

Resolution and Pebbling Games

Nicola Galesi* and Neil Thapen**

Abstract. We define a collection of Prover-Delayer games to characterise some subsystems of propositional resolution. We give some natural criteria for the games which guarantee lower bounds on the resolution width. By an adaptation of the size-width tradeoff for resolution of [10] this result also gives lower bounds on proof size.

We also use games to give upper bounds on proof size, and in particular describe a good strategy for the Prover in a certain game which yields a short refutation of the Linear Ordering principle.

Using previous ideas we devise a new algorithm to automatically generate resolution refutations. On bounded width formulas, our algorithm is at least as good as the width based algorithm of [10]. Moreover, it finds short proofs of the Linear Ordering principle when the variables respect a given order.

Finally we approach the question of proving that a formula F is hard to refute if and only if is “almost” satisfiable. We prove results in both directions when “almost satisfiable” means that it is hard to distinguish F from a satisfiable formula using limited pebbling games.

1 Introduction

Propositional resolution is one of the most intensively studied logical systems. It is important both from applied and from theoretical points of view. On one hand it provides the logical basis for almost all of the more important and efficiently implemented automatic theorem provers (see [6]). On the other hand, it has probably been the most studied proof system in the area of proof complexity ([16, 7, 5, 10, 22] among others).

Most of the work to understand the strength of resolution has been concentrated on proving lower bounds for the length of refutations. Recently Ben-Sasson and Wigderson [10] based on ideas of [7] gave a unified approach to obtaining lower bounds. They showed that if a bounded width formula has a short refutation, then it has a narrow refutation. Using this relationship they give an algorithm to generate resolution refutations based on the width measure.

Simple pebbling games were initially introduced into the world of resolution in [20] to study size lower bounds in the subsystem of resolution where the proofs are treelike. Later the study of the space measure for resolution (see [13, 1, 8]) suggested the use of a more complex pebbling game.

It was soon clear that games also play an important role also in the study of the width limit for the full resolution system. Atserias and Dalmau [3] have

* Departament LSI. Universitat Politècnica Catalunya. Barcelona, Spain.

** St Hilda's College. University of Oxford. Oxford, UK.

given a finite model-theoretic characterization of the bounded width formulas with narrow refutations, using a pebbling game. Atserias, Kolaitis and Vardi in [4] gave a characterization, by combinatorial games from finite model theory, of refutational width in a refutational system tailored for constraint satisfaction problems.

In this paper we carry further the idea of using pebble games to study resolution. In section 3 we define a new modification of the resolution system (narrow resolution) for which we give a unified way of proving lower bounds similar to those for bounded width resolution of [10], but without having the restriction that the original formula has bounded width.

We define a “witnessing” pebble game, played between a Prover and a Delayer, and show that proofs in this system correspond to strategies for the Prover. On the other hand, a good strategy for the Delayer corresponds to the formula being extended dynamically satisfiable (EDS), an extension of an idea used in [12, 14] to prove lower bounds on space and on treelike size.

In section 4 we give some sufficient conditions for a formula to be k -EDS. In particular we show that if there is a satisfiable formula G which looks similar to F in that G cannot be distinguished from F with $k + 1$ or fewer pebbles, then F is k -EDS. We also adapt a criterion of Riiis’ [21] to show that if F is a translation of a combinatorial principle on some finite structure, and this principle is satisfiable on a larger structure, then F is $\Omega(n)$ -EDS. This gives us a useful sufficient condition for proving width lower bounds.

Starting from this witnessing pebble game we explore in two directions. In Section 5 we study some more general “structured” games, and generalize the width concept. In the earlier game the restriction on width corresponded to the Prover only being able to remember the values of a limited number of variables. In the structured game, the number of pebbles still limits how much the Prover can remember, so the games have useful properties in common with bounded width proofs; but each pebble can be labelled with information about several variables, which means the Prover can refute things that would be impossible if he was limited by width.

We show how to recover Resolution refutations from Prover strategies (theorem 20). In particular, since we show that the Prover has a winning strategy for a certain “ordered” structure game with only three pebbles over the Linear Ordering Principle LOP (theorem 26), we can recover a polynomial size refutation of this principle.

The LOP principle was used in [11] to prove the optimality of the width-size tradeoff for Resolution. It is one of the very few examples of CNF formulas having polynomial size resolution refutations but on which the Ben-Sasson Wigderson algorithm takes a very long (just subexponential) time to recover a refutation. Moreover it is also hard for DPLL-based theorem provers. Motivated by the previous result we then look at the the question of automatically generating resolution refutations using strategies for the game.

We devise a new algorithm based on reconstructing strategies for the Prover which allows us to extend the Ben-Sasson and Wigderson algorithm to an al-

gorithm which, with some extra information about an ordering of the variables, generates short refutations of LOP in polynomial time. Moreover we show that for bounded width formula our algorithm is at least as good as the Ben-Sasson and Wigderson algorithm.

Finally in section 6 we explore further the idea used in section 4, that if a CNF F is similar (in some finite model theoretic sense) to a satisfiable formula G , then F can be thought of as “almost satisfiable” and is hard to refute. While this last result can be seen as a soundness theorem, in this section we also give a kind of “completeness” theorem, and show that if F is hard to refute using a certain game, then F is similar to a satisfiable CNF G .

