

WIDTH AND SIZE OF REGULAR RESOLUTION PROOFS

ALASDAIR URQUHART

University of Toronto
e-mail address: urquhart@cs.toronto.edu

ABSTRACT. This paper discusses the topic of the minimum width of a regular resolution refutation of a set of clauses. The main result shows that there are examples having small regular resolution refutations, for which any regular refutation must contain a large clause. This forms a contrast with corresponding results for general resolution refutations.

INTRODUCTION

Recent results [1, 13] showing near-exponential separations between the size of regular and general refutations of certain sets of clauses also show a separation of general and regular resolution width. That is to say, the examples used in showing the size separation have large regular resolution width, but bounded general resolution width.

This observation suggests that it might be possible to prove results for regular resolution similar to those of Ben-Sasson and Wigderson [5] for tree resolution and general resolution. The main theorem below shows that this hope is bound to be disappointed; it exhibits examples having small regular resolution size, but large regular width.

The first part of the paper gives a characterization of regular resolution width, in the style of Atserias and Dalmau [3]. The second part discusses the relationship between the size and width of regular resolution refutations.

1. RESOLUTION PROOFS AND THEIR WIDTH

A *literal* is a propositional variable x or its negation $\neg x$. A *clause* is a set of literals, interpreted as the disjunction of the set. For clauses containing exactly one positive literal, we use the implication $p_1, \dots, p_k \rightarrow q$ as alternative notation for the clause $\neg p_1 \vee \dots \vee \neg p_k \vee q$. For notational convenience, we shall also allow the case where the positive literal q is replaced by the propositional constant \perp . For any assignment σ , $\sigma(\perp) = 0$, so that the expression “ $p_1, \dots, p_k \rightarrow \perp$ ” is an alternative notation for the purely negative clause $\neg p_1 \vee \dots \vee \neg p_k$.

1998 ACM Subject Classification: F2.2, F4.1.

Key words and phrases: regular resolution proofs, size of proofs, width of proofs.

The author gratefully acknowledges the support of the Natural Sciences and Engineering Research Council of Canada.

If p is a variable, and C a clause, then we say that p has a *positive occurrence* in C if p is in C , and a *negative occurrence* in C if $\neg p$ is in C . In addition, we shall say that \perp has a positive occurrence in the purely negative Horn clause $p_1, \dots, p_k \rightarrow \perp$. If Σ is a set of clauses, and x, y are variables in Σ , or the propositional constant \perp , then we say that there is an *implicational chain from x to y in Σ* if there is a sequence $x = x_0, \dots, x_k = y$ of variables (or constants) and a sequence C_1, \dots, C_k of clauses so that for all i , $0 < i \leq k$, x_{i-1} occurs negatively and x_i positively in C_i .

The *resolution rule* allows us to derive the *resolvent* $C \vee D$ from the clauses $C \vee x$ and $D \vee \neg x$ by *resolving on* the variable x . A *resolution derivation* of a clause C from a set of clauses Σ consists of a sequence of clauses in which each clause is either a clause of Σ , or derived from earlier clauses by resolution, and C is the last clause in the sequence; it is a *refutation* of Σ if C is the empty clause Λ .

The *size* of a resolution proof is the number of occurrences of clauses in the proof – that is to say, the length of the proof considered as a sequence. For a contradictory set of clauses Σ , we write $S(\Sigma)$ for the minimum size of a resolution refutation of Σ . A resolution proof is a *tree-style* proof if every clause in the proof is used at most once as a premiss in a resolution inference. We write $S_T(\Sigma)$ for the minimum size of a tree-style refutation of a contradictory set of clauses Σ . A resolution refutation of a set of clauses Σ is an *input proof* if in every application of the resolution rule in it, at least one premiss of the application is an input clause in Σ . Every input refutation is automatically a tree-style refutation.

An *irregularity* in a resolution proof is a sequence of clauses C_1, \dots, C_k so that C_{i+1} is derived from C_i (that is, C_i is one of the premisses of a resolution inference in which the conclusion is C_{i+1}), and there is a variable that occurs in C_1 and C_k , but not in any intermediate clause C_j , $1 < j < k$. A resolution proof is *regular* if it contains no irregularity.

If V is a set of propositional variables, then an *assignment* is a Boolean function defined on a subset of V , that is, an assignment of $\{0, 1\}$ to some or all of the variables in V . If α is an assignment, then we write $|\alpha|$ for the cardinality of α , the number of variables to which α assigns values. The result of restricting a clause C by setting a literal l is defined as follows. If the literal l occurs in C , then $C[l := 1] = 1$, while $C[l := 0]$ is $C \setminus \{l\}$. If Σ is a set of clauses, and $v \in \{0, 1\}$, then $\Sigma[l := v]$ is the set of clauses $\{C[l := v] : C \in \Sigma\} \setminus \{1\}$.

The *width* of a clause is the number of literals in it. The width $w(\Sigma)$ of a set of clauses is the maximum width of a clause in Σ , while the width of a resolution proof is the maximum width of a clause in it. If Σ is a contradictory set of clauses, then we define the *refutation width* of Σ , written $w(\Sigma \vdash 0)$, to be the minimum width of a resolution refutation of Σ . If \mathcal{F} is a family of resolution proofs, we define a restricted notion of refutation width, the *\mathcal{F} -refutation width*, $w(\Sigma \vdash^{\mathcal{F}} 0)$, to be the minimum width of a refutation of Σ that belongs to \mathcal{F} . In particular, we define the *regular refutation width*, $w(\Sigma \vdash^{\mathcal{R}} 0)$, of a contradictory set of clauses Σ to be the \mathcal{R} -refutation width where \mathcal{R} is the class of all regular resolution proofs.

The notation $\log x$ stands for the base two logarithm of x , and $\log^k x$ for $(\log x)^k$.

2. CHARACTERIZATION OF GENERAL RESOLUTION WIDTH

In this section, we give a proof of a result of Atserias and Dalmau [3] characterizing the width of general resolution refutations. The characterization is in terms of a two player

game, that we shall call the *k-width game*, played by the *Prover* and the *Adversary*.¹ The rules of the game are as follows.

