Narrow Proofs May Be Spacious: Separating Space and Width in Resolution

[Extended Abstract] *

Jakob Nordström
Royal Institute of Technology (KTH)
SE-100 44 Stockholm, Sweden
jakobn@kth.se

ABSTRACT

The width of a resolution proof is the maximal number of literals in any clause of the proof. The space of a proof is the maximal number of clauses kept in memory simultaneously if the proof is only allowed to infer new clauses from clauses currently in memory. Both of these measures have previously been studied and related to the resolution refutation size of unsatisfiable CNF formulas. Also, the refutation space of a formula has been proven to be at least as large as the refutation width, but it has been open whether space can be separated from width or the two measures coincide asymptotically. We prove that there is a family of k-CNF formulas for which the refutation width in resolution is constant but the refutation space is non-constant, thus solving a problem mentioned in several previous papers.

Categories and Subject Descriptors: F.4.1[Mathematical Logic and Formal Languages]: Mathematical Logic—*proof theory*; F.1.3[Computation by Abstract Devices]: Complexity Measures and Classes—*Relations among complexity measures*;

General Terms: Theory

Keywords: Proof complexity, resolution, width, space, separation, lower bound, pebble game, pebbling contradiction

1. INTRODUCTION

A proof system for a language L is an algorithm $P(x, \pi)$ which runs in time polynomial in $|x|$ and $|\pi|$ such that for all $x \in L$ there is a string π (a *proof*) for which $P(x, \pi) = 1$. For $x \notin L$, it should hold for all strings π that $P(x, \pi) = 0$. The complexity of a proof system P is the smallest bounding function $g : \mathbb{N} \mapsto \mathbb{N}$ such that $x \in L$ if and only if there is a proof π of size $|\pi| \leq g(|x|)$ for which $P(x, \pi) = 1$. If a proof system is of polynomial complexity, it is said to be polynomially bounded. A *propositional proof system* is a proof system for tautologies in propositional logic.

Studying the complexity of proving tautologies in propositional logic is an important problem both from a theoretical and a practical point of view. On the one hand, propositional proof complexity is closely related to central questions in computational complexity, in view of the result of Cook and Reckhow [15] that separating NP and $\mathsf{co\text{-}NP}$ (which would imply $\mathsf{P} \neq \mathsf{NP}$) is equivalent to obtaining superpolynomial lower bounds for arbitrary propositional proof systems. On the other hand, designing efficient algorithms for proving tautologies, or equivalently refuting unsatisfiable formulas, is a very important problem in applied research and in industry, for instance in the context of formal methods.

Perhaps the most studied propositional proof system is *resolution*, which appeared in [10] and began to be investigated in connection with automated theorem proving in the 1960s [16, 17, 32]. Because of its simplicity—there is only one derivation rule—and because all lines in a proof are clauses, this proof system is well adapted to proof search algorithms. Many real-world automated theorem provers are based on resolution.

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 proofs, i.e., the number of clauses, which is easily seen to be polynomially related to the proof size. In 1968, Tseitin [35] presented a superpolynomial lower bound on proof length for a restricted form of resolution, called *regular* resolution, but it was not until almost 20 years later that Haken [23] proved the first superpolynomial lower bound for general resolution. This exponential lower bound of Haken was later followed by many similar results, for instance in [5, 14, 31, 36].

A second complexity measure for resolution, first made explicit by Galil [21], is the *width*, measured as the maximal size of a clause in the proof. Ben-Sasson and Wigderson [9] showed a strong upper bound on width in terms of length, and used this to rederive and simplify many lower bounds on length by proving lower bounds on width.

The results on width lead to the question of whether other complexity measures could yield interesting insights as well. In [19, 33], Esteban and Torán introduced the concept of proof *space*, transforming a previous definition from [25]. Intuitively, the space of a resolution proof is the maximal number of clauses one needs to keep in memory while verifying the proof. A number of upper and lower bounds for

*A full version of this paper is available as ECCC Technical Report TR05-066 at `www.eccc.uni-trier.de/eccc/`.

proof space in resolution and other proof systems were subsequently presented in for instance [1, 7, 18, 20]. In several of these papers, it was noted that the lower bounds on proof space in resolution for different formula families matched known lower bounds on proof width. Atserias and Dalmau [3] showed that this was not a coincidence, but that the minimal proof space of a k-CNF formula is always at least as large as the minimal proof width minus a constant.

An immediate follow-up question to this is whether proof space and proof width always coincide or there is a formula family that separates the two measures asymptotically.

Another natural question is how space and length are related. Is there a Ben-Sasson-Wigderson-style upper bound on space in terms of length, or can short resolution proofs be arbitrarily complex with respect to space?

A third, intimately connected question is to determine the proof space of *pebbling contradictions* defined in terms of pebble games on directed acyclic graphs. Non-constant lower bounds on the space of refuting pebbling contradictions would separate space and width, and possibly also clarify the relation between space and length if the bounds were good enough. On the other hand, a constant upper bound would improve the trade-off results for different measures in resolution in [6].

The above three questions have been mentioned as interesting open problems in [6, 18, 20, 34].

In this paper, we answer the first question by separating space and width. This is done by proving an asymptotically tight bound on space for pebbling contradictions over binary trees, thus at least partially solving the open problem about the space complexity of pebbling contradictions as well. More precisely, our results are as follows (formal definitions are given in Sections 2 and 4).

THEOREM 1. *Let T_h denote the complete binary tree of height h and $Peb^d_{T_h}$ the pebbling contradiction of degree $d \geq 2$ defined over T_h. Then the space of refuting $Peb^d_{T_h}$ in resolution is $Sp\big(Peb^d_{T_h} \vdash 0\big) = \Theta(h)$.*

COROLLARY 2. *For all $k \geq 4$, there exists a family of k-CNF formulas $\big\{F_n\big\}_{n=1}^{\infty}$ of size $\mathrm{O}(n)$ with refutation width $W(F_n \vdash 0) = \mathrm{O}(1)$ and refutation space $Sp(F_n \vdash 0) = \Theta(\log n)$.*

The organization of this paper is as follows. We start by presenting the resolution proof system in Section 2. Section 3 gives a short introduction to pebble games, and in Section 4 we review some previous results connecting resolution and pebbling. The bound on refutation space which separates space and width is then proven in three steps.

- First, we define a modified pebble game and establish a lower bound for this game in terms of the standard black-white pebble game (Section 5).
- Next, we show that a resolution refutation of a pebbling contradiction induces a pebbling of the underlying graph in our modified pebble game (Section 6).
- Finally, we prove that if a set of clauses induces many pebbles, the set must contain at least as many clauses (Section 7).

Since a resolution proof induces a pebbling, and such a pebbling must contain many pebbles at some point, we deduce that the space of any resolution proof must be large. We conclude in Section 8 by giving suggestions for further research.

All proofs omitted below can be found in the full-length version of this paper [29].

2. THE RESOLUTION PROOF SYSTEM

Any propositional logic formula can be converted to a formula in conjunctive normal form (CNF) that is only linearly larger and is unsatisfiable if and only if the original formula is a tautology. Therefore, any sound and complete system which produces refutations of unsatisfiable CNF formulas can be considered as a general propositional proof system. One such proof system is *resolution*.

A *literal* is either a propositional logic variable x or its negation $\overline{x}$. We define $\overline{\overline{x}} = x$. A *clause* $C = a_1 \vee \ldots \vee a_k$ is a set of literals. We say that C is a *subclause* of D if $C \subseteq D$. A clause containing at most k literals is called a *k-clause*. A *CNF formula* $F = C_1 \wedge \ldots \wedge C_m$ is a set of clauses. A *k-CNF formula* is a CNF formula consisting of k-clauses. We let $Vars(C)$ denote the set of variables and $Lit(C)$ the set of literals in a clause C. This notation is extended to sets of clauses by taking unions.

When reasoning about the space of resolution derivations, it is convenient to use the definition of resolution introduced by [1]. We employ the standard notation $[n] = \{1, 2, \ldots, n\}$.

DEFINITION 3 (RESOLUTION). *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 a clause $C \in F$ (an* axiom*).*

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

Inference *$\mathbb{C}_t = \mathbb{C}_{t-1} \cup \{C \vee D\}$ for a clause $C \vee D$ inferred by the* resolution rule *from $C \vee x, D \vee \overline{x} \in \mathbb{C}_{t-1}$.*

A resolution derivation $\pi : F \vdash A$ of a clause A from 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 $\pi : F \vdash 0$ of the empty clause 0 (the clause with no literals) from F.*

For a formula F and a set of formulas $\mathcal{G} = \{G_1, \ldots, G_n\}$, we say that $\mathcal{G}$ *implies* F, denoted $\mathcal{G} \vDash F$, if every truth value assignment satisfying all formulas $G_i \in \mathcal{G}$ satisfies F as well. Resolution is sound and implicationally complete. That is, if there is a resolution derivation $\pi : F \vdash A$ then $F \vDash A$, and if $F \vDash A$ then there is a resolution derivation $\pi : F \vdash A'$ for some $A' \subseteq A$. In particular, F is unsatisfiable if and only if there is a resolution refutation of F.

The *length* $L(F)$ of a CNF formula F is $|F|$, and the length $L(\pi)$ of a derivation π is the number of distinct clauses in π. The length of deriving A from F is $L(F \vdash A) = \min_{\pi : F \vdash A}\big\{L(\pi)\big\}$, where the minimum is taken over all derivations of A. The *width* $W(C)$ of a clause C is $|C|$. The width of a set of clauses $\mathbb{C}$ is $W(\mathbb{C}) = \max_{C \in \mathbb{C}}\big\{W(C)\big\}$, the width of a derivation π is $W(\pi) = \max_{\mathbb{C}_t \in \pi}\big\{W(\mathbb{C}_t)\big\}$, and the width of deriving A from F is $W(F \vdash A) = \min_{\pi : F \vdash A}\big\{W(\pi)\big\}$.