Due to space limitations many proofs are shortened or omitted. A full version is available as an ECCC technical report [15].

2 Preliminary Definitions

Resolution is a refutation proof system for formulas in CNF form based on the following rule: $\frac{C \vee x \quad \bar{x} \vee D}{C \vee D}$ where if C and D have common literals, they appear only once in $C \vee D$. A resolution refutation of a CNF formula F is a derivation of the empty clause from the clauses defining F , using the above inference rule. As usual a refutation can be view as a directed acyclic graph.

The *size* of a refutation is the number of clauses in it. The *width* of clause is the number of literals in it; the width of a refutation is the maximal width of a clause in the proof; the width of refuting a formula is the minimal width of a refutation of that formula.

We consider two standard families of CNF, the *Pigeon Hole Principle*, PHP_n^m , with variables p_{ij} expressing “pigeon i goes to hole j ”:

$$\bigwedge_{i \in [m]} \bigvee_{j \in [n]} p_{ij} \quad \wedge \quad \bigwedge_{i, i' \in [m], j \in [n], i \neq i'} (\neg p_{ij} \vee \neg p_{i'j})$$

and the linear ordering principle LOP_n , that expresses (the negation of) that every linear ordering of n elements has a least element, with variables $x_{i,j}$ expressing “ i is less than j ”:

$$\bigwedge_{i, j, k \in [n]} (\neg x_{i,j} \vee \neg x_{j,k} \vee x_{i,k}) \quad \wedge \quad \bigwedge_{i, j \in [n]} (x_{i,j} \vee x_{j,i}) \quad \wedge \\ \bigwedge_{i, j \in [n]} (\neg x_{i,j} \vee \neg x_{j,i}) \quad \wedge \quad \bigwedge_{i \in [n]} \bigvee_{j \in [n], i \neq j} x_{j,i}.$$

3 The Witnessing Game

Definition 1 Fix $k \in \mathbb{N}$. Call a clause narrow if it has width k or less; otherwise it is wide. A width k narrow resolution refutation of a CNF F is a sequence of narrow clauses, beginning with the narrow clauses of F and finishing with the empty clause. There are three ways that clauses can be introduced into the sequence:

1. From B we can derive Bx (weakening);

2. From Bx and $C\bar{x}$ we can derive BC (resolution);
3. If $x_1 \dots x_m$ is a (usually wide) clause in F , then from $B\bar{x}_1, \dots, B\bar{x}_m$ we can derive B (resolution by cases).

Lemma 2 *If F is a r -CNF with a width k narrow resolution refutation, then F has a width $r + k - 2$ “normal” resolution refutation.*

We introduce a pebble game to accompany this proof system.

Definition 3 *Let F be a CNF. The witnessing pebble game on F is played between a Prover and a Delayer on the set of literals arising from the variables in F . A pebble can never appear on both a literal and its negation. Normally we will limit the game to some number k of pebbles, and call this the k -pebble witnessing game. In each turn, one of three things can happen.*

1. The Prover lifts a pebble from the board; the Delayer makes no response.
2. (Querying a variable) The Prover gives a pebble to the Delayer and names an empty variable x (that is, neither x nor \bar{x} can be pebbled already). The Delayer must put the pebble on either x or \bar{x} .
3. (Querying a clause) The Prover gives a pebble to the Delayer and names a clause C from F . The Delayer must place the pebble on one of the literals in C , without contradicting any pebble already on the board. If this is not possible then the Prover wins.

Notice that the Prover can win exactly when the pebbles on the board falsify some clause of F and the Prover has one pebble left over.

The next theorem describes the exact relationships between winning strategies for the Prover in the witnessing game, and refutational size and space.

Theorem 4 *Let F be a CNF and $k \in \mathbb{N}$.*

1. *If there is a winning strategy for the Prover in the k -pebble witnessing game for F , then there is a narrow resolution refutation of F of width k .*
2. *If there is a narrow resolution refutation of F of width k then there is a winning strategy for the Prover in the $(k + 1)$ -pebble witnessing game for F .*
3. *If F has a (normal) resolution proof of width k , then there is a winning strategy for the Prover in the $(k + 1)$ -pebble witnessing game for F .*
4. *If F has a (normal) resolution proof of clause space k , then there is a winning strategy for the Prover in the k -pebble witnessing game for F .*

Proof We include a proof of 1, as it illustrates the correspondence between a strategy and a proof which we will be using throughout this paper. Consider the Prover’s strategy as a tree, with each node labelled with the set of literals falsified under the assignment given by the pebbles currently in play. Then the root will be the empty clause, and the leaves will be (some subset of) the narrow clauses of F . If we read this tree from the leaves down to the root, we get precisely a narrow resolution refutation of F . Removing a pebble corresponds to weakening,

the Prover querying a variable corresponds to a resolution step, and the Prover querying a clause corresponds to a resolution-by-cases step. \square

We can now adapt the Ben-Sasson Wigderson result that “short proofs are narrow” to talk about games rather than proofs. This allows us to apply it directly to CNFs of unbounded width.