The players are given a set of clauses Σ , on a set V of variables, and an integer parameter $k \geq 0$. The players together construct a succession of assignments to the variables in V . Initially, the assignment is empty. Each round of the game proceeds as follows, starting from a current assignment. First, the Prover queries an unassigned variable, and the Adversary assigns a value to it. Second, the Prover is allowed to delete some of the values of the variables in V from the assignment resulting from the Adversary's reply; the result is the new current assignment.

The Adversary can win in two ways. First, if the current assignment (after deletions) assigns values to more than k variables; second, if an earlier assignment is repeated during the play of the game. The Prover wins if the current assignment falsifies an initial clause in Σ . Clearly every play of the game must eventually terminate with a win for the Prover or for the Adversary (Atserias and Dalmau define their game so that when the Adversary wins, the game can continue infinitely).

Definition 2.1. If Σ is a set of clauses on a set V of variables, then a non-empty family \mathcal{A} of V -assignments is an *extendible k-family* for Σ if it satisfies the following conditions:

- (1) No assignment in \mathcal{A} falsifies a clause in Σ ;
- (2) Each assignment α in \mathcal{A} satisfies the condition $|\alpha| \leq k$;
- (3) If $\alpha \in \mathcal{A}$, and $\beta \subseteq \alpha$, then $\beta \in \mathcal{A}$;
- (4) If $\alpha \in \mathcal{A}$, $|\alpha| < k$, and $x \in V$, then there is a $\beta \in \mathcal{A}$, so that $\alpha \subseteq \beta$, and $\beta(x)$ is defined.

The next theorem shows that a resolution refutation of width k constitutes a winning strategy for the Prover, while an extendible $k + 1$ -family provides a winning strategy for the Adversary.

Theorem 2.2. [Atserias and Dalmau 2003] *Let Σ be a set of clauses, and $k \geq w(\Sigma)$. Then the following are equivalent:*

- (1) *There is no resolution refutation of Σ of width k ;*
- (2) *There is an extendible $k + 1$ -family for Σ ;*
- (3) *The Adversary wins the $k + 1$ -width game based on Σ .*

Proof. First, let us suppose that there is no resolution refutation of Σ of width k . Let \mathcal{C} be the set of all clauses having a resolution proof from Σ of width at most k ; since $w(\Sigma) \leq k$, $\Sigma \subseteq \mathcal{C}$. Let \mathcal{A} be the set of all assignments of size at most $k + 1$ that do not falsify any clause in \mathcal{C} . We claim that \mathcal{A} is an extendible $k + 1$ -family for Σ . First, \mathcal{A} is non-empty, because it contains the empty assignment (since \mathcal{C} does not contain the empty clause). Second, \mathcal{A} satisfies the first three conditions of Definition 2.1, by construction. To prove the fourth condition, let $\alpha \in \mathcal{A}$, and $|\alpha| \leq k$, $x \in V$, but there is no extension β of α in \mathcal{A} with $\beta(x)$ defined. It follows that there is a clause $D \in \mathcal{C}$ that is falsified if we extend α by setting x to 0. Then $D = E \vee x$ for some E , since otherwise α would falsify D . Similarly, there is a clause $F \vee \neg x$ in \mathcal{C} that is falsified by the extension of α that sets x to 1. But then α must falsify $E \vee F$, showing that $E \vee F$ has width at most k , since $|\alpha| \leq k$. Hence, it follows that $E \vee F$ is in \mathcal{C} , contradicting our assumption that α is in \mathcal{A} .

Second, let us suppose that there is an extendible $k + 1$ -family for Σ . Then the Adversary can play the k -width game on Σ by responding to the Prover's queries with the appropriate

¹Atserias and Dalmau, following the tradition of finite model theory, call their players the *Spoiler* and the *Duplicator*, but our terminology seems clearer in the present context.

assignment from the family, starting with the empty assignment. Since no assignment in the family falsifies an initial clause, this strategy must eventually end in a win for the Adversary, no matter how the Prover plays.

Finally, let us suppose that there is a resolution refutation of Σ of width k . Then the refutation provides the Prover with a winning strategy in the $k + 1$ -width game based on Σ . Starting from the empty clause at the root, the Prover follows a path in the refutation so that at each round, the assignment (after appropriate deletions) is a minimal assignment falsifying the current clause. The variable queried is the variable resolved upon to derive the current clause. This strategy must result in a win for the Prover when the path reaches a clause in Σ . \square

3. CHARACTERIZATION OF REGULAR RESOLUTION WIDTH

In the present section, we modify the result of Atserias and Dalmau to characterize the width of regular resolution refutations. The characterization is again in terms of a two player game, that we shall call the *regular k -width game*. The game is exactly the same as that described in the previous section, but with the added condition that the Prover can never query a previously queried variable.

As in the case of general resolution width, we can characterize the regular resolution width in terms of extendible families of assignments. However, we need to redefine the notion of an assignment. In the earlier notion of assignment, a variable could be in three states: *positive* (1), *negative* (0), and *unassigned* (*). For the case of regular resolution, we define an *extended assignment* to be an assignment of values in which each variable can be in four states: *positive* (1), *negative* (0), *unassigned* (*), or *forgotten* (\boxtimes). The *empty extended assignment* to a set V of variables consists of the assignment in which all variables in V are unassigned (*) (this should be distinguished from assignments in which all of the variables are unassigned or forgotten (\boxtimes)).

If α is an extended assignment, then those variables that are assigned the values 0 or 1 are the *live* variables in α , and we write $|\alpha|$ for the number of live variables in α . If α and β are extended assignments to a set of variables V , then we write $\alpha \subseteq \beta$ if β results from α by replacing some unassigned variables by live variables. We also write $\alpha \sqsubseteq \beta$ if α results from β by forgetting some variables, that is, changing the value of a live (0 or 1) variable to \boxtimes .

As in the case of the earlier k -width game, the players are given a set of clauses Σ , on a set V of variables, and an integer parameter $k \geq 0$. Together, they construct a succession of extended assignments to the variables in V . Initially, the assignment is empty. Each round of the game proceeds as follows. First, the Prover queries an unassigned variable, and the Adversary assigns a value to it. Next, the Prover is allowed to forget some of the variables in the resulting assignment, that is, to change the value of a live variable from 0 or 1 to \boxtimes (forgotten); the result is the new current assignment.