The width of refuting a CNF formula F over n variables in resolution was proven in [9] to be bounded from above by $W(F \vdash 0) \leq W(F) + \mathrm{O}\big(\sqrt{n \log L(F \vdash 0)}\big)$. In [12], it was shown that this bound on width in terms of length is essentially optimal.

The *space* [1, 6] of a resolution derivation $\pi = \{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\}$ is the maximal number of clauses in any clause configuration in π, or formally $Sp(\pi) = \max_{t \in [\tau]}\{|\mathbb{C}_t|\}$. The minimal space of deriving a clause A from a CNF formula F is $Sp(F \vdash A) = \min_{\pi : F \vdash A}\{Sp(\pi)\}$.

Any unsatisfiable CNF formula can be refuted in space linear in the formula size, or more precisely in space at most $\min\{L(F) + 1, |\mathit{Vars}(F)| + 2\}$ (see [19]). Thus the interesting question is which formulas demand this much space, and which formulas can be refuted in for instance logarithmic or even constant space. It was shown in [1, 33] that there are polynomial-size formulas that meet the upper bounds on space in terms of the number of clauses and variables up to a multiplicative constant.

Lower bounds on space have been presented for a number of different CNF formula families [1, 7, 33]. In these papers, it was observed that the lower bounds on refutation space coincided with previously proven lower bounds on refutation width, which lead to the conjecture that width is a lower bound for space. This conjecture was confirmed in [3], where it was proven that $Sp(F \vdash 0) - 3 \geq W(F \vdash 0) - W(F)$.[1] In other words, the extra clause space exceeding the minimum 3 needed for any resolution derivation is bounded from below by the extra width exceeding the formula width.

As was mentioned in the introduction, a very natural question, which has remained open, is what holds in the other direction. Do the space and width measures coincide asymptotically, or is there a formula family separating space from width? We remark that in order for this question to be interesting, we should restrict our attention to families of k-CNF formulas. Any resolution refutation of an unsatisfiable CNF formula F with *minimum* clause width k can be shown to require clause space at least $k + 2$ (see [19]), so it is easy to find CNF formulas $\{F_n\}_{n=1}^{\infty}$ of growing width such that $W(F_n \vdash 0) - W(F_n) = \mathrm{O}(1)$ but $Sp(F_n \vdash 0) = \Omega(n)$.

In this paper, we settle the open question of the relationship between space and width by exhibiting for any fixed $k \geq 4$ a k-CNF formula family $\{F_n\}_{n=1}^{\infty}$ such that $W(F_n \vdash 0) = \mathrm{O}(1)$ but $Sp(F_n \vdash 0) = \Theta(\log L(F_n)) = \omega(1)$.

3. PEBBLE GAMES

Pebble games were devised for studying programming languages and compiler construction, but have found a variety of applications in computational complexity theory. In connection with resolution, pebble games have been employed both to analyze resolution derivations with respect to how much memory they consume (using the original definition of space in [19]) and to construct CNF formulas which are hard for different variants of resolution in various respects (see for example [2, 8, 11, 13]). An excellent survey of pebbling up to 1980 is [30].

The black pebbling price of a directed acyclic graph 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. In the following, we let $V(G)$ denote the vertices of G, and say that vertices with indegree 0 are *sources* and vertices with outdegree 0 *targets*.

Definition 4 (Pebble game). *Suppose that G is a directed acyclic graph (DAG) with sources S and unique target z. A* pebble configuration *on G is a pair of subsets $\mathbb{P} = (B, W)$ of $V(G)$, comprising the black-pebbled vertices B and white-pebbled vertices W. A* legal black-white pebbling reaching *(B, W) in G is a sequence of configurations $\mathcal{P} = \{\mathbb{P}_0, \ldots, \mathbb{P}_\tau\}$ such that $\mathbb{P}_0 = (\emptyset, \emptyset)$, $\mathbb{P}_\tau = (B, W)$, and for all $t \in [\tau]$, $\mathbb{P}_t$ follows from $\mathbb{P}_{t-1}$ by one of the following rules:*

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.*

The cost *of a pebble configuration $\mathbb{P} = (B, W)$ is $\mathsf{cost}(\mathbb{P}) = |B \cup W|$ and the cost of a legal pebbling $\mathcal{P} = \{\mathbb{P}_0, \ldots, \mathbb{P}_\tau\}$ is $\max_{t \in [\tau]}\{\mathsf{cost}(\mathbb{P}_t)\}$. The* black-white pebbling price *of a pebble configuration (B, W), denoted $\mathsf{BW\text{-}Peb}(B, W)$, is the minimal cost of any legal pebbling reaching (B, W), and the* black-white pebbling price *of the DAG G is $\mathsf{BW\text{-}Peb}(G) = \mathsf{BW\text{-}Peb}(\{z\}, \emptyset)$.*

A legal black pebbling of G *is a pebbling reaching $(\{z\}, \emptyset)$ using black pebbles only, i.e., $W_t = \emptyset$ for all t, and the* (black) pebbling price *of G, denoted $\mathsf{Peb}(G)$, is the minimal cost of any legal black pebbling of G.*

We let T (T_h) denote a complete binary tree (of height h) considered as a DAG with edges directed towards the root. The black pebbling price of T_h is $\mathsf{Peb}(T_h) = h + 2$, which can be established by an easy induction over the tree height. In [28], general bounds for the black-white pebbling price of trees of any arity were proven, which can be simplified to the tight result $\mathsf{BW\text{-}Peb}(T_h) = \lfloor \frac{h}{2} \rfloor + 3$ for complete binary trees (see Section 4 of [29] for a proof).

4. PEBBLING CONTRADICTIONS

A *pebbling contradiction* defined on a DAG G encodes the pebble game on G by defining the sources to be true and the target false, and specifying that truth propagates through the graph according to the pebbling rules.

Definition 5 (Pebbling contradiction [9]). *Let G be a DAG with sources S and a unique target z and with all vertices having indegree 0 or 2, and let $d \in \mathbb{N}^+$. Associate d distinct variables $x(v)_1, \ldots, x(v)_d$ with every vertex $v \in V(G)$. The dth degree* pebbling contradiction *over G, denoted Peb_G^d, is the conjunction of the following clauses:*

- $\bigvee_{i=1}^{d} x(s)_i$ *for all $s \in S$ (*source axioms*),*
- $\overline{x(z)_i}$ *for all $i \in [d]$ (*target axioms*),*

[1] The theorem in [3] is $Sp(F \vdash 0) \geq W(F \vdash 0) - W(F)$, but this can be sharpened by a constant if one does the calculations in the proofs carefully.

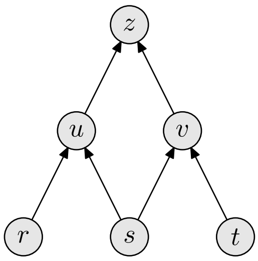

$$
\begin{aligned}
& (x(r)_1 \vee x(r)_2) \\
\wedge\ & (x(s)_1 \vee x(s)_2) \\
\wedge\ & (x(t)_1 \vee x(t)_2) \\
\wedge\ & (\overline{x(r)}_1 \vee \overline{x(s)}_1 \vee x(u)_1 \vee x(u)_2) \\
\wedge\ & (\overline{x(r)}_1 \vee \overline{x(s)}_2 \vee x(u)_1 \vee x(u)_2) \\
\wedge\ & (\overline{x(r)}_2 \vee \overline{x(s)}_1 \vee x(u)_1 \vee x(u)_2) \\
\wedge\ & (\overline{x(r)}_2 \vee \overline{x(s)}_2 \vee x(u)_1 \vee x(u)_2) \\
\wedge\ & (\overline{x(s)}_1 \vee \overline{x(t)}_1 \vee x(v)_1 \vee x(v)_2) \\
\wedge\ & (\overline{x(s)}_1 \vee \overline{x(t)}_2 \vee x(v)_1 \vee x(v)_2)
\end{aligned}
\qquad
\begin{aligned}
\wedge\ & (\overline{x(s)}_2 \vee \overline{x(t)}_1 \vee x(v)_1 \vee x(v)_2) \\
\wedge\ & (\overline{x(s)}_2 \vee \overline{x(t)}_2 \vee x(v)_1 \vee x(v)_2) \\
\wedge\ & (\overline{x(u)}_1 \vee \overline{x(v)}_1 \vee x(z)_1 \vee x(z)_2) \\
\wedge\ & (\overline{x(u)}_1 \vee \overline{x(v)}_2 \vee x(z)_1 \vee x(z)_2) \\
\wedge\ & (\overline{x(u)}_2 \vee \overline{x(v)}_1 \vee x(z)_1 \vee x(z)_2) \\
\wedge\ & (\overline{x(u)}_2 \vee \overline{x(v)}_2 \vee x(z)_1 \vee x(z)_2) \\
\wedge\ & \overline{x(z)}_1 \\
\wedge\ & \overline{x(z)}_2
\end{aligned}
$$

Figure 1: The pebbling contradiction $Peb^2_{\Pi_2}$ for the pyramid graph Π_2 of height 2.

- $\overline{x(u)}_i \vee \overline{x(v)}_j \vee \bigvee_{l=1}^{d} x(w)_l$ *for all* $i,j \in [d]$ *and all* $w \in V(G) \setminus S$*, where* u,v *are the two predecessors of* w *(*pebbling axioms*).*

The pebbling contradiction Peb^d_G is a (2+d)-CNF formula with $\mathrm{O}\bigl(d^2 \cdot |V(G)|\bigr)$ clauses over $d \cdot |V(G)|$ variables. See Figure 1 for an example pebbling contradiction.

