Electronic Colloquium on Computational Complexity, Comment 1 on Report No. 66 (2005)

A Tight Bound on the Refutation Clause Space of Pebbling Contradictions over Binary Trees

Jakob Nordström

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

September 13, 2005

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 memory cells used if the proof is only allowed to resolve on clauses kept in memory. Both of these measures have previously been studied and related to the refutation size of unsatisfiable CNF formulas. Also, the resolution refutation space of a formula has been proven to be at least as large as the refutation width. In ECCC Technical Report TR05-066, we proved that space can be separated from width, answering an open question in several previous papers. In these notes we sharpen this result to a tight bound on the refutation space of pebbling contradictions over binary trees.

This is a strenghening of the analysis in Section 8 of [5]. We refer to [5] for any non-standard notation or terminology used in these notes.

1 A Tight Bound on the Refutation Space of $Peb^d_{T_h}$

In Section 7 of [5] we proved that $Sp\left(Peb^d_{T_h} \vdash 0\right) = Sp\left({}^*Peb^d_{T_h} \vdash \bigvee_{l=1}^{d} z_l\right)$, and that each resolution derivation $\pi : {}^*Peb^d_{T_h} \to \bigvee_{i=1}^{d} z_i$ induces a legal L-pebbling $\mathcal{L}$ of T_h such that $\max_{\mathbb{C}\in\pi}\left\{\mathsf{cost}(\mathbb{L}(\mathbb{C}))\right\} = \Omega\left(\mathsf{cost}(\mathcal{L})\right)$. From Sections 5 and 6 of [5] we know that $\mathsf{cost}(\mathcal{L}) = \Omega\left(\mathsf{BW\text{-}Peb}(T)\right)$. The final component needed to piece together the proof of our lower bound on the refutation space of pebbling contradictions is to show 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.

We cannot expect a proof of this fact to work regardless of the pebbling degree d. The induced L-pebbling in Section 7 of [5] makes no assumptions about d, but we know that $Sp\left({}^*Peb^1_G \vdash z_1\right) = Sp\left(Peb^1_G \vdash 0\right) = \mathrm{O}(1)$. If we look at the resolution refutation π of Peb^1_G in constant space sketched at the end of Section 4 of [5], we see that the induced L-pebbling starts by placing white pebbles on $pred(z)$ and a black pebble on z, i.e., introducing $z\langle pred(z)\rangle$, and then pushes the white pebbles downwards by introducing $v\langle pred(v)\rangle$ for all v

ISSN 1433-8092

in reverse topological order and merging until it reaches $z\langle S\rangle$ for S the source vertices of G. Finally, the white pebbles $s \in S$ are eliminated one by one by introducing $s\langle\emptyset\rangle$ and merging. The reason that Peb_G^1 can be refuted in constant space is that one single clause $\bigvee_{v\in V} \overline{v}_1 \vee z_1$ can induce an arbitrary number $|V|$ of white pebbles, or, phrasing it differently, that white pebbles are free for $d = 1$.

Below, we prove a lower bound $|\mathbb{C}| \geq N$ for N induced pebbles in Theorem 1.12. As we just observed, we will need $d \geq 2$ if some of these N pebbles are white. Black pebbles are *not* free for $d = 1$, however, but instead of showing a separate bound for them we assume $d \geq 2$ and give a simple, unified proof for N simultaneous black or white pebbles. We conclude the section by combining the bound in Theorem 1.12 with previous theorems from [5] to obtain a tight bound on the refutation clause space of pebbling contradictions over binary trees in Theorem 1.13 and a separation of space and width in Corollary 1.14.

In the proofs, we will use the following definitions.

Definition 1.1. We say that a vertex v is *represented positively* in a clause C if $\{v_1, \ldots, v_d\} \cap \mathit{Lit}(C) \neq \emptyset$ and *negatively* if $\{\overline{v}_1, \ldots, \overline{v}_d\} \cap \mathit{Lit}(C) \neq \emptyset$, and that C *mentions* v *positively* or *negatively*, respectively. This definition is extended to sets of vertices and clauses by taking unions.