Again, the Adversary can win in two ways. First, if the current assignment assigns values to more than k variables; second, if the Prover has not won up to this point, but there are no unqueried variables, so the Prover has no legal move. The Prover wins if the current assignment falsifies an initial clause in Σ (if this clause contains more than k variables, then we count this as a win for the Adversary). As before, every play of the game must eventually terminate with a win for the Prover or for the Adversary.

Definition 3.1. If Σ is a set of clauses on a set V of variables, then a family \mathcal{A} of extended V -assignments is a *regular extendible k -family for Σ* if it satisfies the following conditions:

- (1) The empty assignment belongs to \mathcal{A} ;
- (2) No assignment in \mathcal{A} falsifies a clause in Σ ;
- (3) Each assignment α in \mathcal{A} satisfies the condition $|\alpha| \leq k$;
- (4) If $\alpha \in \mathcal{A}$, and $\beta \sqsubseteq \alpha$, then $\beta \in \mathcal{A}$;
- (5) If $\alpha \in \mathcal{A}$, $|\alpha| < k$, $x \in V$, and $\alpha(x) = *$, then there is a $\beta \in \mathcal{A}$, so that $\alpha \sqsubseteq \beta$, and $\beta(x) = 0$ or $\beta(x) = 1$.

The next theorem is the analogue of Theorem 2.2 for regular resolution.

Theorem 3.2. *Let Σ be a set of clauses, and $k \geq w(\Sigma)$. Then the following are equivalent:*

- (1) *There is no regular resolution refutation of Σ of width k ;*
- (2) *There is a regular extendible $k + 1$ -family for Σ ;*
- (3) *The Adversary wins the regular $k + 1$ -width game based on Σ .*

Proof. (1 \Rightarrow 2): Let us suppose that there is no regular resolution refutation of Σ of width k . Define \mathcal{C} to be the set of all clauses having a regular resolution proof from Σ of width at most k ; since $w(\Sigma) \leq k$, $\Sigma \subseteq \mathcal{C}$. Let \mathcal{A} be the set of all extended assignments of size at most $k + 1$ that do not falsify any clause in \mathcal{C} . We claim that \mathcal{A} is an extendible $k + 1$ -family for Σ .

Since the empty clause Λ does not belong to \mathcal{C} , the empty assignment is in \mathcal{A} , so the first condition in Definition 3.1 is satisfied. The second condition holds because $\Sigma \subseteq \mathcal{C}$, and the third condition by definition. The fourth condition also follows from the definition of \mathcal{A} .

It remains to prove the fifth condition. Assume that $\alpha \in \mathcal{A}$, and $|\alpha| \leq k$, $x \in V$, and $\alpha(x) = *$, but there is no extension $\beta \supseteq \alpha$ in \mathcal{A} with $\beta(x)$ defined. Let α^0 and α^1 be the extended assignments obtained from α by setting x to 0 and 1, respectively. Since neither α^0 nor α^1 belong to \mathcal{A} , it follows that there are regular resolution derivations \mathcal{R}^0 and \mathcal{R}^1 of clauses C_0 and C_1 , each having width at most k , so that for $i = 0, 1$, $\alpha^i(C_i) = 0$. Since $\alpha \in \mathcal{A}$, it follows that $C_0 = D \vee x$, and $C_1 = E \vee \neg x$, for some clauses D and E . However, if we extend the regular resolution derivations \mathcal{R}^0 and \mathcal{R}^1 by resolving on x , so that the final clause is $D \vee E$, the result is a regular resolution derivation of $D \vee E$, where $\alpha(D \vee E) = 0$. Since $|\alpha| \leq k$, $D \vee E$ has at width at most k , showing that $D \vee E$ is in \mathcal{C} ; this contradicts our assumption that $\alpha \in \mathcal{A}$.

(2 \Rightarrow 3): Second, let us suppose that there is a regular extendible $k + 1$ -family for Σ . Then the Adversary can play the k -width game on Σ by responding to the Prover's queries with the appropriate assignment from the family, starting with the empty assignment. Since no assignment in the family falsifies an initial clause, this strategy must eventually end in a win for the Adversary, no matter how the Prover plays.

(3 \Rightarrow 1): Finally, let us suppose that there is a regular resolution refutation of Σ of width k . Then the refutation provides the Prover with a winning strategy in the regular $k + 1$ -width game based on Σ . Starting from the empty clause at the root, the Prover follows a path in the refutation so that at the end of each round, after the Prover has forgotten certain live variables, the remaining live variables are the domain of a minimal assignment falsifying the current clause. The variable queried is the variable resolved upon to derive the current clause. This strategy must result in a win for the Prover. \square

Corollary 3.3. *The question “Is there a regular resolution refutation of the set of clauses Σ with width k ?” is in PSPACE.*

Proof. Theorem 3.2 shows that this question can be answered by an alternating Turing machine operating in polynomial time. \square

In the case of general resolution width, it is not clear whether the corresponding problem is in PSPACE, because there is no polynomial upper bound on how long the k -width game might last.

4. SIZE AND WIDTH OF REGULAR RESOLUTION PROOFS

4.1. The width and size of resolution proofs. Recent results on size separation between regular and general resolution also show a width separation.

Theorem 4.1. *For each $n > 0$, there is a contradictory set of clauses with $O(n^2)$ variables and $O(n^3)$ clauses for which the general resolution width is bounded, but the regular resolution width is $\Omega(n)$.*

Proof. The paper [1] implicitly contains such a separation. More specifically the family of clauses $GT'_{n,\rho}$ defined in §3 of [1] fulfil the conditions of the theorem. The $\Omega(n)$ lower bound on regular resolution width is proved (implicitly) in Theorem 3.10 of that paper, which shows an exponential ($2^{n/200}$) lower bound on the size of regular resolution refutations of $GT'_{n,\rho}$. \square

The author’s paper [13] demonstrates an improved size separation between regular and general resolution; it also shows a width separation between the two forms of proof system. The main theorem shows that for infinitely many n , there is a set Π_n of $O(n \log^5 n)$ clauses containing $O(n \log \log n)$ variables, where the maximum width of a clause in Π_n and the general resolution width are both $O(\log \log n)$, while the regular resolution width is $\Omega(n/\log n)$.