Definition 5 *If F is a CNF and x is a literal, we obtain $F|x$ from F by removing all clauses containing x and removing \bar{x} from any clause in which it appears.*

So from a resolution refutation of F we can obtain a refutation of $F|x$ of equal size or smaller, by substituting in a value of “true” for x and simplifying.

Lemma 6 *If the Prover has a winning strategy for the k -pebble witnessing game on $F|x$, then in the k -pebble witnessing game on F the Prover can force the Delayer to either lose the game or place a pebble on \bar{x} .*

Proof Let S be the k -pebble winning strategy for $F|x$. We will make this into a strategy for the game on F as follows. Whenever a clause C is queried in S such that $C \in F|x$ but $C \notin F$, it must be that $C\bar{x} \in F$. So replace this query with a query of $C\bar{x}$. Then either the Delayer must eventually place a pebble on \bar{x} , or the play of pebbles must be exactly the same as given in strategy S so that the Delayer must eventually lose. \square

Theorem 7 ([10]) *Fix $d, n \in \mathbb{N}$ and let $\beta = (1 - \frac{d}{2n})^{-1}$. Say that a clause is fat if it has width greater than d . Then for any $m \leq n$ and any b , if F is a CNF on m variables and has a (normal) resolution refutation Π containing $< \beta^b$ many fat clauses, then the Prover has a winning strategy in the witnessing pebble game on F , using $d + b + 1$ pebbles. (See corollary 11 for an application of this.)*

Proof The proof is by induction on m . The base case $m = 0$ is trivial, so suppose $m > 0$.

If $b = 0$, then every clause in Π (and also every clause in F) has width $\leq d$, so by an earlier observation there is a strategy for the Prover using $d + 1$ pebbles.

Otherwise, let Π^* be the set of fat clauses appearing in Π . Then there must be some literal x appearing in at least $\frac{d}{2n}|\Pi^*|$ fat clauses, since otherwise $|\Pi^*|d \leq |\{(y, C) : y \text{ is a literal in } C \in \Pi^*\}| < 2m\frac{d}{2n}|\Pi^*|$.

The first part of the Prover’s strategy is to force the Delayer to put a pebble on x . Now $F|\bar{x}$ contains only $m - 1$ variables and has a refutation with fewer than β^b fat clauses, so by the inductive hypothesis the Prover has a strategy for the game on $F|\bar{x}$ using $b + d + 1$ pebbles. Hence by the lemma the Prover can force the Delayer to satisfy x . Setting x to true will make all the clauses in Π containing x vanish, so $F|x$ contains only $m - 1$ variables and has a refutation with fewer than $(1 - \frac{d}{2n})|\Pi^*| \leq \beta^{b-1}$ fat clauses. Hence the Prover has a winning strategy T for $F|x$ with only $b + d$ pebbles.

The Prover now leaves one pebble on x and uses the remaining $b + d$ pebbles to carry out strategy T on the remaining variables. As in the lemma, he must

change T slightly to make it into a strategy for the game on F , by replacing queries to $C \in F|x \setminus F$ with queries to $C\bar{x}$. But the Delayer can never put a pebble on \bar{x} , because there is already a pebble on x . So the game plays just like the $F|x$ game with strategy T , and the Prover wins. \square

4 Extended dynamic satisfiability

Along similar lines to those used in [3] for resolution width and in [12] for resolution space, we characterize the Delayer's strategy in the witnessing game, in terms of families of partial assignments for the formula. We also generalize this result by studying sufficient conditions that imply good strategies for the Delayer, along the lines of the approaches in [21] and [18] using first order model theory. The following definition was introduced in [12].

Definition 8 *A CNF F is k -dynamically satisfiable (k -DS) if there is a class R of partial assignments to the variables of F with the following properties:*

1. R is closed under subset;
2. If $\alpha \in R$, $|\alpha| < k$ and C is any clause of F , then there is an extension $\beta \in R$ of α that satisfies C (in the sense that it makes at least one of the literals in C true, but does not necessarily assign a value to all the literals).

To characterize good strategies for the delayer in our witnessing game, we alter the definition of dynamic satisfiability by adding a case to deal with queries made to variables:

Definition 9 *A CNF F is k extended-dynamically satisfiable (k -EDS) if there is a class R of partial assignments satisfying the conditions of definition 8, with the extra case:*

3. If $\alpha \in R$, $|\alpha| < k$ and x is any variable appearing in F , then there is an extension $\beta \in R$ of α that assigns some value to the variable x .