For a set of vertices U, we let $\mathit{Vars}(U) = \{u_1, \ldots, u_d \mid u \in U\}$ denote the set of all variables representing vertices in U. For a set of clauses $\mathbb{C}$, we use $V(\mathbb{C}) = \{u \in U \mid \mathit{Vars}(u) \cap \mathit{Vars}(\mathbb{C}) \neq \emptyset\}$ to denote all vertices represented (positively or negatively) in $\mathbb{C}$, and we write $\mathbb{C}[U] = \{C \in \mathbb{C} \mid V(C) \cap U \neq \emptyset\}$ to denote the subset of all clauses in $\mathbb{C}$ mentioning vertices in U.

Definition 1.2. For v a vertex in T and α a truth value assignment, v is said to be *true under* α if $\alpha\left(\bigvee_{i=1}^{d} v_i\right) = 1$ and *false under* α if $\alpha\left(\bigvee_{i=1}^{d} v_i\right) = 0$. We define

$$\alpha^{v=\nu}(u_i) = \begin{cases} \alpha(u_i) & \text{if } u \neq v, \\ \nu & \text{if } u = v \end{cases}$$

and say that $\alpha^{v=0}$ *flips* v *to false*.

Definition 1.3. A *restriction* ρ is a partial truth value assignment. We represent a restriction as the set of literals $\rho = \{a_1, \ldots, a_m\}$ set to true by ρ. For a clause C, the *ρ-restriction* of C is

$$C|_\rho = \begin{cases} 1 & \text{if } \rho \cap \mathit{Lit}(C) \neq \emptyset, \\ C \setminus \{\overline{a} \mid a \in \rho\} & \text{otherwise,} \end{cases}$$

where 1 denotes the trivially true clause, and the ρ-restriction $\mathbb{C}|_\rho$ of a set of clauses $\mathbb{C}$ is the union of the ρ-restrictions $C|_\rho \neq 1$ for $C \in \mathbb{C}$.

We write $\rho(\neg C) = \{\overline{a} \mid a \in \mathit{Lit}(C)\}$ to denote the unique minimal restriction that falsifies C.

Definition 1.4. We say that a set of clauses $\mathbb{C}$ implies a clause D *minimally* if $\mathbb{C} \vDash D$ but for all $\mathbb{C}' \subsetneqq \mathbb{C}$ it holds that $\mathbb{C}' \nvDash D$. If $\mathbb{C} \vDash 0$ minimally, we say that $\mathbb{C}$ is *minimally unsatisfiable*.

We prove a pair of technical lemmas about minimally implicating clause sets.

Lemma 1.5. *Suppose for $\mathbb{C}$ a set of clauses and D a clause that $\mathbb{C} \vDash D$ minimally, and let $\rho = \rho(\neg D)$. Then $\mathbb{C}|_\rho$ is minimally unsatisfiable. Also, it holds that $\left|\mathbb{C}|_\rho\right| = \left|\mathbb{C}\right|$, i.e., no literal $a \in Lit(D)$ occurs negated in $\mathbb{C}$.*

Proof. Pick any $\mathbb{C}' \subseteq \mathbb{C}$ such that $\mathbb{C}'|_\rho$ is minimally unsatisfiable. If there was a truth value assignment α such that $\alpha\big(\mathbb{C}'\big) = 1$ and $\alpha\big(D\big) = 0$, this α would satisfy $\mathbb{C}'|_\rho$, which is contrary to assumption. Hence $\mathbb{C}' \vDash D$, and again by assumption we must have $\mathbb{C}' = \mathbb{C}$. This also shows that $\left|\mathbb{C}|_\rho\right| = \left|\mathbb{C}\right|$, for if ρ satisfied some clause $C \in \mathbb{C}$ this would imply that $(\mathbb{C} \setminus \{C\})|_\rho$ was minimally unsatisfiable for $\mathbb{C} \setminus \{C\} = \mathbb{C}' \subsetneqq \mathbb{C}$. $\square$