As shown in [8], Peb^d_G can be refuted in resolution by deriving $\bigvee_{i=1}^{d} x(v)_i$ for all $v \in V(G)$ inductively in topological order and then resolving with the target axioms $\overline{x(z)}_i$, $i \in [d]$. Writing down this resolution proof, it is easy to verify that $L(F \vdash 0) = \mathrm{O}\bigl(d^2 \cdot |V(G)|\bigr)$ and $W(F \vdash 0) \leq \mathrm{O}(d)$.

As to refutation space, it is not too difficult to see that $Sp\bigl(Peb^d_G \vdash 0\bigr) \leq \textsf{\textit{Peb}}(G) + \mathrm{O}(1)$ using an optimal black pebbling of G together with the resolution refutation from [8] sketched above. For binary trees, [20] improved this bound to $Sp\bigl(Peb^2_{T_h} \vdash 0\bigr) \leq \bigl\lceil \frac{2h+1}{3} \bigr\rceil + 3 = \frac{2}{3}\textsf{\textit{Peb}}(T_h) + \mathrm{O}(1)$. No lower bounds on $Sp\bigl(Peb^d_G \vdash 0\bigr)$ have been shown, however, and for one variable per vertex, we know from [6] that $Sp\bigl(Peb^1_G \vdash 0\bigr) = \mathrm{O}(1)$. At the same time, the resolution refutation in [20] is structurally quite similar to the optimal black-white pebbling of T_h presented in [28], and it seems hard to come up with a way how any resolution refutation could do better for $d \geq 2$ variables per vertex. This raises the suspicion that the black-white pebbling price $\textsf{\textit{BW-Peb}}(T_h) = h/2 + \mathrm{O}(1)$ might be a lower bound for $Sp\bigl(Peb^d_{T_h} \vdash 0\bigr)$, and in general that it might hold that $Sp\bigl(Peb^d_G \vdash 0\bigr) \geq \textsf{\textit{BW-Peb}}(G)$ for any DAG G and $d \geq 2$.

To prove such a bound, we would like to establish a connection between resolution refutations of Peb^d_G and pebblings of the underlying graph G. Let us say that a vertex v is "true" if $\bigvee_{i=1}^{d} x(v)_i$ has been derived and "false" if $\overline{x(v)}_i$ has been derived for all $i \in [d]$. Any resolution proof refutes a pebbling contradiction by deriving that some vertex v is both true and false, and then resolves to get 0. If we let w be any vertex with predecessors u, v, we see that if we have derived that u and v are true, by downloading $\overline{x(u)}_i \vee \overline{x(v)}_j \vee \bigvee_{l=1}^{d} x(w)_l$ for all $i,j \in [d]$ we can derive $\bigvee_{l=1}^{d} x(w)_l$. This appears analogous to the rule that if u and v are black-pebbled we can place a black pebble on w. In the opposite direction, if we know $\overline{x(w)}_l$ for all $l \in [d]$, using the axioms $\overline{x(u)}_i \vee \overline{x(v)}_j \vee \bigvee_{l=1}^{d} x(w)_l$ we can derive that either u or v is false. This looks similar to eliminating a white pebble on w by placing white pebbles on the predecessors u, v and then removing the pebble from w. Generalizing this loose, intuitive reasoning, we argue that a set of black-pebbled vertices V should correspond to the derived conjunction of truth of all $v \in V$, and that a set of white-pebbled vertices W should correspond to the derived disjunction of falsity of one of the $w \in W$.

Suppose that we could show that as the resolution derivation proceeds, the black and white pebbles corresponding to different clause configurations as outlined above move about on the vertices of G in accordance with the rules of the pebble game. If so, we would get that there is some clause configuration $\mathbb{C}$ corresponding to a lot of pebbles. This could in turn hopefully yield a non-constant lower bound for the refutation space. For if $\mathbb{C}$ induces N black pebbles, i.e., implies N disjoint clauses, it seems likely that $|\mathbb{C}|$ should be linear in N. And if $\mathbb{C}$ induces N white pebbles, $|\mathbb{C}|$ should grow with N if $d \geq 2$, since $\mathbb{C}$ has to force d literals false simultaneously for one out of N vertices.[2] This is the guiding intuition behind the result proven in this paper.

5. A MODIFIED PEBBLE GAME

To prove a lower bound on the refutation space of pebbling contradictions, we want to interpret resolution derivations in terms of pebble placements on the corresponding graph. The problem is that it turns out to be hard to make the pebble configurations induced by a resolution derivation obey the rules of the black-white pebble game. Instead, we alter the pebbling rules.

In this section, we present the modified pebble game used for analyzing resolution derivations. We then argue that for binary trees, we get essentially the same bound on pebbling price in this new pebble game as in the black-white pebble game of Definition 4.

Our first modification of the pebble game is to change the rule for white pebble removal so that a white pebble can be removed from a vertex when a black pebble is placed on that same vertex. This does not really change anything.

Our second, and far more substantial, modification of the pebble game is motivated by the fact that a resolution derivation might choose to erase "the wrong clause" from

[2] Note, though, that for $d = 1$ a single clause $\overline{x(v_1)}_1 \vee \overline{x(v_2)}_1 \vee \ldots \vee \overline{x(v_N)}_1$ can imply the falsity of one of an arbitrary number of vertices $v_1, v_2, \ldots, v_N$, which, incidentally, is one way of explaining why $Sp\bigl(Peb^1_G \vdash 0\bigr) = \mathrm{O}(1)$.

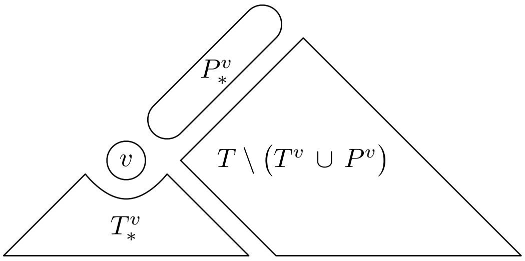

Figure 2: Referencing sets of vertices of a tree T relative to a vertex $v \in V(T)$.

the point of view of our induced pebble configurations. This can lead to pebble configurations reverting to weaker configurations, or to large chunks of black and white pebbles just disappearing. In order to prove lower bounds for a pebble game allowing for such moves, we have to keep track of exactly which white pebbles have been used to get a black pebble on a vertex. Loosely put, removing a white pebble from a vertex v without placing a black pebble on the same vertex should be in order, provided that all black pebbles placed on vertices above v with the help of this white pebble are removed as well.

To define this modified pebble game formally, we need some notation and terminology. We use z to denote the unique target vertex of a complete binary tree T, i.e., the root, p, q, r, u, v, w to denote arbitrary vertices, and U, V, W to denote arbitrary subsets of vertices. We let $succ(v)$ denote the immediate successor of v and $pred(v)$ the immediate predecessors. For a leaf v we have $pred(v) = \emptyset$, and for the root z we have $succ(z) = \emptyset$. We let T^v denote the vertices in the complete binary subtree of T rooted at v and P^v the vertices in the path from v to the root z of T. We set $T_*^v = T^v \setminus \{v\}$ and $P_*^v = P^v \setminus \{v\}$ (see Figure 2). The set W is *below* the vertex u if $W \subseteq T_*^u$. If $P_*^w \cap W = \emptyset$ for all $w \in W$, we say that the vertex set W is *simple*.

We now define the concept used to "label" each black pebble with the set of white pebbles (if any) this black pebble is dependent on. The intuition behind the next definition is that $v\langle W\rangle$ should denote a black pebble on v together with the white pebbles W below v with the help of which we have been able to place the black pebble on v.

Definition 6 (Pebble subconfiguration). *For v a vertex and $W \subseteq T_*^v$ a simple set below v, we say that $v\langle W\rangle$ is a* pebble subconfiguration *with a black pebble on v supported by white pebbles on $w \in W$. The black pebble on v in $v\langle W\rangle$ is said to be* dependent *on the white pebbles in W. We refer to $v\langle\emptyset\rangle$ as an* independent black pebble.

We write $v\langle V\rangle \preceq u\langle U\rangle$ if $T^v \setminus \bigcup_{w\in V} T^w \subseteq T^u \setminus \bigcup_{w\in U} T^w$. If $v\langle V\rangle \preceq u\langle U\rangle$ and $v\langle V\rangle \neq u\langle U\rangle$, we write $v\langle V\rangle \prec u\langle U\rangle$.

We use $\mathbb{L}$ to denote a set of pebble subconfigurations and refer to such a set as a labelled pebble configuration *or an* L-configuration.

Definition 6 is illustrated in Figure 3. Note that $\preceq$ is an order relation and that the minimal elements with respect to $\preceq$ are pebble subconfigurations $v\langle pred(v)\rangle$.

Our modified pebble game is defined in terms of moves not of individual pebbles, but of entire pebble subconfigurations.

Definition 7 (Labelled pebble game). *For T a binary tree with root z, a* labelled black-white pebbling, *or*

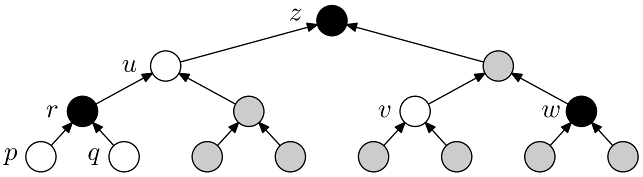

Figure 3: The pebble subconfigurations $z\langle u, v\rangle$, $r\langle p, q\rangle$ and $w\langle\emptyset\rangle$. Note that $w\langle\emptyset\rangle \prec z\langle u, v\rangle$.