The results just described suggest a natural conjecture that a good lower bound on the regular width of a set of clauses leads to a good lower bound on the size of a regular refutation of them. For both general and tree resolution, Ben-Sasson and Wigderson [5] have proved strong results along these lines.

Theorem 4.2. [Ben-Sasson and Wigderson 2001] *Let Σ be a contradictory set of clauses with an underlying set of variables V . Then:*

- (1) $S_T(\Sigma) \geq 2^{w(\Sigma+0)-w(\Sigma)}$;
- (2) $S(\Sigma) = \exp\left(\Omega\left(\frac{(w(\Sigma+0)-w(\Sigma))^2}{|V|}\right)\right)$

Given the width and size separation results between regular and general resolution cited above, it seems reasonable to conjecture that the second lower bound proved by Ben-Sasson and Wigderson might hold, in the form where we replace “resolution size” by “regular resolution size,” and “resolution width” by “regular resolution width.” In the remainder of the paper, we show that this conjecture fails.

Before proceeding to the main constructions, it may be helpful to the reader to clarify the relations between the various forms of resolution discussed here.

If we consider the size measure alone, then it is not hard to see that regular resolution is at least as powerful as tree resolution. This is because a pruning procedure [11] [12, p. 436] can be applied to a tree refutation to remove any irregularities while decreasing the size of the tree. On the other hand, the lower bound on width for regular refutations proved below does not apply to tree resolution, since the minimum width of a tree-style refutation of a set of clauses is the same as that of a general resolution refutation (we can convert any general resolution proof into a tree-style proof by repeating subderivations).

However, if we insist on restricting our attention to tree-style refutation *of minimum size*, then the lower bounds on width do apply, since such refutations are necessarily regular.

5. PEBBLING GAMES AND PEBBLING FORMULAS

5.1. The pebbling game. A *pointed graph* G is a directed acyclic graph where all vertices have indegree at most two, having a unique sink, or target vertex, to which there is a directed path from all the vertices in G . It is *binary* if all vertices except for the source vertices have indegree two. If v is a vertex in a pointed graph G , then $G|v$ is the subgraph of G restricted to the vertices from which there is a directed path to v .

The *pebbling game* played on a pointed graph G is a one-player game in the course of which pebbles are placed on or removed from vertices in G . The rules of the game are as follows;

- (1) A pebble may be placed on a source vertex at any time.
- (2) If all predecessors of a vertex are marked with pebbles, then a pebble may be placed on the vertex itself.
- (3) A pebble may be removed from a vertex at any time.

A *move* in the game consists of placing or removing one of the pebbles in accordance with one of the three rules. The *configuration* at a given stage in the game is the set of vertices in G that are marked with a pebble. A play of the game begins with no pebbles on G . The goal of the game is to place a pebble on the sink vertex t , while minimizing the number of pebbles used (that is, minimizing the number of pebbles on the graph at any stage of the game). Thus a successful play of the game can be presented as a sequence of configurations C_0, \dots, C_k , where $C_0 = \emptyset$ and $t \in C_k$, where C_{j+1} is obtained from C_j by one of the three rules.

A *strategy* for the game is a sequence of moves following the rules of the game that ends in pebbling the target vertex. The *cost* of such a strategy is the minimum number of pebbles required in order to execute it, that is to say, the size of the largest configuration in the sequence of configurations produced by following the strategy. The *pebbling number* of G , written as $\#G$, is the minimum cost of a strategy for the pebbling game played on G .

5.2. Pebbling formulas. We associate a contradictory set of clauses $\text{Peb}(G)$ with every pointed graph G . Each vertex in G except the target t is assigned a distinct variable; to simplify notation, we identify a vertex with the variable associated with it, and use the notation $\text{Var}(G)$ for the set of these variables. We associate the constant \perp (falsum) with the target vertex t , and make the identification $t = \perp$.

Definition 5.1. If G is a pointed graph, $\text{Peb}(G)$ is a set of clauses expressed in terms of the variables $\text{Var}(G)$, so that $\text{Peb}(G) = \{\text{Clause}(v) : v \in G\}$.

- (1) If v is a source vertex of G , then $\text{Clause}(v) = v$.
- (2) If v is a vertex in G , with predecessor u , then $\text{Clause}(v) = u \rightarrow v$.
- (3) If v is a vertex in G , with predecessors u, w , then $\text{Clause}(v) = u, w \rightarrow v$.

If we set some variables in $\text{Peb}(G)$, then the resulting set of clauses is not necessarily of the form $\text{Peb}(G')$, where G' is a subgraph of G . We shall focus on a family of special assignments, called *pebbling assignments*, that preserve this property. If $v \in G$, $v \neq t$, then we define the assignment $\llbracket v := 1 \rrbracket$ to be the assignment defined by first setting the variable v to 1, and then setting to 1 any variable u for which there is no implicational chain from u to \perp in the resulting clause set. The assignment $\llbracket v := 0 \rrbracket$ is defined as follows: first, choose a directed path $\pi = (v, \dots, t)$ from v to the target t , set all the vertices in the path to 0, and in addition set any vertex from which v is not reachable, but not in the path π , to 1. The assignment $\llbracket v := 0 \rrbracket$ is not uniquely determined by this construction, since it depends on the path chosen – however, this is not important, since the set of clauses $\text{Peb}(G) \upharpoonright \llbracket v := 0 \rrbracket$ resulting from the restriction is independent of the path. A *pebbling assignment* results from a sequence of restrictions of the form $\llbracket v := 0 \rrbracket$ and $\llbracket w := 1 \rrbracket$.

The effect of the restrictions just defined can be described directly as an operation on the underlying graph. If G is a pointed graph, and $v \in G$, $v \neq t$, $G[v := 1]$ is the graph resulting from G by first removing v , together with all edges entering or leaving v , and then restricting the resulting graph to the vertices from which the target vertex t is accessible. $G[v := 0]$ is the pointed graph $G \upharpoonright v$.