Lemma 10 *A CNF F is k -extended dynamically satisfiable if and only if the Delayer has a winning strategy for the k -pebble witnessing game on F .*

Proof Suppose F is k -EDS. Then the Delayer can guarantee that after every turn the assignment α given by the k pebbles is in R . Hence by parts 2 and 3 of the definition of extended dynamic satisfiability, the Delayer is always able to consistently satisfy any variable or clause of F that the Prover queries. Conversely, suppose that the Delayer has a winning strategy. Let R be the set of all assignments corresponding to the configurations of pebbles that can appear in a game in which the Delayer uses this strategy. Then R witnesses that F is k -EDS. \square

Corollary 11 *(to theorem 7) For any $\varepsilon > 0$, if a CNF F has n variables and is $(\sqrt{8n}^{\frac{1+\varepsilon}{2}} + 1)$ -EDS, then it has no resolution refutation of size 2^{n^ε} .*

4.1 A sufficient condition for extended dynamic satisfiability

We will treat CNFs as two sorted structures, with a clause sort and a variable sort and two binary relations “variable x appears positively in clause C ” and “variable x appears negatively in clause C ”. We describe a kind of pebble game, the $(1, k)$ -embedding game. The game is played between a Spoiler and a Duplicator on the set of clauses and variables of two CNFs F and G . Each player has k variable pebbles and one clause pebble.

Each turn the Spoiler either plays a clause pebble on one of the clauses of F , in which case the Duplicator must play his corresponding pebble on one of the clauses of G ; or the Spoiler plays a variable pebble on one of the variables of either CNF, in which case the Duplicator must place his corresponding pebble on one of the variables of the other CNF.

At the end of each turn, the pebbles in play give a correspondence between some of the clauses and variables of F and those of G . The aim of the Duplicator is to make sure that this is a partial isomorphism. If at any point it is not, then the Spoiler has won.

Theorem 12 *Suppose G is satisfiable, and there is a winning strategy for the Duplicator in this game. Then there is a winning strategy for the Delayer in the k pebble witnessing game on F . Hence F is k -EDS.*

For an example, let F be PHP_n^{n+1} and G be PHP_n^n . Then G is satisfiable, and the Duplicator wins the $(1, n - 2)$ -embedding game on F and G (using a strategy generated by the set of partial injective mappings from the pigeons of F to the pigeons of G). This “almost satisfiability” of PHP seems to be related to the idea of critical assignments that is often used in hardness proofs for it.

In general this theorem is not easy to use, because it is not necessarily easy to prove that two CNFs F and G are in this relationship (although a sufficient condition is for F and G to be indistinguishable in the normal $k + 1$ pebble game of finite model theory). We give an easier-to-use, but weaker, sufficient condition for extended dynamic satisfiability below, for the case where F and G are propositional translations of some first order principle. The argument is a standard one, see Krajicek [18, 17] or Riis [21] although we do not insist that the structure in which the principle is satisfied is infinite.

Theorem 13 *Let F be a CNF. Let ϕ be a first order quantifier free formula in a relational language L with equality. Let Φ be the formula $\forall x_1 \in [n_1] \dots \forall x_k \in [n_k] \exists x_{k+1} \in [n_{k+1}] \dots \exists x_l \in [n_l] \phi(\bar{x})$, where $n_1, \dots, n_l \in \mathbb{N}$.*

Suppose that Φ is satisfiable “on a larger domain”, that is, there is an interpretation in \mathbb{N} of the relation symbols from L such that, with this interpretation, $\mathbb{N} \models \forall x_1 \in S_1 \dots \forall x_k \in S_k \exists x_{k+1} \in S_{k+1} \dots \exists x_l \in S_l \phi(\bar{x})$, where S_1, \dots, S_l are (possibly infinite) subsets of \mathbb{N} with $|S_j| \geq n_j$ for each j and $n_i \geq n_j \rightarrow S_i \supseteq S_j$ for each i, j .

Let $\langle \Phi \rangle$ be the formula $\bigwedge_{i_1 \in [n_1], \dots, i_k \in [n_k]} \bigvee_{i_{k+1} \in [n_{k+1}], \dots, i_l \in [n_l]} \phi(\bar{i})$. We treat this as a propositional formula in the usual fashion, thinking of atomic sentences involving a relation symbol as propositional variables, and of atomic sentences

involving equality between numbers as the appropriate one of the connectives $\{T, F\}$. Let $n = \min\{n_1, \dots, n_l\}$ and let r be the maximum arity of any relation symbol. Then $\langle \Phi \rangle$ is $(\frac{n}{r} - l + 1)$ -EDS (l is the number of variables).

It follows that if F is a CNF, and we can fix a one-to-one renaming of the propositional variables of $\langle \Phi \rangle$ with respect to which every clause of F is implied by some clause of $\langle \Phi \rangle$, and every variable in F appears in $\langle \Phi \rangle$, then F is $(\frac{n}{r} - l + 1)$ -EDS. (We call this condition “ F is covered by Φ ”.)

Corollary 14 LOP_n is $(\frac{n}{2} - 3)$ -EDS. (Note that LOP_n has n^2 variables, so corollary 11 does not apply.)

