symbolic Logic*, 13(4):482–537, Dec. 2007.

[34] J. Torán. Lower bounds for space in resolution. In *Proceedings of the 13th International Workshop on Computer Science Logic (CSL '99)*, volume 1683 of *Lecture Notes in Computer Science*, pages 362–373. Springer, 1999.

[35] J. Torán. Space and width in propositional resolution. *Bulletin of the European Association for Theoretical Computer Science*, 83:86–104, June 2004.

[36] G. Tseitin. On the complexity of derivation in propositional calculus. In A. O. Silenko, editor, *Structures in Constructive Mathematics and Mathematical Logic, Part II*, pages 115–125. Consultants Bureau, New York-London, 1968.

[37] A. Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, Jan. 1987.