L-pebbling, *of T is a sequence of labelled pebble configurations $\mathcal{L} = \{\mathbb{L}_0, \ldots, \mathbb{L}_\tau\}$ such that $\mathbb{L}_0 = \emptyset$ and $\mathbb{L}_t$ is obtained from $\mathbb{L}_{t-1}$ by one of the following rules:*

Introduction $\mathbb{L}_t = \mathbb{L}_{t-1} \cup \{v\langle pred(v)\rangle\}$ *for* $v\langle pred(v)\rangle \notin \mathbb{L}_{t-1}$.

Merger $\mathbb{L}_t = \mathbb{L}_{t-1} \cup \{v\langle(V \cup W) \setminus \{w\}\rangle\}$ *for* $v\langle V\rangle, w\langle W\rangle \in \mathbb{L}_{t-1}$ *with* $w \in V$. *We write* $u\langle U\rangle = \mathsf{merge}(v\langle V\rangle, w\langle W\rangle)$ *for* $u\langle U\rangle = v\langle(V \cup W) \setminus \{w\}\rangle$ *and refer to this as a* merger *on* w.

Reversal $\mathbb{L}_t = \mathbb{L}_{t-1} \cup \{v\langle V\rangle\}$ *if* $v\langle V\rangle \prec u\langle U\rangle$ *for some* $u\langle U\rangle \in \mathbb{L}_{t-1}$.

Erasure $\mathbb{L}_t = \mathbb{L}_{t-1} \setminus \{v\langle V\rangle\}$ *for* $v\langle V\rangle \in \mathbb{L}_{t-1}$.

A legal L-pebbling *of a binary tree T is an L-pebbling $\mathcal{L}$ ending in $\mathbb{L}_\tau = \{z\langle\emptyset\rangle\}$.*

Let $Bl(\mathbb{L}_t) = \{v \mid \exists v\langle W\rangle \in \mathbb{L}_t\}$ denote the black pebbles and $Wh(\mathbb{L}_t) = \{w \mid \exists v\langle W\rangle \in \mathbb{L}_t \text{ s.t. } w \in W\}$ the white pebbles in $\mathbb{L}_t$. Then the cost of an L-configuration $\mathbb{L}$ is $\mathsf{cost}(\mathbb{L}) = |Bl(\mathbb{L}) \cup Wh(\mathbb{L})|$ and the cost of an L-pebbling $\mathcal{L}$ is $\mathsf{cost}(\mathcal{L}) = \max_{t\in[\tau]}\{\mathsf{cost}(\mathbb{L}_t)\}$. The L-pebbling price *of a subconfiguration $v\langle W\rangle$, denoted $\textsf{L-Peb}(v\langle W\rangle)$, is the minimum cost of any L-pebbling such that $\mathbb{L}_\tau = \{v\langle W\rangle\}$, and the L-pebbling price of T is $\textsf{L-Peb}(T) = \textsf{L-Peb}(z\langle\emptyset\rangle)$.*

In the L-pebble game, one can remove a white pebble without placing a black pebble on the same vertex, but if so the rule for erasure makes sure that any black pebble dependent on the removed white pebble is removed as well. A "traditional" removal of a white pebble from w corresponds to merging $v\langle V\rangle$ and $w\langle W\rangle$ into $v\langle(V \cup W) \setminus \{w\}\rangle$ and then erasing $v\langle V\rangle$ and $w\langle W\rangle$. The "backward" pebbling moves to weaker pebble configurations mentioned in the beginning of this section are moves according to the reversal rule.

The L-pebble game without reversal moves is essentially just a disguised version of the ordinary black-white pebble game. Arguing very informally, it seems plausible that making reversals in an L-pebbling should only "weaken" the pebble configurations (for example, reversing from $z\langle u, v\rangle$ to $w\langle\emptyset\rangle$ in Figure 3), and that it should therefore be possible to eliminate all reversal moves from a pebbling without affecting the pebbling cost. Unfortunately, this intuition does not hold in general. If the L-pebble game is extended from trees to arbitrary DAGs in the natural way, the above statement is false. For instance, it fails for the pyramids Π_h of Figure 1. Klawe [24] proved that $\textsf{BW-Peb}(\Pi_h) = h/2 + \mathrm{O}(1)$, but it can be shown that Π_h can be L-pebbled with $\mathrm{O}(1)$ pebbles if we allow reversal moves of black pebbles downwards. For binary trees, however, we can prove that the L-pebbling price and the black-white pebbling price coincide asymptotically.

Theorem 8. *For T_h a complete binary tree of height h, $\textsf{L-Peb}(T_h) = \Theta(\textsf{BW-Peb}(T_h)) = \Theta(h)$.*

The technically quite complicated proof of this fact, which is a cornerstone of our result, can be found in Sections 6 and 7 of [29].

6. RESOLUTION DERIVATIONS INDUCE LABELLED PEBBLINGS

The next step in our proof is to show that sets of clauses can be interpreted in terms of pebble configurations in such a way that resolution derivations induce legal L-pebblings.

For simplicity, from now on we write $v_1, \ldots, v_d$ instead of $x(v)_1, \ldots, x(v)_d$ for the d variables associated with v in a dth degree pebbling contradiction. We write $*\mathit{Peb}^d_G = \mathit{Peb}^d_G \setminus \bigl\{\overline{z}_1, \ldots, \overline{z}_d\bigr\}$ to denote the pebbling contradiction over G with target axioms removed. If $\mathit{pred}(r) = \bigl\{p, q\bigr\}$, the *axioms for* r is the set $\mathit{Ax}^d(r) = \bigl\{\overline{p}_i \vee \overline{q}_j \vee \bigvee_{l=1}^d r_l \mid i, j \in [d]\bigr\}$, and for r a source we let $\mathit{Ax}^d(r) = \bigl\{\bigvee_{i=1}^d r_i\bigr\}$. For V a set of vertices, we define $\mathit{Ax}^d(V) = \bigl\{\mathit{Ax}^d(v) \mid v \in V\bigr\}$.

Our first observation is that instead of refutations of Peb^d_G we may study derivations of $\bigvee_{i=1}^d z_i$ from $*\mathit{Peb}^d_G$.

Lemma 9. $\mathit{Sp}\bigl(\mathit{Peb}^d_G \vdash 0\bigr) = \mathit{Sp}\bigl(*\mathit{Peb}^d_G \vdash \bigvee_{i=1}^d z_i\bigr)$.

Proof. For any derivation $\pi^* : *\mathit{Peb}^d_G \vdash \bigvee_{l=1}^d z_l$, we can get a refutation of Peb^d_G from π^* in the same space by resolving $\bigvee_{l=1}^d z_l$ with all $\overline{z}_l$, $l \in [d]$, in space 3. In the other direction, for $\pi : \mathit{Peb}^d_G \vdash 0$ we can extract a derivation of $\bigvee_{l=1}^d z_l$ in at most the same space by simply omitting all downloads of and resolution steps on $\overline{z}_l$ in π, leaving the literals z_l in the clauses. Instead of the final empty clause 0 we get some clause $D \subseteq \bigvee_{l=1}^d z_l$, and since $*\mathit{Peb}^d_T \nvDash D \subsetneqq \bigvee_{l=1}^d z_l$ and resolution is sound, we have $D = \bigvee_{l=1}^d z_l$. □

We now try to provide some intuition for how clause configurations are translated into pebble configurations. In the black-white pebble game, if at some time t there is an independent black pebble on v, an optimal pebbling will not place any pebbles on T^v after time t. As an analogy, if $\mathbb{C}_t \vDash \bigvee_{i=1}^d v_i$ it is easy to see that no axioms from $\mathit{Ax}^d(T^v)$ need be used after time t to derive $\bigvee_{i=1}^d z_i$, so we let $\mathbb{C}_t$ induce a subconfiguration $v\langle\emptyset\rangle$ in this case. One way of looking at a dependent black pebble on v supported by white pebbles on W, i.e., a subconfiguration $v\langle W\rangle$, is that given independent black pebbles on all $w \in W$, the white pebbles could be eliminated to yield $v\langle\emptyset\rangle$. By analogy, $\mathbb{C}_t$ should induce $v\langle W\rangle$ if $\mathbb{C}_t \cup \bigl\{\bigvee_{i=1}^d w_i \mid w \in W\bigr\} \vDash \bigvee_{i=1}^d v_i$. As an example, we would like the clause configuration

$$\bigl\{\overline{u}_i \vee \overline{v}_j \vee \textstyle\bigvee_{l=1}^d z_l,\ \overline{p}_i \vee \overline{q}_j \vee \bigvee_{l=1}^d r_l,\ \bigvee_{l=1}^d w_l \mid i, j \in [d]\bigr\}$$

to induce the pebble subconfigurations $z\langle u, v\rangle$, $r\langle p, q\rangle$ and $w\langle\emptyset\rangle$ of Figure 3.

Since our resolution derivations have no reason to be this well-behaved, we need to add a number of technical details in the formal definitions. In the following, $\mathbb{B}\bigl(V\bigr)$ can be thought of as "truth of all vertices in V" and A_V as "truth of some vertex in V".

Definition 10. *Let* $\mathbb{B}\bigl(V\bigr) = \bigl\{\bigvee_{i=1}^d v_i \mid v \in V\bigr\}$ *and* $A_V = \bigvee_{v \in V} \bigvee_{i=1}^d v_i$. *Given a set of clauses* $\mathbb{C}$ *and a vertex* v, *if* $V \subseteq T \setminus P^v$ *is such that* $\mathbb{C} \cup \mathbb{B}\bigl(V\bigr) \vDash A_{P^v}$ *we say that* V *is a* support *for* v *with respect to* $\mathbb{C}$. *If there is no* $V' \subsetneqq V$ *such that* $\mathbb{C} \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^v}$ *the support is* minimal. *For* V *a support for* v *with respect to* $\mathbb{C}$ *such that* $\mathbb{C} \cup \mathbb{B}\bigl(V\bigr) \nvDash A_{P^v_*}$, *we say that* v *is* maximal *with respect to* $\mathbb{C}$ *and* V. *We define the* supporting white pebbles *in the set* V *of the vertex* v *to be* $\mathit{swp}(v, V) = \bigl\{w \in V \cap T^v_* \mid P^w_* \cap V = \emptyset\bigr\}$.