Proof LOP_n is covered by the formula $\forall x, y, z \in [n] \exists w \in [n], (\neg P(x, y) \vee \neg P(y, x)) \wedge (P(x, y) \vee P(y, x)) \wedge (P(x, y) \wedge P(y, z) \rightarrow P(x, z)) \wedge P(w, x)$, and this is satisfiable if we let the quantifiers range over all of \mathbb{N} and interpret P as any total ordering with no least element. \square

5 Pebbling games and subsystems of resolution

We represent (usually partial) assignments by the following notation: $[x_1 \mapsto 1, \dots, x_n \mapsto 0]$. Given an assignment α , we denote by $\bar{\alpha}$ the negation of α , that is, if α is the assignment $[x \mapsto 1, y \mapsto 0, z \mapsto 1]$ then $\bar{\alpha}$ is the clause $(\neg x \vee y \vee \neg z)$. Given a set S of assignments, we denote by \bar{S} the set of clauses obtained by the negation of all assignments in S .

Definition 15 (Structure) Let F be a CNF formula. For each clause C of F , let S_C be a set of partial assignments to the variables of C , such that

1. each assignment in S_C satisfies C
2. C implies the disjunction of the assignments (in other words, $C \cup \bar{S}_C$ is a contradictory set of clauses).

We call $\mathcal{S} = \bigcup_{C \in F} S_C$ a structure for F .

The *Structured Witnessing Game* $SWG(F, \mathcal{S}_F)$, over a CNF F and with a structure \mathcal{S}_F , is a two player (Prover and Delayer) game defined as follows.

At each round the Prover either

- puts a pebble on some clause C of F . Then the Delayer answers by choosing one assignment $\alpha \in S_C$, and labelling the pebble with it; the Delayer is not allowed to choose an assignment inconsistent with an assignment already in play
- or removes a pebble, together with with its label.

The game ends when the Delayer is unable to choose an assignment consistently. This only happens when the the assignments labelling all the pebbles in play together falsify some initial clause of F .

We say that a formula F is (k, \mathcal{S}) -easy if the Prover has a winning strategy for the game $\text{SWG}(F, \mathcal{S})$ using at most k pebbles simultaneously. Otherwise the Delayer has a winning strategy, and we say that F is (k, \mathcal{S}) -hard.

We will look at three structures in this paper.

- In the *unary* structure \mathcal{U}_F , for each $C \in F$, $S_C = \{[l \mapsto 1] : l \in C\}$. In other words, assignments in S_C satisfy exactly one literal in C . Note that the game for this structure is similar to the witnessing game described above, except that here the Prover is not allowed to query the value of individual variables.
- In the *ordered* structure \mathcal{O}_F , for each clause $C \in F$ we first fix a total order \prec on the variables of C (extended to literals by ignoring negations), then define $S_C = \{(\bigcup_{r \in C, r \prec l} [r \mapsto 0]) \cup [l \mapsto 1] : l \in C\}$.
- In the *full structure* \mathcal{F}_F , for each clause C we let S_C be the set of all possible assignments to all of the variables in C .

In lemma 24 we show that the PHP_n^m is $(n/2 - 1, \mathcal{F})$ -hard. On the other hand we also prove that LOP_n formulas are $(3, \mathcal{O})$ -easy, and that this gives us an upper bound on the size of refutations of LOP_n .

Clearly the more complex the structure is, the more information the Prover can get using fewer pebbles, and the easier it will be to force the Delayer into a contradiction. This is captured by the following lemma.

Definition 16 *Let \mathcal{S}_F and \mathcal{T}_F be two structures for F . We write $\mathcal{S}_F \leq \mathcal{T}_F$ if for all $C \in F$, it holds that: (1) for all $\alpha \in S_C$ there is a $\beta \in T_C$ such that $\beta \subseteq \alpha$; (2) for all $\beta \in T_C$ there is $\alpha \in S_C$ such that $\beta \subseteq \alpha$.*

Lemma 17 *If F is (k, \mathcal{S}_F) -hard and $\mathcal{S}_F \leq \mathcal{T}_F$, then F is (k, \mathcal{T}_F) -hard.*

5.1 Feasible Structured Games and Short Refutations

Definition 18 (Feasible Structure) *Let F be a CNF formula. We say that a structure $\mathcal{S}_F = \bigcup_{C \in F} S_C$ is feasible if there are two polynomials p and q , such that for all $C \in F$:*

- $|S_C| \leq p(|F|)$;
- there is a resolution refutation of $C \cup \overline{S}_C$ of size bounded by $q(|F|)$.

Lemma 19 *The unary and the ordered structures are feasible.*

From a winning strategy for the Prover in a feasible structured game, we can construct a resolution refutation of F .

Theorem 20 *Let F be CNF with m clauses and let \mathcal{S}_F be a feasible structure for F . If F is (k, \mathcal{S}_F) -easy, then there is a resolution refutation of F of size bounded by $O(m^k |\mathcal{S}_F|^k) q(|F|)$.*

Proof (sketch) Think of the Prover’s strategy as a dag. Each node is labelled with a position in the game and the query that the Prover makes from that position, hence there are at most $O(m^k |\mathcal{S}|^k)$ nodes. There is an edge going out from a node for each possible reply that the Delayer makes.

We make this strategy into a resolution refutation as in the proof of theorem 4. We negate the assignments at each node so that the nodes are now labelled with clauses, and we reverse the direction of all the edges. Finally we replace all the edges now coming into a node with the resolution refutation (of $C \cup \bar{S}_C$ from the definition of feasible structure) corresponding to the clause C being queried at that node. \square

5.2 Structured Games as Subsystems of Resolution

In [14] it is proved that k -dynamic satisfiability is a sufficient condition for a CNF F to require exponential size treelike resolution refutations. In the next theorem we prove that it completely characterizes Delayer strategies for the unary structure.