Lemma 1.6. *Suppose for $\mathbb{C}$ a set of clauses and D a clause that $\mathbb{C} \vDash D$ minimally and that $a \in Lit(\mathbb{C})$ but $\overline{a} \notin Lit(\mathbb{C})$. Then $a \in Lit(D)$.*

Proof. Suppose not. Let $\mathbb{C}_1 = \big\{C \in \mathbb{C} \mid a \in Lit(C)\big\}$ and $\mathbb{C}_2 = \mathbb{C} \setminus \mathbb{C}_1$. Since $\mathbb{C}_2 \nvDash D$ there is an α such that $\alpha(\mathbb{C}_2) = 1$ and $\alpha(D) = 0$. Note that $\alpha(a) = 0$, since otherwise $\alpha(\mathbb{C}_1) = 1$. It follows that $\overline{a} \notin Lit(D)$. Flip a to true. By construction $\alpha^{a=1}\big(\mathbb{C}_1\big) = 1$, but $\mathbb{C}_2$ and D are not affected since $\{a, \overline{a}\} \cap \big(Lit(\mathbb{C}_2) \cup Lit(D)\big) = \emptyset$, so $\alpha^{a=1}\big(\mathbb{C}_2\big) = 1$ and $\alpha^{a=1}\big(D\big) = 0$. Contradiction. $\square$

The fact that a minimally unsatisfiable CNF formula F must have more clauses than variables seems to have been proven independently a number of times (see e.g. [1, 2, 4]). We generalize this result to subsets of variables and the clauses containing variables in these subsets.

Lemma 1.7. *For F a minimally unsatisfiable CNF formula and $V \subseteq Vars(F)$ any subset of variables, let $F_V = \{C \in F \mid Vars(C) \cap V \neq \emptyset\}$. Then $|F_V| > |V|$.*

Proof. By induction over $V \subseteq Vars(F)$.

If $|V| = 1$, then $|F_V| \geq 2$, since any $x \in V$ must occur both positively and negatively in F. For suppose x occurs only positively or only negatively. Then because of the minimality of F we can satisfy $F' = F \setminus \{C \in F \mid x \in Vars(C)\}$ with a partial truth value assignment ρ to $Vars(F) \setminus \{x\}$ since x does not occur in F', and then extend ρ to a satisfying assignment for all of F by setting x to the right value. (This is basically the proof of Lemma 1.6.)

The inductive step generalizes this idea. Suppose that $|F_{V'}| > |V'|$ for all strict subsets $V' \subsetneqq V$ and consider V. Since $F_{V'} \subseteq F_V$ if $V' \subseteq V$, choosing any V' of size $|V| - 1$, we see that $|F_V| \geq |F_{V'}| \geq |V'| + 1 = |V|$.

If $|F_V| > |V|$ there is nothing to prove, so assume that $|F_V| = |V|$. Consider the bipartite graph with the variables V and the clauses in F_V as vertices, and edges between variables and clauses for all variable occurrences. Since for all $V' \subseteq V$ the set of neighbours $N(V') = F_{V'} \subseteq F_V$ satisfies $|N(V')| \geq |V'|$, by Hall's Marriage Theorem there is a perfect matching between V and F_V. Use this matching to satisfy F_V assigning values to variables in V only.

The clauses in $F' = F \setminus F_V$ are not affected by this partial truth value assignment, since they do not contain any occurrences of variables in V. Furthermore, by the minimality of F it must hold that F' can be satisfied by assigning values to variables in $Vars(F') \setminus V$.

The two partial truth value assignments for F_V and F' can be combined to a satisfying assignment for all of F, which contradicts the fact that F is unsatisfiable. Thus $|F_V| > |V|$. The lemma follows. $\square$

The next lemma is needed to show that if a clause set $\mathbb{C}$ implies black pebbles on a set of vertices V, then these vertices must be represented positively in $\mathbb{C}$.