Definition 11 (Induced L-configuration). *The* induced L-configuration $\mathbb{L}(\mathbb{C})$ *of a set of clauses* $\mathbb{C}$ *consists of all subconfigurations* $v\langle V\rangle$ *such that*

1. *there is a minimal support* $V' \subseteq T \setminus P^v$ *for* v *with respect to* $\mathbb{C}$,
2. v *is maximal with respect to* $\mathbb{C}$ *and* V',
3. $V = \mathit{swp}(v, V')$.

Suppose that $\pi = \bigl\{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\bigr\}$ is a resolution derivation of $\bigvee_{i=1}^d z_i$ from $*\mathit{Peb}^d_T$. For $\mathbb{C}_0 = \emptyset$ we have $\mathbb{L}(\mathbb{C}_0) = \emptyset$, and $\mathbb{C}_\tau = \bigl\{\bigvee_{i=1}^d z_i\bigr\}$ induces a single independent black pebble $\mathbb{L}(\mathbb{C}_\tau) = \bigl\{z\langle\emptyset\rangle\bigr\}$ on the root of T. Hence, we are done if we can show that $\bigl\{\mathbb{L}(\mathbb{C}_0), \ldots \mathbb{L}(\mathbb{C}_\tau)\bigr\}$ is (essentially) a legal L-pebbling.

The rest of this section is devoted to proving this. We start by stating three technical lemmas. The first lemma relates subset containment of supporting sets and the order relation between corresponding subconfigurations.

Lemma 12. *For a vertex* $v \in V(T)$, *if* $u \in P^v$ *is a vertex and* $U', V' \subseteq T \setminus P^v$ *are vertex sets such that* $U' \cap T^v_* \subseteq V' \cap T^v_*$, *then* $u\langle \mathit{swp}(u, U')\rangle \succeq v\langle \mathit{swp}(v, V')\rangle$.

Proof. This follows from Definitions 6 and 10 by noting that $T^v \setminus \bigcup_{w \in V'} T^w \subseteq T^u \setminus \bigcup_{w \in U'} T^w$. □

A second handy lemma is that if V' is not minimal or v maximal with respect to $\mathbb{C}$, this just means that $\mathbb{C}$ induces something stronger than $v\langle \mathit{swp}(v, V')\rangle$.

Lemma 13. *If* $\mathbb{C} \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^v}$ *for* $V' \subseteq T \setminus P^v$, *then there is an induced subconfiguration* $u\langle U\rangle \in \mathbb{L}(\mathbb{C})$ *such that* $v\langle \mathit{swp}(v, V')\rangle \preceq u\langle U\rangle$.

Proof. Minimize $U' \subseteq V'$ and then pick $u \in P^v$ of maximal height so that $\mathbb{C} \cup \mathbb{B}\bigl(U'\bigr) \vDash A_{P^u}$. Set $U = \mathit{swp}(u, U')$ and use Lemma 12. □

The following easy lemma will be used repeatedly. We omit the proof.

Lemma 14. *Suppose that* C, D *are clauses and* $\mathbb{C}$ *is a set of clauses. Then* $\mathbb{C} \cup \bigl\{C\bigr\} \vDash D$ *if and only if* $\mathbb{C} \vDash \overline{a} \vee D$ *for all* $a \in \mathit{Lit}(C)$.

Using these lemmas, we can prove that resolution derivations induce labelled pebblings. By the L-pebbling rules in Definition 7, any subconfiguration $v\langle V\rangle$ may be erased freely at any time. Consequently, we need not worry about subconfigurations disappearing during the transition from $\mathbb{C}_t$ to $\mathbb{C}_{t+1}$. What we do need to check, though, is that no $v\langle V\rangle$ appears inexplicably in $\mathbb{L}(\mathbb{C}_{t+1})$ as a result of a derivation step $\mathbb{C}_t \rightsquigarrow \mathbb{C}_{t+1}$, but that we can always derive any $v\langle V\rangle \in \mathbb{L}(\mathbb{C}_{t+1}) \setminus \mathbb{L}(\mathbb{C}_t)$ from $\mathbb{L}(\mathbb{C}_t)$ by the L-pebbling rules.

Theorem 15. *Let* $\pi = \bigl\{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\bigr\}$ *be any resolution derivation of* $\bigvee_{i=1}^d z_i$ *from* $*\mathit{Peb}^d_T$. *Then the induced L-configurations* $\bigl\{\mathbb{L}(\mathbb{C}_0), \ldots, \mathbb{L}(\mathbb{C}_\tau)\bigr\}$ *form the "backbone" of a legal L-pebbling* $\mathcal{L}$ *of* T *in the following sense: all transitions* $\mathbb{L}(\mathbb{C}_t) \rightsquigarrow \mathbb{L}(\mathbb{C}_{t+1})$ *can be accomplished in accordance with the L-pebbling rules in such a way that the cost of* $\mathcal{L}$ *is bounded by* $\mathsf{cost}(\mathcal{L}) = \mathrm{O}\bigl(\max_{t \in [\tau]}\bigl\{\mathsf{cost}(\mathbb{L}(\mathbb{C}_t))\bigr\}\bigr)$.

PROOF. Induction over the derivation by case analysis of the derivation step $\mathbb{C}_t \rightsquigarrow \mathbb{C}_{t+1}$.

If $\mathbb{C}_{t+1}$ is derived from $\mathbb{C}_t$ by inference, as our definition of induced pebbles is semantical it is immediate that $\mathbb{L}(\mathbb{C}_{t+1}) = \mathbb{L}(\mathbb{C}_t)$ and there is nothing to prove.

By Lemma 13, all new induced subconfigurations resulting from an erasure $\mathbb{C}_{t+1} = \mathbb{C}_t \setminus \{C\}$ can be obtained from $\mathbb{L}(\mathbb{C}_t)$ by reversal moves.

The case of axiom download $\mathbb{C}_{t+1} = \mathbb{C}_t \cup \{C\}$ is more complicated. Let us fix a vertex $v \in V(T)$ and an axiom $C \in Ax^d(r)$. If $v\langle V\rangle$ is a pebble subconfiguration induced at time $t+1$, by assumption there is a minimal $V' \subseteq T \setminus P^v$ with $V = swp(v, V')$ such that $\mathbb{C}_t \cup \{C\} \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^v}$.

Our intuition is that C should not yield any interesting new subconfigurations $v\langle V\rangle$ if $r \in T \setminus T^v$, and for $r \in T^v$ we should be able to explain new subconfigurations with the help of an introduction of $r\langle pred(r)\rangle$ in our L-pebbling. We prove this by a case analysis over r.

$r \in T \setminus \bigl(T^v \cup P^v\bigr)$: Observing that $\mathbb{B}\bigl(r\bigr) \vDash C$ (this will be used repeatedly), we get that $\mathbb{C}_t \cup \mathbb{B}\bigl(V' \cup \{r\}\bigr) \vDash A_{P^v}$ for $V' \cup \{r\} \subseteq T \setminus P^v$. Lemma 13 tells us that there is a $u\langle U\rangle \in \mathbb{L}(\mathbb{C}_t)$ such that $v\langle V\rangle = v\langle swp(v, V')\rangle = v\langle swp(v, V' \cup \{r\})\rangle \preceq u\langle U\rangle$, so we can get $v\langle V\rangle$ from $\mathbb{L}(\mathbb{C}_t)$ by a reversal move.

$r \in P^v_*$: Write $C = \overline{p}_i \vee \overline{q}_j \vee \bigvee_{l=1}^{d} r_l$ for $\{p, q\} = pred(r) \neq \emptyset$ and assume without loss of generality that p is the vertex in $P^v \cap pred(r)$. Using Lemma 14 to move p_i to the right of the implication sign yields $\mathbb{C}_t \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^v} \vee p_i = A_{P^v}$, and since V' is minimal it follows that $v\langle V\rangle \in \mathbb{L}(\mathbb{C}_t)$.

$r = v$: Note first that we are prepared to accept the introduction of $r\langle pred(r)\rangle$ as a legal pebbling move, so if $\mathbb{C}_t \cup \{C\} \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^r}$ for $pred(r) \subseteq V'$ no further analysis is needed for $r\langle swp(r, V')\rangle = r\langle pred(r)\rangle$. In particular, this is always the case if $pred(r) = \emptyset$, i.e., if r is a source.

Suppose that $v\langle V\rangle = r\langle swp(r, V')\rangle \in \mathbb{L}(\mathbb{C}_{t+1})$ for $V \neq pred(r) = \{p, q\}$, and write $C = \overline{p}_i \vee \overline{q}_j \vee \bigvee_{l=1}^{d} r_l$. We want to derive $r\langle V\rangle$ by the pebbling rules from $\mathbb{L}(\mathbb{C}_t) \cup \bigl\{r\langle pred(r)\rangle\bigr\}$. By symmetry, we get two subcases.