Theorem 21 *F is k -dynamically satisfiable if and only if it is (k, \mathcal{U}) -hard.*

Extended dynamic satisfiability corresponds to the witnessing game, in which the Prover is allowed to query variables. That is not allowed in our structured games, but we do have the following relationship.

Theorem 22 *If F is (k, \mathcal{F}) -hard (that is, hard for the full game) then F is k -EDS.*

Now we will consider strategies for the Prover and Delayer for two standard families of CNFs, the pigeonhole principle PHP_n^m and LOP_n :

From definition 16 and lemma 17, it is straightforward to observe that:

Lemma 23 $\mathcal{F} \leq \mathcal{O} \leq \mathcal{U}$.

It is quite easy to prove that the PHP_n^m is hard for all these games:

Lemma 24 PHP_n^m is $(n/2 - 1, \mathcal{F})$ -hard.

On the other hand, using the fact that LOP_n is $(\frac{n}{2} - 3)$ -dynamically satisfiable (see [12] or corollary 14), we have:

Corollary 25 LOP_n is $(\frac{n}{2} - 3, \mathcal{U})$ -hard.

The main result of this subsection follows from the next theorem.

Theorem 26 LOP_n is $(3, \mathcal{O})$ -easy.

Proof Consider the following order on all variables, that in particular defines an order on each clause: $x_{i,j} \prec x_{h,k}$ iff either $j < k$ or $j = k$ and $i < h$.

We describe the strategy of the Prover by stages: we prove that at each stage the Prover, using only three pebbles, either wins or will force the Delayer to answer to a clause of the form $\bigvee_{j \in [n], j \neq r} x_{j,r}$ with an assignment assigning strictly more literals to 0 than the previous stage. Hence, if he does not win sooner, after at most n stages he will force a clause of this form to be falsified.

Assume w.l.o.g. that at the beginning of a stage the Prover pebbles the clause $C_1 = \bigvee_{j \in [n], j \neq 1} x_{j,1}$. Let $\alpha \in S_{C_1}$ be the assignment chosen by the Delayer. α will be of the form $[x_{2,1} \mapsto 0, \dots, x_{j-1,1} \mapsto 0, x_{j,1} \mapsto 1]$ for some $j \in [n], j \neq 1$. The Prover then pebbles the clause $C_j = \bigvee_{k \in [n], k \neq j} x_{k,j}$. Let β be the assignment chosen by the Delayer. β is of the form $[x_{1,j} \mapsto 0, \dots, x_{k-1,j} \mapsto 0, x_{k,j} \mapsto 1]$ for some $k \in [n], k \neq j$.

Now if $k < j$, then $\alpha(x_{k,1}) = 0$. But then all literals in the clause $(\neg x_{k,j} \vee \neg x_{j,1} \vee x_{k,1})$ are false, and the Prover can pebble this clause and win.

Assume then that $k > j$. The Prover then pebbles the clause $\neg x_{k,j} \vee \neg x_{j,k}$. Since $\beta(x_{k,j}) = 1$, The Delayer must answer with the assignment γ setting $x_{j,k}$ to 0 (and obviously $x_{k,j}$ to 1). At this point the Prover removes the pebbles from C_1 and C_j and places a pebble on C_k . The Delayer must answer with an assignment δ of the form $[x_{1,k} \mapsto 0, \dots, x_{l-1,k} \mapsto 0, x_{l,k} \mapsto 1]$ where clearly $l > j$, to not contradict the assignment γ . \square

5.3 Automatic Generation of Refutations

These results show that the *ordered structure*, via theorem 20 and lemma 19, gives rise to a subsystem of daglike resolution (corresponding to a strategy for the Prover with $O(1)$ pebbles): (1) powerful enough to obtain polynomial size refutations of important families of contradictions like LOP_n ; and (2) where PHP (and other classes of formulas we omit in this version) are hard to refute and the hardness proof is relatively easy.

Since LOP_n is an example of formulas known to be hard for many automatic theorem provers (see [6]), the previous properties suggest that it is worth investigating algorithms for generating winning strategies for the Prover, as this gives a way of generating resolution refutations.

We present such an algorithm below. It is analogous to the Ben-Sasson Wigderson algorithm based on width. In our case, rather than limiting the width, we limit the number of pebbles used in the strategy.

We first describe a subroutine. The input is a formula F , and the description of the structure \mathcal{S}_F (this may be given in a simple way, e.g. as an ordering if we are dealing with the ordered game), and a number k of pebbles. The output is a winning strategy for the Prover using k pebbles, if one exists, otherwise “NONE”.