Lemma 5.2.

- (1) For $b = 0, 1$, $\text{Peb}(G) \upharpoonright \llbracket v := b \rrbracket = \text{Peb}(G[v := b])$.
- (2) If G is a pointed graph, and $v \in G$, then

$$\#G \leq \max\{\#G[v := 0], \#G[v := 1] + 1\}.$$

Proof. The first part of the lemma follows straightforwardly from the definitions. For the second part, we employ the following strategy in the pebble game on G ; the strategy is the same as the one used in Lemma 15 of [4].

First, follow a minimum cost strategy to pebble v in $G[v := 0]$. Second, leaving a pebble on v , but removing all other pebbles, follow a minimum cost strategy in the pebbling game on $G[v := 1]$ to pebble the target vertex in G , using the extra pebble for any moves where a pebble is needed on v to justify a placement. The cost of this strategy is at most $\max\{\#G[v := 0], \#G[v := 1] + 1\}$. \square

Example 5.3. If G is the pyramid graph shown in Figure 1, then $\text{Peb}(G)$ is the set of clauses

$$\{a, b, c, d, (a, b \rightarrow e), (b, c \rightarrow f), (c, d \rightarrow g), (e, f \rightarrow h), (f, g \rightarrow i), (h, i \rightarrow \perp)\}.$$

The assignment $\llbracket e := 1 \rrbracket$ sets the variables a and e to 1; $\text{Peb}(G) \upharpoonright \llbracket e := 1 \rrbracket$ is $\text{Peb}(G[e := 1])$ where $G[e := 1]$ results from G by removing the vertices a and e . If we choose the path $f \rightarrow h \rightarrow \perp$, then the assignment $\llbracket f := 0 \rrbracket$ sets the variables f and h to 0, while the variables a, d, e, g, i are all set to 1. The set of clauses $\text{Peb}(G) \upharpoonright \llbracket f := 0 \rrbracket$ is $\text{Peb}(G[f := 0])$, where $G[f := 0]$ is the subgraph of G containing only b, c and f .

Lemma 5.4. *If G is a pointed graph with n vertices, then $\text{Peb}(G)$ has a tree resolution refutation with size $2n - 1$.*




Figure 1: A pyramid graph

Proof. Starting with the clause $\text{Clause}(\perp)$ associated with the sink of G , construct a sequence of purely negative clauses, working from the sink to the sources, by successive inferences using input resolution. Let $C \vee \neg w$ be the last clause in the sequence constructed so far, where w is not a source vertex, and C is purely negative. Resolve $C \vee \neg w$ against the clause $\text{Clause}(w) = (u, v \rightarrow w)$ to produce the next purely negative clause in the sequence, $C \vee \neg u \vee \neg v$. The sequence must end in a purely negative clause in which all the literals are of the form $\neg s$, where s is a source vertex. Now resolve each of these negative literals against the one-literal positive clauses corresponding to the sources. In this input refutation, each variable is resolved upon exactly once, so that the refutation has size $2n - 1$. \square

If Σ is a set of clauses, then a *C-critical assignment* is a total assignment to the variables in Σ that makes all the clauses true, except C . In the case of $\text{Peb}(G)$, we are interested in a particular family of critical assignments. Let v be a vertex in G , and $\pi = (v, \dots, t)$ a directed path in G from v to the target vertex t . Set all the vertices in the path π to 0, and all other vertices in G to 1. This assignment makes all of the clauses in $\text{Peb}(G)$ true, except for $\text{Clause}(v)$. An assignment determined by the path π we shall call a *v-critical assignment*, since the clause that it falsifies is associated with the vertex v . Since we have assumed that G is a pointed graph, such v -critical assignments exist for all the vertices v in G , so that $\text{Peb}(G)$ is minimally inconsistent.

Lemma 5.5. *If G is a pointed graph with $\sharp G = p$, then there are at least p vertices v in G for which there is a v -critical assignment for $\text{Peb}(G)$.*

Proof. Every pebbling strategy for G must contain a configuration with p pebbles, so there must be at least p vertices in G . For every vertex in G , we can construct a v -critical assignment for $\text{Peb}(G)$ by choosing a path from v to the target vertex. \square

6. CONSTRUCTING CLAUSE SETS WITH LARGE REGULAR WIDTH

6.1. The basic construction. To produce clause sets requiring large regular width, we start from the set of clauses $\text{Peb}(G)$, where G is a pointed graph with n vertices. We use the abbreviation V for the set of variables $\text{Var}(G)$, and V^p for the set of all sequences of variables in V of length p .

Let σ be a function from G to V^p , that is to say, a function associating a sequence of length p with every $v \in G$. Thus, for each $v \in G$, we have an associated sequence $\sigma(v) = \sigma_1(v), \dots, \sigma_p(v)$, where each $\sigma_j(v)$ is a variable in V ; the sequence may contain repetitions. Now for $v \in G$, define the set $\text{Clauses}^\sigma(v)$ to be the set of all clauses having the form

$$\text{Clause}(v) \vee \pm\sigma_1(v) \vee \dots \vee \pm\sigma_p(v),$$

where $\pm r$, for $r \in V$, is either r or $\neg r$. $\text{Clauses}^\sigma(v)$ contains 2^p clauses of width at most $p + 3$. In addition, for $A \subseteq G$, define

$$\text{Clauses}^\sigma(A) = \bigcup \{ \text{Clauses}^\sigma(v) \mid v \in A \}.$$

The construction just described can be considered as an iteration of the method used to construct the family of clauses $GT'_{n,\rho}$ defined in §3 of [1]. A key difference from the earlier construction is that the $GT'_{n,\rho}$ examples begin from a set of clauses GT'_n that is hard for tree resolution (though easy for general resolution), while the present construction begins from a set of clauses $\text{Peb}(G)$ that is easy for tree resolution.

The clause sets that we construct in this section are of the form $\text{Clauses}^\sigma(G)$, for G a pointed graph with n vertices. To ensure that these clause sets require large regular width, the map σ must satisfy a combinatorial condition that can be stated roughly as follows: the image of any large set of vertices in G has a large intersection with any large set of variables. In the next subsection, we give a precise meaning to the term “large,” and prove the existence of a function σ satisfying the condition, by a probabilistic construction.