1. $p \in V, q \notin V$: By Definition 10, we have $p \in V'$ and $q \notin V'$. Observe that this implies that $V' \subseteq T \setminus P^q$. Also, we can use Lemma 14 to move q_j to the right-hand side of the implication sign and get $\mathbb{C}_t \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^r} \vee q_j \subseteq A_{P^r} \vee \bigvee_{j=1}^{d} q_j = A_{P^q}$. Plugging this into Lemma 13 shows that there is a $w\langle W\rangle \in \mathbb{L}(\mathbb{C}_t)$ such that $q\bigl\langle V \setminus \{p\}\bigr\rangle = q\bigl\langle swp(q, V')\bigr\rangle \preceq w\langle W\rangle$. Thus we can derive $q\bigl\langle V \setminus \{p\}\bigr\rangle$ from $\mathbb{L}(\mathbb{C}_t)$ by reversal and then merge $r\langle pred(r)\rangle = r\langle p, q\rangle$ with $q\bigl\langle V \setminus \{p\}\bigr\rangle$ to obtain $r\bigl\langle\bigl(\{p, q\} \cup (V \setminus \{p\})\bigr) \setminus \{q\}\bigr\rangle = r\langle V\rangle$.
2. $p, q \notin V$: Again, by Definition 10 we have $p, q \notin V'$. If we use Lemma 14 twice we get $\mathbb{C}_t \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^p} \wedge A_{P^q}$, and noting that $V' \subseteq T \setminus \bigl(P^p \cup P^q\bigr)$ we can apply Lemma 13 to derive $p\bigl\langle V \cap T^p_*\bigr\rangle$ and $q\bigl\langle V \cap T^q_*\bigr\rangle$ from $\mathbb{L}(\mathbb{C}_t)$ by reversals. Merging these subconfigurations with $r\langle p, q\rangle$, we get the desired subconfiguration $r\bigl\langle\bigl(V \cap T^p_*\bigr) \cup \bigl(V \cap T^q_*\bigr)\bigr\rangle = r\langle V\rangle$.

$r \in T^v_*$: By assumption, $\mathbb{C}_t \cup \{C\} \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^v}$, and since $r \in T^v_*$ and $\mathbb{B}\bigl(r\bigr) \vDash C$ we have $\mathbb{C}_t \cup \mathbb{B}\bigl(V' \cup \{r\}\bigr) \vDash A_{P^v}$ for $V' \cup \{r\} \subseteq T \setminus P^v$. If $P^r \cap V' \neq \emptyset$, it holds that $swp(v, V' \cup \{r\}) = swp(v, V')$ and we can obtain $v\langle V\rangle$ from $\mathbb{L}(\mathbb{C}_t)$ by reversal according to Lemma 13, so suppose $P^r \cap V' = \emptyset$.

Pick $U' \subseteq V' \cup \{r\}$ minimal and then $u \in P^v$ maximal with respect to U' such that $\mathbb{C}_t \cup \mathbb{B}\bigl(U'\bigr) \vDash A_{P^u}$. By the minimality of V' we have $r \in U'$, and since $P^r_* \cap U' \subseteq P^r_* \cap V' = \emptyset$ it holds that $r \in swp(u, U')$. Consequently, we cannot use $u\langle U\rangle = u\bigl\langle swp(u, U')\bigr\rangle \in \mathbb{L}(\mathbb{C}_t)$ to derive $v\langle V\rangle \not\preceq u\langle U\rangle$ by reversal. However, since $U' \subseteq V' \cup \{r\}$, Lemma 12 says that $v\bigl\langle(V \cup \{r\})\setminus T^r_*\bigr\rangle = v\bigl\langle swp(v, V' \cup \{r\})\bigr\rangle \preceq u\langle U\rangle$ can be derived by reversal from $\mathbb{L}(\mathbb{C}_t)$. If we could also derive $r\bigl\langle V \cap T^r_*\bigr\rangle$ from $\mathbb{L}(\mathbb{C}_t) \cup \bigl\{r\langle pred(r)\rangle\bigr\}$, we could do a merger to get $v\bigl\langle\bigl(\bigl((V \cup \{r\}) \setminus T^r_*\bigr) \cup \bigl(V \cap T^r_*\bigr)\bigr) \setminus \{r\}\bigr\rangle = v\langle V\rangle$.

Hence, we are done if we can derive the subconfiguration $r\bigl\langle V \cap T^r_*\bigr\rangle = r\bigl\langle swp(v, V') \cap T^r_*\bigr\rangle = r\bigl\langle swp(r, V')\bigr\rangle$ from $\mathbb{L}(\mathbb{C}_t) \cup \bigl\{r\langle pred(r)\rangle\bigr\}$. But $A_{P^r} \supseteq A_{P^v}$, so by assumption we have $\mathbb{C}_t \cup \{C\} \cup \mathbb{B}\bigl(V'\bigr) \vDash A_{P^r}$ for $V' \subseteq T \setminus P^r$. This is almost exactly the case $r = v$ above, where we proved that $r\bigl\langle swp(r, V')\bigr\rangle$ is derivable from $\mathbb{L}(\mathbb{C}_t) \cup \bigl\{r\langle pred(r)\rangle\bigr\}$. The only difference is that now it is not necessarily true that V' is a minimal support and that r is maximal with respect to V'. But these assumptions were not used in the derivation of $r\bigl\langle swp(r, V')\bigr\rangle$ from $\mathbb{L}(\mathbb{C}_t) \cup \bigl\{r\langle pred(r)\rangle\bigr\}$ anyway, so we can reuse exactly the same proof to get $r\bigl\langle swp(r, V')\bigr\rangle$. This concludes the analysis for $r \in T^v_*$.

Studying the pebbling moves in the case analysis above, we see that all subconfigurations $v\langle V\rangle \in \mathbb{L}(\mathbb{C}_{t+1}) \setminus \mathbb{L}(\mathbb{C}_t)$ resulting from an axiom download can be obtained from $\mathbb{L}(\mathbb{C}_t) \cup r\langle pred(r)\rangle$ by a (possibly empty) sequence of reversals from $\mathbb{L}(\mathbb{C}_t)$, followed by a (possibly empty) sequence of mergers on $\{r\} \cup pred(r)$.

Finally, since a single resolution derivation step $\mathbb{C}_t \rightsquigarrow \mathbb{C}_{t+1}$ can induce several consecutive L-pebbling moves to get from $\mathbb{L}(\mathbb{C}_t)$ to $\mathbb{L}(\mathbb{C}_{t+1})$, we have to take care of the possibility that the maximal pebbling cost in $\mathcal{L}$ might be reached in some intermediate L-configuration $\mathbb{L}'$ in between $\mathbb{L}(\mathbb{C}_t)$ and $\mathbb{L}(\mathbb{C}_{t+1})$. This is why we get only an asymptotic bound $\boldsymbol{cost}(\mathcal{L}) = \mathrm{O}\bigl(\max_{t \in [\tau]}\bigl\{\boldsymbol{cost}(\mathbb{L}(\mathbb{C}_t))\bigr\}\bigr)$. We refer to Section 8 of [29] for the omitted technical details. □

7. A SEPARATION OF SPACE AND WIDTH

The final component needed to piece together the proof of our lower bound on the refutation space of pebbling contradictions is that the number of pebbles in an induced L-configuration $\mathbb{L}(\mathbb{C})$ and the number of of clauses in $\mathbb{C}$ are somehow connected. Note that we have to assume $d \geq 2$ here, since $Sp\bigl({*Peb}^1_G \vdash z_1\bigr) = Sp\bigl(Peb^1_G \vdash 0\bigr) = \mathrm{O}(1)$.

We say that a set of clauses $\mathbb{C}$ implies a clause D *minimally* if $\mathbb{C} \vDash D$ but for all $\mathbb{C}' \subsetneq \mathbb{C}$ it holds that $\mathbb{C}' \nvDash D$. We have the following two easy lemmas (the proofs of which can be found in Section 3 of [29]).

LEMMA 16. *Let $\mathbb{C}$ be a set of clauses and D a clause such that $\mathbb{C} \vDash D$ minimally. Suppose that it holds that $a \in Lit(\mathbb{C})$ but $\overline{a} \notin Lit(\mathbb{C})$. Then we must have $a \in Lit(D)$.*

Lemma 17. *Suppose for a set of clauses $\mathbb{C}$ and clauses D_1 and D_2 with $\mathit{Vars}(D_1) \cap \mathit{Vars}(D_2) = \emptyset$ that $\mathbb{C} \vDash D_1 \lor D_2$ but $\mathbb{C} \nvDash D_2$. Then there is a literal $a \in \mathit{Lit}(\mathbb{C}) \cap \mathit{Lit}(D_1)$.*

With every resolution derivation $\pi : F \vdash D$, we can associate a DAG G_π with the clauses in π labelling the vertices and with edges from the assumption clauses to the resolvent for each inference. There might be several different occurrences of a clause C in π, but if so we can label each occurrence of C with a timestamp when it was derived or downloaded and keep track of which copy of C is used where.

For a derivation $\pi = \{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\}$ of D from F, we say that the clause $C \in \mathbb{C}_t$ is *superfluous* if there is no path in G_π from the vertex for $C \in \mathbb{C}_t$ to the vertex for D. A derivation is *non-superfluous* if it contains no superfluous clauses. For every derivation $\pi : F \vdash D$, we can get a corresponding non-superfluous derivation $\pi' : F \vdash D$ in at most the same length, width and space by only considering the vertices in G_π from which D is reachable.

The clause $C \in F$ is *used* in a resolution derivation π if C occurs in π. The derivation $\pi : F \vdash D$ *uses exactly* $F' \subseteq F$ if π uses all clauses $C' \in F'$ but no clauses $C \in F \setminus F'$. We say that a CNF formula F *syntactically precisely yields* the clause D if there is a non-superfluous resolution derivation $\pi : F \vdash D$ that uses exactly F, and denote this by $F \vDash_{\triangledown} D$.