The subroutine first builds up a dag A as follows: The nodes of A are all of the positions possible in the game, ie. all the possible ways of labelling k or fewer pebbles plus a source node. There are $\leq O(m^k |\mathcal{S}_F|^k)$ of them, where m is the

number of clauses in F . At the start of the algorithm all the positions that are self contradictory or falsify clauses of F are *marked*, and the graph has no edges. While the source node has not been marked the algorithm does the following: it checks each not marked node X in A . If by removing a pebble the Prover can move from X to a node Y already marked in A , then the algorithm marks X and labels it with the pebble to be removed and adds an edge from X to Y . If there is a clause C that is not pebbled at X , and is such that whatever label the Delayer could choose to give to C , it would lead to a position corresponding to a node already marked in A , then the algorithm marks X , labels it with C , and adds edges going to all the nodes corresponding to the Delayer's possible answers.

The subroutine stops when the source node has been marked, or when there are no more edges to be added to the graph. If the source node has been marked, the subroutine outputs only the marked subgraph of A , which is a winning strategy for the Prover. Otherwise it outputs NONE.

The proof search algorithm works by calling this subroutine for increasing values of k , until the subroutine outputs a strategy (which it will do eventually, when the Prover is able to query all clauses at once).

The running time of the algorithm is $O(m^{r+1}|\mathcal{S}_F|^{r+1})$ where r is the minimal number of pebbles used by the Prover to win $SWG(F, \mathcal{S}_F)$.

For feasible structures, using theorem 20, this algorithm allows us to generate daglike resolution refutations of size polynomial in the size of the formula.

For the case of ordered structures we could add a preprocessing phase to try all possible orders. Although in the worst case this can increase the running time to exponential we notice that for LOP_n , a good ordering is the one described in theorem 26, and this (or one equally good) arises very naturally from the first order combinatorial principle corresponding to the LOP_n formula, together with *any* ordering of the numbers $1, \dots, n$.

Moreover we observe also that although the full structure is not feasible in general, it become so in the case of formulas with $O(1)$ initial width. Now by theorem 4, lemma 10 and theorem 22, if F has width k refutations in resolution then F is $(k + 1, \mathcal{F})$ -easy, so if there is a constant width formula with a narrow refutation then our algorithm (looking for strategies for the full game) works at least as well on it as the Ben-Sasson and Wigderson algorithm.

6 Satisfiability vs Hardness

In Section 4 we showed how the existence of a satisfiable formula G similar to F gives a good strategy for the Delayer in a certain game, which shows that F has no narrow resolution refutation.

This suggests an attractive idea, that we should look for a sort of soundness and completeness theorem for polynomial size resolution. It would have the following form: a CNF F has no small resolution refutation if and only if F is "almost" satisfiable, in that it is hard to distinguish from a satisfiable formula G .

In this section we give two results in this direction. We first show that if F is hard to distinguish from G in the sense that it is hard to prove in resolution that F and G are different, then F is hard to refute. We then show a converse, although this talks about games rather than proofs: if there is a good strategy for the Delayer in the full (structured) pebble game on F , then there exists a satisfiable CNF G that looks similar to F , in a certain sense.

Definition 27 *Let F and G be CNFs, considered as two-sorted structures. We define a new CNF, $ISO(F, G)$, that expresses the statement “ F is isomorphic to G ”. This is intended to be used when F is not isomorphic to G , as a generalization of the pigeonhole principle. Let $Var(F)$ and $Cl(F)$ be the sets of variables and clauses in a CNF.*

We take the conjunction of

1. $\bigvee_{D \in Cl(G)} \sigma_{CD}$ for each $C \in Cl(F)$; $\bigvee_{C \in Cl(F)} \sigma_{CD}$ for each $D \in Cl(G)$;
2. $\neg\sigma_{CD} \vee \neg\sigma_{C'D'}$ and $\neg\sigma_{CD} \vee \neg\sigma_{C'D}$ for all $C \neq C' \in Cl(F)$, $D \neq D' \in Cl(G)$;
3. $\bigvee_{y \in Var(G)} \sigma_{xy}$ for each $x \in Var(F)$; $\bigvee_{x \in Var(F)} \sigma_{xy}$ for each $y \in Var(G)$;
4. $\neg\sigma_{xy} \vee \neg\sigma_{xy'}$ and $\neg\sigma_{xy} \vee \neg\sigma_{x'y}$ for all $x \neq x' \in Var(F)$, $y \neq y' \in Var(G)$;
5. If x appears positively in C in F , but y does not appear positively in D in G , then we include the clause $\neg\sigma_{xy} \vee \neg\sigma_{CD}$; similarly for appearing negatively or not appearing at all.

So 1 and 2 say that σ is a bijection on clauses, 3 and 4 say it is a bijection on variables, and 5 says it preserves the structure.

Theorem 28 *Suppose that G is satisfiable and F has a small resolution refutation. Then $ISO(F, G)$ has a small resolution refutation.*

Proof (sketch) Let α be a satisfying assignment to G . α partitions $Var(G)$ into a set A of true variables and a set B of false variables. For $x \in Var(F)$, let X be the clause $\bigvee_{y \in A} \sigma_{xy}$ and let \bar{X} be the clause $\bigvee_{z \in B} \sigma_{xz}$. Let F^* be F with every variable x replaced by X and every negated variable $\neg x$ replaced by \bar{X} .