Lemma 1.8. *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 $\mathbb{C} \vDash D_1 \vee D_2$ but $\mathbb{C} \nvDash D_2$. Then there is a literal $a \in \mathit{Lit}(\mathbb{C}) \cap \mathit{Lit}(D_1)$.*

Proof. Pick a truth value assignment α such that $\alpha\big(\mathbb{C}\big) = 1$ but $\alpha\big(D_2\big) = 0$. By assumption $\alpha\big(D_1\big) = 1$. Let α' be the same assignment except that all satisfied literals in D_1 are flipped to false. Then $\alpha'\big(D_1 \vee D_2\big) = 0$ forces $\alpha'\big(\mathbb{C}\big) = 0$, so the flip must have falsified some previously satisfied clause in $\mathbb{C}$. □

We also need to show that white-pebbled vertices are represented in $\mathbb{C}$. Now if $\mathbb{C}$ induces a white pebble on a vertex w, it follows immediately by Lemma 1.6 that all literals $\overline{w}_i$, $i \in [d]$, are represented in $\mathit{Lit}(\mathbb{C})$. But we can say something stronger.

Lemma 1.9. *Suppose for a clause set $\mathbb{C}$ and a vertex w that there is a $v \in P^w_*$ and a $V \subseteq T \setminus P^w_*$ such that $\mathbb{C} \cup \mathbb{B}(V) \vDash A_{P^v}$ but $\mathbb{C} \cup \mathbb{B}(V \setminus \{w\}) \nvDash A_{P^v}$. Then there is a subset $\big\{\overline{w}_i \vee C_i \mid i \in [d]\big\} \subseteq \mathbb{C}$ for which $\overline{w}_j \notin \mathit{Lit}(C_i)$ if $j \neq i$.*

Proof. Pick α such that $\alpha\big(\mathbb{C}\big) = \alpha\big(\mathbb{B}(V \setminus \{w\})\big) = 1$ but $\alpha\big(A_{P^v}\big) = 0$. Then it must be the case that $\alpha\big(\bigvee_{i=1}^{d} w_i\big) = 0$. For all $i \in [d]$ we have $\alpha^{w_i=1}\big(\mathbb{B}(V)\big) = 1$ but $\alpha^{w_i=1}\big(A_{P^v}\big) = 0$, so flipping w_i while keeping w_j false for $j \neq i$ must falsify some clause in $\mathbb{C}$. This establishes that there are clauses $\overline{w}_i \vee C_i \in \mathbb{C}$ for all $i \in [d]$ such that $\overline{w}_j \notin \mathit{Lit}(C_i)$ for $j \neq i$. □

Lemma 1.9 tells us that one white pebble costs d clauses. We are convinced that the correct bound for N white pebbles should be dN clauses if $d \geq 2$.

We next prove a couple of lemmas to try to argue why the intuition for a bound dN, or at least $(d-1)N$, is strong. At the same time, the proofs of these lemmas indicate why such a bound appears hard to get. Loosely speaking, the problem seems to be that Lemma 1.7 does not really use any structural information about the CNF formula in question. Since very different formulas can yield the same clauses-variable occurrences bipartite graph, perhaps it should not be not very surprising if the lemma does not always yield optimal bounds.

For N white pebbles in one common subconfiguration, the cost is at least $(d-1)N$ clauses.

Lemma 1.10. *If a clause set $\mathbb{C}$ derived from *Peb^d_T induces a subconfiguration $v\langle W\rangle$ then $|\mathbb{C}| > (d-1)|W|$.*

Proof. Pick $\mathbb{C}_v \subseteq \mathbb{C}$ and $V \subseteq T \setminus P^v$ minimal such that $W = \mathit{swp}(v, V)$ and $\mathbb{C}_v \cup \mathbb{B}(V) \vDash A_{P^v}$. Note that $\mathit{swp}(v, V) \subseteq V$ by definition. For $\rho = \rho(\neg A_{P^v})$, Lemma 1.5 says that $(\mathbb{C}_v \cup \mathbb{B}(V))\big|_\rho = \mathbb{C}_v\big|_\rho \cup \mathbb{B}(V)$ is a minimally unsatisfiable clause set. Since $\mathbb{B}(V)$ contains $d|V|$ variables but only $|V|$ clauses, Lemma 1.7 yields that $|\mathbb{C}| \geq \big|\mathbb{C}_v|_\rho\big| > (d-1)|V| \geq (d-1)|W|$. □

Also, two white pebbles always cost at least $2d-1$ clauses, although here the argument starts to become pretty involved...