To obtain the bound on clause set size in terms of the number of induced pebbles, we first generalize a result in [4] (again we refer to Section 3 of [29] for a proof).

Theorem 18. *Suppose that F is a CNF formula and D a clause such that $F \vDash_{\triangledown} D$, and for $V \subseteq \mathit{Vars}(F)$ let $F_V = \{C \in F \mid \mathit{Vars}(C) \cap V \neq \emptyset\}$. Then if $V \subseteq \mathit{Vars}(F) \setminus \mathit{Vars}(D)$, it holds that $|F_V| > |V|$. In particular, if $F \vDash D$ minimally then $|F_V| > |V|$ for all $V \subseteq \mathit{Vars}(F) \setminus \mathit{Vars}(D)$.*

Let us define $\mathit{Vars}^d(u) = \{u_1, \ldots, u_d\}$. We say that a vertex u is *represented* in a clause C derived from $^*\!\mathit{Peb}^d_T$, or that C *mentions* u, if $\mathit{Vars}^d(u) \cap \mathit{Vars}(C) \neq \emptyset$. We write

$$V(C) = \{u \in V(T) \mid \mathit{Vars}^d(u) \cap \mathit{Vars}(C) \neq \emptyset\}$$

to denote all vertices represented in C and

$$\mathbb{C}[U] = \{C \in \mathbb{C} \mid V(C) \cap U \neq \emptyset\}$$

to denote the subset of all clauses in $\mathbb{C}$ mentioning vertices in U. Using Theorem 18 and induction over subsets of induced pebbles $U \subseteq \mathit{Bl}(\mathbb{L}(\mathbb{C})) \cup \mathit{Wh}(\mathbb{L}(\mathbb{C}))$, we get that $|\mathbb{C}[U]| \geq |U|$ which yields the desired bound.

Theorem 19. *Suppose that $\mathbb{C}$ is a set of clauses derived from $^*\!\mathit{Peb}^d_T$ for $d \geq 2$ and that $V \subseteq V(T)$ is a set of vertices such that $\mathbb{C}$ induces a black or white pebble on each $v \in V$. Then $|\mathbb{C}| \geq |V|$.*

Proof. Suppose that $\mathbb{C}$ induces a subconfiguration $v\langle W \rangle$. By Definition 11, there is a minimal support $V_v \subseteq T \setminus P^v$ with $W = \mathit{swp}(v, V_v) \subseteq V_v$ such that $\mathbb{C} \cup \mathbb{B}(V_v) \vDash A_{P^v}$ but $\mathbb{C} \cup \mathbb{B}(V_v) \nvDash A_{P^v_*}$ and $\mathbb{C} \cup \mathbb{B}(V'_v) \nvDash A_{P^v}$ for all $V'_v \subsetneqq V_v$. Fix for each $v\langle W \rangle$ a subset $\mathbb{C}_v \subseteq \mathbb{C}$ such that $\mathbb{C}_v \cup \mathbb{B}(V_v) \vDash A_{P^v}$ minimally. Since by definition $\mathit{Vars}(\mathbb{B}(V_v)) \cap \mathit{Vars}(A_{P^v}) = \emptyset$, using Lemma 17 with $D_1 = \bigvee_{i=1}^d v_i$ and $D_2 = A_{P^v_*}$ we see that the vertex v must be represented in $\mathbb{C}_v$ by some positive literal v_i. For the white pebbles in $W \subseteq V_v$, it follows for the same reason from Lemma 16 that all literals $\overline{w}_j$, $j \in [d]$ must be present in $\mathbb{C}_v$.

We prove by induction over $U \subseteq V$ that $|\mathbb{C}[U]| \geq |U|$, from which the theorem clearly follows. The base case $|U| = 1$ is immediate, since we just proved that all pebbled vertices $v \in V$ are represented in $\mathbb{C}$.

For the induction step, suppose that $|\mathbb{C}[U']| \geq |U'|$ for all $U' \subsetneqq U$. Pick a "topmost" vertex $u \in U$, i.e., such that $P^u_* \cap U = \emptyset$, and look at the subconfiguration $v\langle W \rangle$ containing u (with $v = u$ if u is black) and the associated subset $\mathbb{C}_v \subseteq \mathbb{C}$. Note that $\mathit{Vars}^d(U) \cap \mathit{Vars}(A_{P^v}) \subseteq \{u\}$. Let $S = U \cap V(\mathbb{C}_v)$ be the set of all vertices in U mentioned by $\mathbb{C}_v$. We claim that $|\mathbb{C}_v[S]| \geq |S|$.

To show this, note first that it was proven above that $u \in S$, and if $S = \{u\}$ we trivially have $|\mathbb{C}_v[S]| \geq 1 = |S|$. Suppose therefore that $S \supsetneqq \{u\}$. We want to apply Theorem 18 on the formula $F = \mathbb{C}_v \cup \mathbb{B}(V_v)$. Let $S' = S \setminus \{u\}$, write $S' = S_1 \dot\cup S_2$ for $S_1 = S' \cap V_v$ and $S_2 = S' \setminus S_1$, and consider

$$\begin{aligned} F_{S'} &= \{C \in (\mathbb{C}_v \cup \mathbb{B}(V_v)) \mid V(C) \cap S' \neq \emptyset\} \\ &= \mathbb{C}_v[S'] \cup \mathbb{B}(S_1). \end{aligned}$$

For each $w \in S_1$, the clauses in $\mathbb{B}(S_1)$ contain d literals $w_1, \ldots, w_d$, and these literals must all occur negated in $\mathbb{C}_v$ by Lemma 16. For each $w \in S_2$, the clauses in $\mathbb{C}_v[S']$ contain at least one variable w_i. Appealing to Theorem 18 with the subset of variables $\mathit{Vars}^d(S') \cap \mathit{Vars}(\mathbb{C}_v)$, we get

$$\begin{aligned} |F_{S'}| &= |\mathbb{C}_v[S'] \cup \mathbb{B}(S_1)| \\ &> |\mathit{Vars}^d(S') \cap \mathit{Vars}(\mathbb{C}_v)| \\ &\geq d|S_1| + |S_2|, \end{aligned}$$

and rewriting this as

$$\begin{aligned} |\mathbb{C}_v[S]| &\geq |\mathbb{C}_v[S']| \\ &= |F_{S'}| - |\mathbb{B}(S_1)| \\ &\geq (d-1)|S_1| + |S_2| + 1 \\ &\geq |S| \end{aligned}$$

proves the claim (this is where we need $d \geq 2$).

Note that $\mathbb{C}_v[S] \subseteq \mathbb{C}[U]$, since $\mathbb{C}_v \subseteq \mathbb{C}$ and $S \subseteq U$. Also, by construction $\mathbb{C}_v[S]$ does not mention any vertices in $U \setminus S$ since $S = U \cap V(\mathbb{C}_v)$. Thus, $\mathbb{C}[U \setminus S] \subseteq \mathbb{C}[U] \setminus \mathbb{C}_v[S]$, and using the induction hypothesis for $U \setminus S \subsetneqq U$ we get

$$|\mathbb{C}[U]| \geq |\mathbb{C}_v[S]| + |\mathbb{C}[U \setminus S]| \geq |S| + |U \setminus S| = |U|.$$

The theorem follows by induction. □

We can now prove a tight bound for the refutation space of pebbling contradictions over binary trees.

(Restated) Theorem 1. *Let T_h denote the complete binary tree of height h and $\mathit{Peb}^d_{T_h}$ the pebbling contradiction of degree $d \geq 2$ defined over T_h. Then the space of refuting $\mathit{Peb}^d_{T_h}$ in resolution is $\mathit{Sp}(\mathit{Peb}^d_{T_h} \vdash 0) = \Theta(h)$.*

Proof. The upper bound $\mathit{Sp}(\mathit{Peb}^d_G \vdash 0) \leq \textsf{\textit{Peb}}(G) + \mathrm{O}(1)$ for any DAG G is fairly obvious: given an optimal black pebbling of G, derive $\bigvee_{i=1}^d v_i$ inductively from $\mathit{Ax}^d(v)$ and $\{\bigvee_{i=1}^d u_i \mid u \in \mathit{pred}(v)\}$ when vertex v is pebbled. With a little care, this can be done in space independent of d. Hence $\mathit{Sp}(\mathit{Peb}^d_{T_h} \vdash 0) = \mathrm{O}(\textsf{\textit{Peb}}(T_h)) = \mathrm{O}(h)$.

For the lower bound, let $\pi = \{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\}$ be a resolution derivation of $\bigvee_{i=1}^d z_i$ from $^*\!\mathit{Peb}^d_{T_h}$ in minimal space.

Combining Lemma 9 with Theorems 8, 15 and 19, we get that $Sp\bigl(Peb^d_{T_h} \vdash 0\bigr) = Sp\bigl({*}Peb^d_{T_h} \vdash \bigvee_{i=1}^{d} z_i\bigr) = Sp(\pi) = \max_{t\in[\tau]}\{|\mathbb{C}_t|\} \geq \max_{t\in[\tau]}\{\mathsf{cost}(\mathbb{L}(\mathbb{C}_t))\} = \Omega\bigl(\mathsf{L\text{-}Peb}(T_h)\bigr) = \Omega(h)$. □

Since $W\bigl(Peb^d_G \vdash 0\bigr) = \mathrm{O}(d)$ by [8], fixing $d \geq 2$ and letting $F_n = Peb^d_{T_h}$ for $h = \lfloor \log(n+1) \rfloor$ in Theorem 1 yields the stated separation of space from width.