For G a pointed graph, and σ a function from G to V^p , define $\text{Peb}^\sigma(G)$ to be $\text{Clauses}^\sigma(G)$. $\text{Peb}^\sigma(G)$ contains $n \cdot 2^p$ clauses of width at most $p + 3$. We shall show in what follows that the sets of clauses $\text{Peb}^\sigma(G)$, for an appropriate family of pointed graphs G and functions σ , require large regular width, but on the other hand have regular tree resolution refutations whose size is linear in $|\text{Peb}^\sigma(G)|$.

6.2. A combinatorial lemma. In this subsection, we formulate and prove the existence result described above, by employing a probabilistic construction. If $\sigma \in X^k$, and $B \subseteq X$, then we use the notation $\sigma \cap B$ for the set of all elements in the sequence σ that also belong to the set B ; similarly, if S is a set of such sequences, then $S \cap B$ is defined to be $\bigcup \{ \sigma \cap B \mid \sigma \in S \}$. For $A \subseteq G$, define $\sigma(A) = \{ \sigma(v) \mid v \in A \}$.

Lemma 6.1. *Let G be a pointed graph with n vertices, $V = \text{Var}(G)$ the set of variables in $\text{Peb}(G)$, and $p = \lceil \log^5 n \rceil$.*

For any $d > 0$, and sufficiently large n , there is a map σ from G to V^p satisfying the condition: For all $A \subseteq G$ and $B \subseteq V$ with $|A| = |B| = \lfloor dn / \log n \rfloor$, $|\sigma(A) \cap B| \geq dn / 2 \log n$.

Proof. Let us associate with each $v \in G$ a random subset of V with size p , chosen with replacement. That is to say, with each $v \in G$, we associate a sequence $\sigma(v) = \sigma_1(v), \dots, \sigma_p(v)$,

where each variable $\sigma_j(v)$ is chosen independently and uniformly at random from the set V of all variables.

In the first part of the proof, let us consider the sets A and B to be fixed subsets of G and V respectively. Define a map σ from G to V^p to be *bad for A and B* if $|\sigma(A) \cap B| \geq dn/2 \log n$; otherwise *good for A and B* . We begin by proving that for fixed sets A and B , a random map σ is bad with exponentially small probability.

To prove this bound on the probability, it is convenient to consider the construction of the map as resulting from a series of independent choices. Divide the sequence $1, \dots, p$ into $q = \lfloor \log^3 n \rfloor$ blocks, so that each block contains at least $\Theta(\log^2 n)$ integers. That is to say, the sequence $1, \dots, p$ can be written as a concatenation $\tau_1 \tau_2, \dots, \tau_q$ of sequences τ_j , each of length at least $\Theta(\log^2 n)$.

Fix a block τ_j , where $1 \leq j \leq q$, and define a random variable Z representing the number of variables in B that are not in the random subset $\tau_j(A)$, that is to say

$$Z(\tau_j) = |\{x \in B \mid x \notin \tau_j(A)\}|.$$

We begin by estimating the expected value of Z .

Let $B = \{b_1, b_2, \dots, b_i, \dots, b_m\}$ where $m = \lfloor dn/\log n \rfloor$. Define an indicator random variable Θ_i by:

$$\Theta_i(\tau_j) = \begin{cases} 1, & \text{if } b_i \notin \tau_j(A) \\ 0, & \text{if } b_i \in \tau_j(A), \end{cases}$$

so that $Z = \Theta_1 + \dots + \Theta_m$. We estimate the expected value of Θ_i by

$$\begin{aligned} E(\Theta_i) &\leq \left(1 - \frac{1}{|V|}\right)^{|A| \cdot |\tau_j|} \\ &\leq \left(1 - \frac{1}{n-1}\right)^{\Theta(n \log n)} \\ &\leq \exp(-\Omega(\log n)), \end{aligned}$$

showing that

$$E(Z) \leq m \cdot \exp(-\Omega(\log n)) = m \cdot o(1).$$

It follows that for any given positive γ , $E(Z) < \gamma m$, for sufficiently large n . For the remainder of the proof, we assume that n is chosen sufficiently large so that $E(Z) < m/8$.

In the second stage of the proof, we need to show that the random variable Z is tightly concentrated around its mean. To do this, we employ a large deviation bound for martingales, following [7].

Order the set A as $\{a_1, \dots, a_m\}$. The sequence $\tau_j(a_1), \tau_j(a_2), \dots, \tau_j(a_m)$ represents a random subset of variables with size $r = m \cdot |\tau_j| = \Theta(n \log n)$. Let \mathcal{R} be the set of all sequences in V of length r . For $\sigma \in \mathcal{R}$, and $1 \leq t \leq r$, define $\sigma \upharpoonright t$ to be the subsequence $\sigma_1, \dots, \sigma_t$. Define an equivalence relation on \mathcal{R} by setting, for $\rho, \sigma \in \mathcal{R}$,

$$\rho \equiv_t \sigma \iff \rho \upharpoonright t = \sigma \upharpoonright t,$$

for $1 \leq t \leq r$, and let \equiv_0 be the universal relation on \mathcal{R} . Let \mathcal{F}_t be the finite Boolean algebra whose atoms are the blocks of the partition of \mathcal{R} induced by \equiv_t , for $0 \leq t \leq r$; the sequence $\mathcal{F}_0, \dots, \mathcal{F}_r$ of Boolean algebras forms a *filtration* over the set \mathcal{R} .

Define a sequence of random variables Z_0, \dots, Z_r by setting $Z_t = E(Z \mid \mathcal{F}_t)$. Then $Z_0 = E(Z)$, $Z_r = Z$, and the sequence Z_0, \dots, Z_r forms a martingale [8, p. 221], the Doob martingale associated with the filtration $\mathcal{F}_0, \dots, \mathcal{F}_r$. The intuitive picture here is that at

time 0, we begin with no specific information about a given sequence σ ; we learn its values one by one at each successive time step t , until we have full information about σ at time r .