Lemma 1.11. *If a clause set $\mathbb{C}$ derived from *Peb^d_T induces two white pebbles on T, then $|\mathbb{C}| \geq 2d-1$.*

Proof. Suppose that $\mathbb{C}$ induces white pebbles on w^1 and w^2.

If w^1 and w^2 are contained in the same subconfiguration we have $|\mathbb{C}| \geq 2d-1$ by Lemma 1.10. Assume for $i = 1, 2$ that the induced subconfigurations are $v^i\langle W^i\rangle$, where $w^i \in W^i$, and let $\mathbb{C}_i \subseteq \mathbb{C}$ and $V^i \subseteq T \setminus P^{v^i}$ be minimal such that $W^i = \mathit{swp}(v^i, V^i)$ and $\mathbb{C}_i \cup \mathbb{B}(V^i) \vDash A_{P^{v^i}}$. If $|V^1| > 1$ or $|V^2| > 1$ we again have $|\mathbb{C}| \geq 2d-1$ by the proof of Lemma 1.10, so suppose that $V^i = W^i = \{w^i\}$.

Let $\rho_i = \rho(\neg A_{P^{v^i}})$ for $i = 1, 2$. By the proof of Lemma 1.9 we know that $\mathbb{C}_i|_{\rho_i}$ contains d clauses $\{D^i_j = \overline{w}^i_j \vee C^i_j \mid j \in [d]\}$ for which $\overline{w}^i_k \notin \mathit{Lit}(C^i_j)$ if $k \neq j$. Let us refer to the literals $\overline{w}^i_j \in \mathit{Lit}(D^i_j)$ for $i = 1, 2$ and $j = 1, \ldots, d$ as *critical occurrences*. If $w^1 \in P^{v^2}$ we are done since ρ_2 kills all d clauses $\overline{w}^1_j \vee C^1_j$ and there are still d clauses $\overline{w}^2_j \vee C^2_j$ left in $\mathbb{C}_2|_{\rho_2}$, and the same holds for w^2 and P^{v^1} by symmetry. Assume therefore that $w^1 \notin P^{v^2}$ and $w^2 \notin P^{v^1}$.

Now if $|\mathbb{C}_1 \cup \mathbb{C}_2| < 2d$, Lemma 1.9 combined with the pigeonhole principle tells us that there is some negative literal, say $\overline{w}^1_1$, which occurs critically in a clause containing a literal $\overline{w}^2_j$. Consider the subset $\mathbb{C}_1[\{w^1, w^2\}]$ of clauses in $\mathbb{C}_1$ mentioning w^1 or w^2, and let $m = \mathit{Vars}(\mathbb{C}_1[\{w^1, w^2\}]) \cap \mathit{Vars}(w^2)$. We know that $\overline{w}^1_1$ occurs critically in $\mathbb{C}_1$ together with some $\overline{w}^2_j$, and that all literals from w^2 in $\mathbb{C}_1$ are present in $\mathbb{C}_1|_{\rho_1}$ as well, since $w^2 \notin P^{v^1}$ by assumption and ρ_1 does not satisfy any clauses in $\mathbb{C}_1$ by Lemma 1.5. Thus $m \geq 1$. By Lemma 1.7 we get $|\mathbb{C}_1[\{w^1, w^2\}]|_{\rho_1} \cup \mathbb{B}(w^1)| > d + m$, that is, $|\mathbb{C}_1[\{w^1, w^2\}]| \geq d + m$.