(Restated) Corollary 2. *For all $k \geq 4$, there is a family $\{F_n\}_{n=1}^{\infty}$ of k-CNF formulas of size $\mathrm{O}(n)$ such that $W(F_n \vdash 0) = \mathrm{O}(1)$ but $Sp(F_n \vdash 0) = \Theta(\log n)$.*

8. CONCLUSION AND OPEN PROBLEMS

We have proven the first lower bound on refutation space in resolution which is not the consequence of a lower bound on the refutation width for the same formulas, but instead separates the two measures. This answers an open question in [6, 18, 20, 34]. However, we believe that it should be possible to strengthen our answer in several interesting ways.

Firstly, we would like to extend the lower bound on the refutation space of pebbling contradictions over binary trees to the k-DNF resolution proof systems $\mathfrak{Res}(k)$ introduced in [26], where the configurations $\mathbb{C}$ consist of k-DNF formulas instead of disjunctive clauses.

Conjecture 1. *For k-DNF resolution refutations of pebbling contradictions over complete binary trees T_h of height h, fixing k it holds that $Sp_{\mathfrak{Res}(k+1)}\bigl(Peb^{k+1}_{T_h} \vdash 0\bigr) = \mathrm{O}(1)$ but $Sp_{\mathfrak{Res}(k)}\bigl(Peb^{k+1}_{T_h} \vdash 0\bigr) = \Omega(h)$.*

Proving this conjecture would establish that the k-DNF resolution proof systems form a strict hierarchy with respect to space, which would be an improvement of the separation in [18] for the restricted case of tree-like k-DNF resolution.

Secondly, it would be nice to generalize the bound on refutation space of pebbling contradictions to DAGs other than trees that have better size-pebbling price trade-off. Our guess is that the black-white pebbling price is a lower bound for pebbling contradictions over any DAG.

Conjecture 2. *For $d \geq 2$ and for G an arbitrary DAG with a unique target and with all vertices having indegree 0 or 2, $Sp\bigl(Peb^d_G \vdash 0\bigr) = \Omega\bigl(\mathsf{BW\text{-}Peb}(G)\bigr)$.*

Since there are DAGs G_n of fan-in 2 and size $\mathrm{O}(n)$ which have black-white pebbling price $\mathsf{BW\text{-}Peb}(G_n) = \Theta\bigl(n/\log n\bigr)$ (see [22]),[3] a proof of Conjecture 2 would immediately yield the corollary that there is a family of unsatisfiable k-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\mathrm{O}(n)$ such that $W\bigl(F_n \vdash 0\bigr) = \mathrm{O}(1)$ but $Sp\bigl(F_n \vdash 0\bigr) = \Omega\bigl(n/\log n\bigr)$.

A third and final question is whether refutation space can be separated from refutation length in the following sense.

Conjecture 3. *There is a family of unsatisfiable k-CNF formulas $\{F_n\}_{n=1}^{\infty}$ over n variables such that $Sp\bigl(F_n \vdash 0\bigr) = \omega\bigl(\sqrt{n \log L(F_n \vdash 0)}\bigr)$.*

This would be an interesting contrast to the relation between length and width $W\bigl(F \vdash 0\bigr) = \mathrm{O}\bigl(\sqrt{n \log L(F \vdash 0)}\bigr)$ for k-CNF formulas shown in [9]. Of course, if we could prove Conjecture 2, we would immediately get a positive answer to Conjecture 3 as well, using the same formula family as in the "corollary" of Conjecture 2.

However, it is not possible to prove Conjecture 2 by a simple generalization of the L-pebble game with reversal moves in Section 5 to general DAGs G. As was observed in Section 5, if we allow reversal moves of black pebbles downwards it is not true that $\mathsf{L\text{-}Peb}(G) = \Omega\bigl(\mathsf{BW\text{-}Peb}(G)\bigr)$.

As a first step, we would therefore have to modify Definition 11 so that a set of clauses $\mathbb{C}$ induces a black pebble on v if there is a minimal subset $\mathbb{C}_v \subseteq \mathbb{C}$ such that $\mathbb{C}_v \cup \mathbb{B}\bigl(V\bigr) \vDash A_{P^v}$ but $\mathbb{C}_v \cup \mathbb{B}\bigl(V\bigr) \nvDash A_{P^v_*}$. Otherwise we could move black pebbles downwards through erasures simply by deriving $\mathbb{C} = \bigl\{\bigvee_{i=1}^{d} v_i,\ \bigvee_{i=1}^{d} succ(v)_i\bigr\}$ and then deleting $\bigvee_{i=1}^{d} succ(v)_i$.

But if we define induced pebbles in terms of minimal subsets $\mathbb{C}_v \subseteq \mathbb{C}$, as a result black pebbles can slide downwards after inference steps, since $\{B \vee C\}$ is weaker than $\{B \vee x, C \vee \overline{x}\}$. This problem can be solved by defining induced pebbles syntactically instead of semantically. Recalling the definitions in the beginning of Section 7, we could say that $\mathbb{C}$ induces a black pebble on v if there is a $\mathbb{C}_v \subseteq \mathbb{C}$ with v represented positively in $\mathbb{C}_v$ and a clause $D_v \subseteq A_{P^v}$ such that $\mathbb{C}_v \cup \mathbb{B}\bigl(V\bigr) \vDash_{\forall} D_v$. With this definition, nothing bad happens during inference or erasure steps, resolution derivations yield legal L-pebblings, and Theorem 18 can be used to show the bound in Theorem 1 for binary trees.

However, because of the fact that the support $\mathbb{B}(w) = \bigvee_{l=1}^{d} w_l$ for a non-leaf w is stronger than the set of axioms $Ax^d(w) = \bigl\{\overline{u}_i \vee \overline{v}_j \vee \bigvee_{l=1}^{d} w_l \mid i, j \in [d]\bigr\}$, we can still get black pebbles moving downwards at axiom download. This can be avoided by defining support in terms of $Ax^d(w)$ instead of $\bigvee_{i=1}^{d} w_i$, which leads to a quite well-behaved pebble game, but then unfortunately the counting argument in Theorem 18 to get a bound on $|\mathbb{C}|$ in terms of the number of induced pebbles breaks down.

These problems arise because we do not a priori have any restrictions on what kind of clauses a resolution derivation from a pebbling contradiction might derive. The counterexample derivations we have found for the definitions sketched in the previous paragraph all seem clearly non-optimal, while all of the definitions yield well-behaved pebblings for "normal" resolution derivations. One way of solving the problems would be if one could define formally what constitutes a "non-optimal" derivation from a pebbling contradiction and then show that each "non-optimal" derivation can be replaced by an "optimal" one in asymptotically the same space. Alternatively, one could try to find new ideas for the connection between the black-white pebble game and resolution derivations from pebbling contradictions, or use the last definition for induced pebbles outlined above but devise new methods for proving bounds on $|\mathbb{C}|$ in terms of the number of induced pebbles.

9. ACKNOWLEDGMENTS

I am grateful to my advisor Johan Håstad for patiently giving feed-back during this work. Also, I would like to thank Douglas Wikström for interesting discussions about theoretical computer science in general and pebble games in particular. Finally, a special thanks to Joel Brynielsson, who introduced me to the mysteries of METAPOST.

[3] Note that in several papers, this result is incorrectly attributed to [27], but [27] itself gives the correct reference.

10. 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] M. Alekhnovich, J. Johannsen, T. Pitassi, and A. Urquhart. An exponential separation between regular and general resolution. In *Proceedings 34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 448–456, May 2002.

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

[4] S. Baumer, J. L. Esteban, and J. Torán. Minimally unsatisfiable CNF formulas. *Bulletin of the European Association for Theoretical Computer Science*, 74:190–192, June 2001.

[5] 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.

[6] E. Ben-Sasson. Size space tradeoffs for resolution. In *Proceedings 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, R. Impagliazzo, and A. Wigderson. Near optimal separation of treelike and general resolution. *Combinatorica*, 24(4):585–603, Sept. 2004.

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

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

[11] 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.

[12] M. L. Bonet and N. Galesi. A study of proof search algorithms for resolution and polynomial calculus. In *Proceedings 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS '99)*, pages 422–431, Oct. 1999.

[13] J. Buresh-Oppenheim and T. Pitassi. The complexity of resolution refinements. In *Proceedings 18th IEEE Symposium on Logic in Computer Science (LICS 03)*, pages 138–147, June 2003.

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

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

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

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

[18] 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.

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

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

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

[22] J. R. Gilbert and R. E. Tarjan. Variations of a pebble game on graphs. Technical Report STAN-CS-78-661, Stanford University, 1978. Available at `http://www-db.stanford.edu/TR/CS-TR-78-661.html`.

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

[24] M. M. Klawe. A tight bound for black and white pebbles on the pyramid. *Journal of the ACM*, 32(1):218–228, Jan. 1985.

[25] H. Kleine Büning and T. Letterman. *Propositional Logic: Deduction and Algorithms.* Cambridge University Press, 1999.

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

[27] T. Lengauer and R. E. Tarjan. Upper and lower bounds on time-space tradeoffs. In *Proceedings 11th Annual ACM Symposium on Theory of Computing (STOC '79)*, pages 262–277, May 1979.

[28] T. Lengauer and R. E. Tarjan. The space complexity of pebble games on trees. *Information Processing Letters*, 10(4/5):184–188, July 1980.

[29] J. Nordström. Narrow proofs may be spacious: Separating space and width in resolution. Technical Report TR05-066, Revision 02, Electronic Colloquium on Computational Complexity (ECCC), Nov. 2005. Available at `http://www.eccc.uni-trier.de/eccc/`.

[30] N. Pippenger. Pebbling. Technical Report RC8258, IBM Watson Research Center, 1980.

[31] R. Raz. Resolution lower bounds for the weak pigeonhole principle. *Journal of the ACM*, 51(2):115–138, 2004.

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

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

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

[35] 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.

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