If ρ and σ are two sequences in \mathcal{R} that differ at most at a single point, then $|Z(\rho) - Z(\sigma)| \leq 1$. In the terminology of Alon and Spencer [2, p. 89], the random variable Z satisfies the Lipschitz condition relative to the filtration $\mathcal{F}_0, \dots, \mathcal{F}_r$. It follows by Theorem 4.1 of Chapter 7 of the monograph by Alon and Spencer [2, p. 90] that $|Z_{t+1} - Z_t| \leq 1$. Consequently, by the martingale tail inequality of Hoeffding and Azuma [8, p. 221] [2, p. 85],

$$\begin{aligned} P(Z \geq m/2) &\leq P(Z - E(Z) > 3m/8) \\ &< \exp(-(3m/8)^2/2r) \\ &\leq \exp(-\Omega(n/\log^3 n)). \end{aligned}$$

Let W be the random variable representing the number of variables in B not in the image of A under σ :

$$W(\sigma) = |\{x \in B \mid x \notin \sigma(A)\}|.$$

Since the maps τ_1, \dots, τ_q are constructed independently, it follows that

$$P(W \geq m/2) \leq [\exp(-\Omega(n/\log^3 n))]^q = \exp(-\Omega(n)).$$

We can now complete the proof of the existence of a map σ satisfying the condition of Lemma 6.1. The probability that a random map $\rho \in \mathcal{R}$ is bad for some A and B is bounded by

$$\binom{n}{m}^2 \exp(-\Omega(n)).$$

Using the simple inequality

$$\binom{n}{k} \leq \left(\frac{en}{k}\right)^k,$$

found in Bollobás's textbook on graph theory [6, p. 216], the binomial coefficient above can be bounded by

$$\begin{aligned} \binom{n}{m} &\leq \left(\frac{en}{m}\right)^m \\ &= \left(e^{O(\log \log n)}\right)^{O(n/\log n)} \\ &= e^{O(n \log \log n / \log n)}. \end{aligned}$$

Hence, the probability can be bounded above by

$$\exp(O(n \log \log n / \log n)) \exp(-\Omega(n)) = \exp(-\Omega(n)).$$

Consequently, the probability that a random map ρ is bad for some A and B is exponentially small for sufficiently large n , showing that a map satisfying the condition of the lemma must exist. \square

7. SEPARATING REGULAR SIZE AND WIDTH

Let G be a pointed graph, V the set of vertices in G (other than the sink) and N and p positive integers. We define a map σ to be *good for G , N , and p* if it satisfies the condition: There is a map σ from G into V^p so that for any $A \subseteq G$ and $B \subseteq V$, if $|A| = |B| = \lfloor N \rfloor$, then $|\sigma(A) \cap B| \geq N/2$. Lemma 6.1 states that for any $d > 0$, given sufficiently large n , $p = \lceil \log^5 n \rceil$, and $N = dn/\log n$, for every pointed graph with n vertices, there is a map σ that is good for G , N , and p .

This lemma allows to construct a set of examples that have polynomial-size regular resolution refutations, but large regular width. The construction is based on the following result of Paul, Tarjan and Celoni.

Theorem 7.1. [9] *There is a sequence of binary pointed graphs G_1, \dots, G_i, \dots with pebbling number at least $cn(i)/\log n(i)$, for sufficiently large i , where $n(i) = |G_i| = O(i2^i)$, and $c > 1/20$.*

It should be mentioned that the graphs H_i constructed by Paul, Tarjan and Celoni, though binary, are not pointed, since they are constructed to have multiple sink nodes. However, in their main theorem, they show that for sufficiently large i , their graph H_i contains a sink node that requires $cn(i)/\log n(i)$ pebbles to pebble it, starting from the empty configuration. Hence, we can construct a pointed graph from H_i by choosing such a sink node, and considering the subgraph G_i containing all the nodes from which this sink is accessible. This subgraph G_i still satisfies the condition $n(i) = |G_i| = O(i2^i)$, so the main theorem of Paul, Tarjan and Celoni continues to hold, if we add the qualifier “pointed.”

Lemma 7.2. *Let G be a pointed graph with n vertices, and σ a map from G to V^p , where $p = \lceil \log^5 n \rceil$. Then the set of clauses $\text{Peb}^\sigma(G)$ contains $n-1$ variables and $n^{O(\log^4 n)}$ clauses, and has a regular tree refutation with size $n^{O(\log^4 n)}$.*

Proof. By Lemma 5.4, $\text{Peb}(G)$ has an input refutation with size at most $2n-1$. For a given vertex v in G , the clause $C(v)$ associated with the vertex can be derived from $\text{Clauses}^\sigma(v)$ by a tree resolution proof with size $2^{O(\log^5 n)} = n^{O(\log^4 n)}$. Consequently, $\text{Peb}^\sigma(G)$ has a tree refutation with size $O(n) \cdot n^{O(\log^4 n)}$, that is, $n^{O(\log^4 n)}$. This tree refutation may not be regular; however, if irregularities are present, it is possible to remove them [11] [12, p. 436] resulting in a smaller regular tree-style refutation. \square

Lemma 7.3. *Let G be a pointed graph with pebbling number $\sharp G = N$ and σ a map that is good for G , $N/2$, and p . Then any regular resolution refutation of $\text{Peb}^\sigma(G)$ must contain a clause with width at least $N/4$.*

Proof. We prove the Lemma by showing that the Adversary wins the regular $N/4$ -width game based on $\text{Peb}^\sigma(G)$. The winning strategy has two stages. In the first stage, the Adversary maintains a pebbling assignment to G ; at the start of the game, this assignment is empty. In the second stage, the Adversary answers according to a fixed v -critical assignment.

Assume that it is the Adversary’s turn, that π is the current pebbling assignment to the variables V in $\text{Peb}^\sigma(G)$, and that x is the variable currently queried by the Prover. The Adversary answers the current query according to these rules.

- (1) If the variable x is already assigned a value by π , then answer the query according to π ;

- (2) If the variable queried is not assigned a value by π , then it must be associated with a node $v \in G \upharpoonright \pi$. Extend π to a pebbling assignment π' so that π' contains $\llbracket v := b \rrbracket$, choosing b so as to maximize the pebbling number of $G \upharpoonright \pi'$.