Since $\mathbb{C}_1[\{w^1, w^2\}] \subseteq \mathbb{C}_1 \cup \mathbb{C}_2$ and $|\mathbb{C}_1 \cup \mathbb{C}_2| < 2d$ we must have $m < d$. Consequently, there are $d - m \geq 1$ variables from w^2, say $w^2_1, \ldots, w^2_{d-m}$, that are not mentioned in $\mathbb{C}_1[\{w^1, w^2\}]$. But all negative literals $\overline{w}^i_j$ for $i = 1, 2$ and $j = 1, \ldots, d$ occur in $\mathbb{C}_1 \cup \mathbb{C}_2$, so the literals $\overline{w}^2_1, \ldots, \overline{w}^2_{d-m}$ can all be found in $(\mathbb{C}_1 \cup \mathbb{C}_2) \setminus \mathbb{C}_1[\{w^1, w^2\}]$. However,

$$\begin{aligned}
\big|(\mathbb{C}_1 \cup \mathbb{C}_2) \setminus \mathbb{C}_1[\{w^1, w^2\}]\big| &= |\mathbb{C}_1 \cup \mathbb{C}_2| - \big|\mathbb{C}_1[\{w^1, w^2\}]\big| \\
&\leq (2d - 1) - (d + m) \\
&= d - (m + 1),
\end{aligned}$$

which contradicts the existence of $d-m$ *distinct* clauses $\{\overline{w}^2_j \vee C^2_j \mid j \in [d - m]\}$ guaranteed by Lemma 1.9. Hence $|\mathbb{C}_1 \cup \mathbb{C}_2| \geq 2d$ and the lemma follows. ☐

We believe that the ideas in the proof of Lemma 1.11 could be pushed further to yield a bound $|\mathbb{C}| \geq (d-1)N$ for N white pebbles. However, to get a simpler proof, and to get a common bound for N simultaneous black *or* white pebbles, we instead opt for the bound $|\mathbb{C}| \geq N$.

Theorem 1.12. *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 $v \in V$, i.e., $V \subseteq \mathit{Bl}(\mathbb{L}(\mathbb{C})) \cup \mathit{Wh}(\mathbb{L}(\mathbb{C}))$. Then $|\mathbb{C}| \geq |V|$.*

Proof. Suppose that $\mathbb{C}$ induces a subconfiguration $v\langle W\rangle$. By Definition 7.6 in [5], there is a minimal support $V_v \subseteq T \setminus P^v$ such that $W = \mathit{swp}(v, V_v) \subseteq V_v$ and $\mathbb{C} \cup \mathbb{B}(V_v) \vDash A_{P^v}$ 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 $V'_v \subsetneqq V_v$.

For the black pebble on v, by Lemma 1.8 we see that v must be represented positively in $\mathbb{C} \cup \mathbb{B}(V_v)$ (write $A_{P^v} = \bigvee_{i=1}^d v_i \vee A_{P^v_*}$), and since $v \notin V_v$ the positive literals from v are found in $\mathbb{C}$. For the white pebbles in W, it follows from

Lemma 1.9 (or even just from Lemma 1.6) that all literals $\{\overline{w}_i \mid w \in W,\, i \in [d]\}$ occur in $\mathbb{C}$.

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 $\left|\mathbb{C}\left[U'\right]\right| \geq \left|U'\right|$ for all $U' \subsetneqq U$. Fix a "topmost" vertex $u \in U$, i.e., such that $P^u_* \cap U = \emptyset$. By definition, the vertex u must be black-pebbled. Choose $\mathbb{C}_u \subseteq \mathbb{C}$ minimal such that $\mathbb{C}_u \cup \mathbb{B}(V_u) \vDash A_{P^u}$. Since $\mathbb{C}_u \cup \mathbb{B}(V_u) \nvDash A_{P^u_*}$ and $V_u \cap P^u = \emptyset$ by definition, the vertex u is represented positively in $\mathbb{C}_u$. Using Lemma 1.5 with the restriction $\rho = \rho(\neg A_{P^u})$, we get that $(\mathbb{C}_u \cup \mathbb{B}(V_u))|_\rho = \mathbb{C}_u|_\rho \cup \mathbb{B}(V_u)$ is minimally unsatisfiable. By the same lemma, all literals over $u' \in U \setminus \{u\}$ in $\mathbb{C}_u$ are present also in $\mathbb{C}_u|_\rho$, since $U \cap P^u = \{u\}$ and ρ does not eliminate any clauses from $\mathbb{C}_u$.

Let $S = U \cap V(\mathbb{C}_u) = \left(U \cap V(\mathbb{C}_u|_\rho)\right) \cup \{u\}$ be the set of all vertices in U mentioned by $\mathbb{C}_u$. We claim that $|\mathbb{C}_u[S]| \geq |S|$.

To show this, note first that it was proven above that $u \in S$, and if $\{u\} = S$ we trivially have $|\mathbb{C}_u[S]| \geq 1 = |S|$. Suppose therefore that $S \supsetneqq \{u\}$. We want to apply Lemma 1.7 on $F = \mathbb{C}_u|_\rho \cup \mathbb{B}(V_u)$. Write $S = S_1 \,\dot{\cup}\, S_2$ for $S_1 = S \cap V_u$ and $S_2 = S \setminus S_1$, and consider $F_S = \left\{C \in \mathbb{C}_u|_\rho \cup \mathbb{B}(V_u)\,;\, V(C) \cap S \neq \emptyset\right\} = \mathbb{C}_u[S \setminus \{u\}]|_\rho \cup \mathbb{B}(S_1)$. For each $v \in S_1$, the clauses in $\mathbb{B}(S_1)$ contain d variables $v_1, \ldots, v_d$, and these variables must all occur negated in $\mathbb{C}_u$ by Lemma 1.6. For each $v \in S_2 \setminus \{u\}$, the clauses in $\mathbb{C}_u[S \setminus \{u\}]|_\rho$ contain at least one variable v_i. Appealing to Lemma 1.7 with $V = \mathit{Vars}(S \setminus \{u\}) \cap \mathit{Vars}(\mathbb{C}_u)$, we get that

$$\begin{aligned}
\left|F_S\right| &= \left|\mathbb{C}_u[S \setminus \{u\}]|_\rho \cup \mathbb{B}(S_1)\right| \\
&> \left|\mathit{Vars}(S \setminus \{u\}) \cap \mathit{Vars}(\mathbb{C}_u)\right| \\
&\geq d\left|S_1\right| + \left|S_2\right| - 1,
\end{aligned}$$

and rewriting this as

$$\left|\mathbb{C}_u[S]\right| \geq \left|\mathbb{C}_u[S \setminus \{u\}]|_\rho\right| = \left|F_S\right| - \left|\mathbb{B}(S_1)\right| \geq (d-1)\left|S_1\right| + \left|S_2\right| \geq \left|S\right|$$

proves the claim.

Note that $\mathbb{C}_u[S] \subseteq \mathbb{C}[U]$, since $\mathbb{C}_u \subseteq \mathbb{C}$ and $S \subseteq U$. Also, by construction $\mathbb{C}_u[S]$ does not mention any vertices in $U \setminus S$. In other words, $\mathbb{C}[U \setminus S] \subseteq \mathbb{C}[U] \setminus \mathbb{C}_u[S]$, and using the induction hypothesis we get

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

The theorem follows by induction. $\square$

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

Theorem 1.13. *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 on T_h. Then the space of refuting $\mathit{Peb}^d_{T_h}$ by resolution is $\mathit{Sp}\left(\mathit{Peb}^d_{T_h} \vdash 0\right) = \Theta(h)$.*

Proof. The upper bound $\mathit{Sp}\left(\mathit{Peb}^d_G \vdash 0\right) = \mathrm{O}\left(\mathsf{Peb}(G)\right)$ for any DAG G is fairly obvious. Given an optimal black pebbling of G, derive $\bigvee_{i=1}^d v_i$ inductively when vertex v is pebbled. With a little care, this can be done in constant extra space independent of d. To see this, suppose for $pred(r) = \{p, q\}$ that a black pebble is placed on r. Then p and q are already black-pebbled, so we have $\bigvee_{i=1}^{d} p_i$ and $\bigvee_{j=1}^{d} q_j$ in memory. It is not hard to verify that $\overline{p}_i \lor \bigvee_{l=1}^{d} r_l$ can be derived in additional space 3 by resolving $\bigvee_{j=1}^{d} q_j$ with $\overline{p}_i \lor \overline{q}_j \lor \bigvee_{l=1}^{d} r_l$ for $j \in [d]$. Resolve $\bigvee_{i=1}^{d} p_i$ with $\overline{p}_1 \lor \bigvee_{l=1}^{d} r_l$ to get $\bigvee_{i=2}^{d} p_i \lor \bigvee_{l=1}^{d} r_l$, and then resolve this clause with $\overline{p}_i \lor \bigvee_{l=1}^{d} r_l$ for $i = 2, \ldots, d$ to get $\bigvee_{l=1}^{d} r_l$ in total extra space 4. Conclude the resolution proof by resolving $\bigvee_{i=1}^{d} z_i$ for the target z with the target axioms $\overline{z}_i$, $i \in [d]$, in space 3. Consequently, $Sp\left(Peb_{T_h}^{d} \vdash 0\right) = \mathrm{O}\big(\mathsf{Peb}(T_h)\big) = \mathrm{O}(h)$.