The Adversary continues to play according to these rules until $\lfloor N/2 \rfloor$ nodes in G have been queried; when this happens, the first stage is completed.

With the first stage completed, let α be the current extended assignment, π the current pebbling assignment maintained by the Adversary; we assume that it is the Prover's turn. By Lemma 5.5, there are at least $N/2$ vertices $v \in G \upharpoonright \pi$ for which there is a v -critical assignment for $\text{Peb}(G \upharpoonright \pi)$. If ϕ is such a critical assignment, then $\pi \cup \phi$ is a v -critical assignment for $\text{Peb}(G)$. Let A be the set of all nodes in G satisfying this condition, and B the set of variables queried in the game so far. Because $|A|, |B| \geq \lfloor N/2 \rfloor$, $|\sigma(A) \cap B| \geq N/4$, since σ is good for G , N , and p .

Since the Prover and Adversary are playing the regular $N/4$ -width game, it follows that $|\alpha| < N/4$ (since the current assignment after the Adversary's reply has width $|\alpha| + 1$). Hence, at least one variable v in $\sigma(A) \cap B$ must be forgotten in α . Let ϕ be a v -critical assignment for $\text{Peb}(G)$; ϕ is also a v -critical assignment for $\text{Peb}^\sigma(G)$. In the second stage of the strategy, the Adversary answers all queries in accordance with the assignment ϕ . Since ϕ makes all of the clauses in $\text{Peb}^\sigma(G)$ true, except for a clause in $\text{Clauses}(v)$ containing the variable v , this strategy results in a win for the Adversary, since the variable v is forgotten, so the Prover cannot query it again. \square

Theorem 7.4. *There is an infinite sequence $\Sigma_1, \Sigma_2, \dots, \Sigma_i, \dots$ of contradictory sets of clauses and a corresponding list of parameters $n(1), n(2), \dots, n(i), \dots$ so that (abbreviating $n(i)$ as n):*

- (1) *Each clause set Σ_i contains $n - 1$ variables and $n^{O(\log^4 n)}$ clauses with width $O(\log^5 n)$;*
- (2) *Σ_i has a regular tree refutation with size $n^{O(\log^4 n)}$;*
- (3) *Any regular refutation of Σ_i must contain a clause with width $\Omega(n/\log n)$.*

Proof. Define $\Sigma_i = \text{Peb}^\sigma(G_i)$, where G_i is one of the sequence of pointed graphs in Theorem 7.1. The theorem follows by Lemmas 6.1, 7.2 and 7.3. \square

Although the clause sets in Theorem 7.4 have size quasi-polynomial in n , they have regular tree refutations that are linear in the size of the clause sets themselves. Furthermore, if we compute the significant quantities in the second part of Theorem 4.2, we find that if the corresponding theorem held for regular size and width, then regular refutations of these clause sets would have to have size exponential in $n/\log^2 n$. This shows that the relations between size and width holding for tree resolution and general resolution cannot be generalized to the case of regular resolution.

ACKNOWLEDGMENTS

I wrote this paper for a five day workshop on proof complexity at the Banff International Research Station in October 2011. I would like to express my thanks to the organizers, Sam Buss, Stephen Cook, Antonina Kolokolova, Toni Pitassi and Pavel Pudlák for a most stimulating workshop, and also to Paul Beame, who, following my talk, pointed out a computational error in the original version of the paper.

REFERENCES

- [1] Michael Alekhnovich, Jan Johannsen, Toniann Pitassi, and Alasdair Urquhart. An exponential separation between regular and general resolution. *Theory of Computing*, 3:81–102, 2007. Preliminary version in Proceedings of the 34th Annual ACM Symposium on Theory of Computing: May 19-21 2002, Montréal, Québec, Canada.
- [2] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. John Wiley, 1992.
- [3] Albert Atserias and Victor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74:323–334, 2008. Preliminary version: 18th IEEE Conference on Computational Complexity, pp. 239-247, 2003.
- [4] Eli Ben-Sasson, Russell Impagliazzo, and Avi Wigderson. Near optimal separation of tree-like and general resolution. *Combinatorica*, 24:585–603, 2004. Preliminary version, ECCC TR00-005, 2000.
- [5] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow – resolution made simple. *Journal of the Association for Computing Machinery*, 48:149–169, 2001. Preliminary version: Proceedings of the 31st Annual ACM Symposium on Theory of Computing, 1999, pp. 517-526.
- [6] Béla Bollobás. *Modern Graph Theory*. Springer-Verlag, 1998. Graduate Texts in Mathematics 184.
- [7] Anil Kamath, Rajeev Motwani, Krishna Palem, and Paul Spirakis. Tail bounds for occupancy and the satisfiability threshold conjecture. *Random Structures and Algorithms*, 7:59–80, 1995.
- [8] Colin McDiarmid. Concentration. In Michel Habib, Colin McDiarmid, Jorge Ramirez-Alfonsin, and Bruce Reed, editors, *Probabilistic Methods for Algorithmic Discrete Mathematics*, pages 195–248. Springer, 1998. Algorithms and Combinatorics 16.
- [9] W.J. Paul, R.E. Tarjan, and J.R. Celoni. Space bounds for a game on graphs. *Mathematical Systems Theory*, 10:239–251, 1977.
- [10] Jörg Siekmann and Graham Wrightson, editors. *Automation of Reasoning*. Springer-Verlag, New York, 1983.
- [11] G.S. Tseitin. On the complexity of derivation in propositional calculus. In A. O. Slisenko, editor, *Studies in Constructive Mathematics and Mathematical Logic, Part 2*, pages 115–125. Consultants Bureau, New York, 1970. Reprinted in [10], Vol. 2, pp. 466-483.
- [12] Alasdair Urquhart. The complexity of propositional proofs. *The Bulletin of Symbolic Logic*, 1:425–467, 1995.
- [13] Alasdair Urquhart. A near-optimal separation of regular and general resolution. *SIAM Journal on Computing*, 40:107–121, 2011.