For the lower bound, according to Observation 7.2 in [5] it holds that $Sp\left(Peb_G^d \vdash 0\right) = Sp\left({}^*Peb_G^d \vdash \bigvee_{i=1}^{d} z_i\right)$. Let $\pi = \left\{\mathbb{C}_0, \ldots, \mathbb{C}_\tau\right\}$ be a resolution derivation of $\bigvee_{i=1}^{d} z_i$ from ${}^*Peb_{T_h}^{d}$ in minimal clause space. Combining Theorems 3.3, 5.12 and 7.13 in [5], we know that the derivation π induces a legal L-pebbling $\mathcal{L}$ of the tree T_h such that there is a clause configuration $\mathbb{C}_t \in \pi$ with $\mathsf{cost}(\mathbb{L}(\mathbb{C}_t)) = \Omega\big(\mathsf{cost}(\mathcal{L})\big) = \Omega\big(\mathsf{BW\text{-}Peb}(T_h)\big) = \Omega\big(h\big)$. Fix such a clause configuration $\mathbb{C}_t$. By Theorem 1.12, $\left|\mathbb{C}_t\right| \geq \mathsf{cost}(\mathbb{L}(\mathbb{C}_t)) = \Omega\big(h\big)$.

It follows that $Sp\left(Peb_{T_h}^{d} \vdash 0\right) = \Theta(h)$ for $d \geq 2$. ☐

Since $W\big(Peb_G^d \vdash 0\big) = \mathrm{O}(d)$ for all pebbling contradictions [3], fixing $d \geq 2$ in Theorem 1.13 yields a separation of clause space from width. Corollary 1.14 follows if we let $F_n = Peb_{T_h}^{2}$ for $h = \lfloor \log(n+1) \rfloor$.

Corollary 1.14. *There is a family* $\left\{F_n\right\}_{n=1}^{\infty}$ *of k-CNF formulas of size* $\mathrm{O}(n)$ *such that* $W\big(F_n \vdash 0\big) = \mathrm{O}(1)$ *but* $Sp\left(F_n \vdash 0\right) = \Theta(\log n)$.

References

[1] Ron Aharoni and Nathan Linial. Minimal non-two-colorable hypergraphs and minimal unsatisfiable formulas. *Journal of Combinatorial Theory*, 43:196–204, 1986.

[2] Sven Baumer, Juan Luis Esteban, and Jacobo Torán. Minimally unsatisfiable CNF formulas. *Bulletin of the European Association for Theoretical Computer Science*, 74:190–192, June 2001.

[3] Eli Ben-Sasson, Russell Impagliazzo, and Avi Wigderson. Near optimal separation of treelike and general resolution. *Combinatorica*, 24(4):585–603, September 2004.

[4] Oliver Kullmann. An application of matroid theory to the SAT problem. In *Proceedings 15th Annual IEEE Conference on Computational Complexity (CCC '00)*, pages 116–124, July 2000.

[5] Jakob Nordström. *Narrow Proofs May Be Spacious: Separating Space and Width in Resolution.* Technical Report TR05-066, Revision 01, Electronic Colloquium on Computational Complexity (ECCC), August 2005. Available at `http://www.eccc.uni-trier.de/eccc/`.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'