We can derive F^* from $ISO(F, G)$, and then use the small refutation of F to give a small refutation of F^* . \square

To use this to get lower bounds for F , we need lower bounds for $ISO(F, G)$. This is a generalization of the pigeonhole principle, so the many existing tools for showing lower bounds for PHP may be useful here. The most tractable case should be when F and G both arise as translations of some first order combinatorial principle, but on structures of slightly different sizes, so we can use the hardness of PHP directly. Krajicek gives an argument like this in [18], relating the proof complexity of the weak pigeonhole principle and the Ramsey theorem. Potentially this will give a sufficient criterion for a formula to have no small daglike refutation, similar to Riis’ criterion for treelike refutations [21].

So far we have given “soundness” theorems, that if F is almost a satisfiable CNF then F is hard to refute (or that it is hard for the Prover to win some game). Now we give a “completeness” theorem.

We think of CNFs as two sorted structures with a clause sort and a variable sort. There are two binary relations, “ x appears positively in C ” and “ x appears negatively in C .”

We define the sense in which our constructed CNF G will be similar to F . Informally, we have “local” embeddings of the clauses of F into the clauses of G . (If we had a global embedding, then the satisfying assignment for G would immediately give a satisfying assignment for F .)

Definition 29 *Let F and G be CNFs. We say that F is strongly k -clause embeddable in G if there is a family H of partial isomorphisms from F to G with the following properties:*

1. *Suppose $f \in H$ maps clauses C_1, \dots, C_k and no other clauses. Then the variables mapped by f are precisely the variables appearing in C_1, \dots, C_k . In other words, f is first of all a partial isomorphism on clauses, but brings with it a partial isomorphism on the variables appearing in those clauses.*
2. *If $f \in H$ maps fewer than k clauses, and C is any other clause in F , then there is $g \in H$ that extends f and maps C somewhere.*

Lemma 30 *If F is strongly k -clause embeddable in G and G is satisfiable, then the Delayer wins the k -pebble full game on F .*

We could apply the following theorem to PHP_n^{n+1} (with $k = n/2 - 1$). Unfortunately the formula G that it gives does not seem to be as elegant as PHP_n^n .

Theorem 31 *Suppose the Delayer wins the k -pebble full game on F . Then there is a satisfiable finite CNF G such that F is strongly k -clause embeddable in G .*

Proof We will build G out of the Delayer’s strategy for the game. Think of this strategy as a dag consisting of positions in the game. Each position P is a tuple $(C_1, \alpha_1), \dots, (C_r, \alpha_r)$ of at most k clauses from F together with assignments to all the variables appearing in each clause, that are consistent and satisfy each clause. Notice that only a finite number of different positions appear in the strategy.

Let G' be the CNF whose clauses are all the pairs (C, α) that appear in the Delayer’s strategy. We give variables to the clauses as follows: if $C \in F$ contains variables x_1, \dots, x_m , then (C, α) contains variables $(x_1, C, \alpha), \dots, (x_m, C, \alpha)$ and (x_i, C, α) appears positively or negatively in (C, α) just as x_i appears positively or negatively in C . That is, G' consists of clauses named by pairs (C, α) and variables named by triples (x, C, α) ; each clause is isomorphic to some clause from F , and no two clauses share any variables. Now define an equivalence relation \sim on the variables of G' as the transitive closure of the relation: (x, C, α) is related to (y, D, β) if and only if x and y are the same variable (in F) and (C, α) and (D, β) appear together in some position in the Delayer’s strategy.

Let G be the CNF obtained from G' by renaming every variable (x, C, α) with its \sim -equivalence class $[(x, C, \alpha)]$.

The proof is completed by showing that G is satisfiable, and that F is strongly k -clause embeddable in G . \square

References

1. M. Alekhovich, E. Ben-Sasson, A. Razborov, A. Wigderson. Space complexity in propositional calculus. *SIAM J. Comput.* 31(4) pp. 1184–1211, 2002.
2. M. Alekhovich, A.A. Razborov. Resolution is not automatizable unless $W[P]$ is not tractable. *42nd IEEE Symposium on Foundations of Computer Science, FOCS 2001*, pp. 210–219.
3. A. Atserias, V. Dalmau. A combinatorial characterization of Resolution Width *18th IEEE Conference on Computational Complexity (CCC)*, pp. 239–247, 2003.
4. A. Atserias, P. Kolaitis, M. Vardi. Constraint Propagation as a Proof System. In *10th International Conference on Principles and Practice of Constraint Programming (CP)*, LNCS vol. 3258, pp. 77–91, 2004.
5. P. Beame, R.M. Karp, T. Pitassi, M.E. Saks. On the Complexity of Unsatisfiability Proofs for Random k -CNF Formulas. *SIAM J. Comput.* 31(4) pp. 1048–1075, 2002.
6. P. Beame, H. Kautz, A. Sabharwal Understanding the power of clause learning *Proceedings IJCAI* pp. 1194–1201, 2003
7. P. Beame, T. Pitassi. Simplified and Improved Resolution Lower Bounds. *37th IEEE Symposium on Foundations of Computer Science, FOCS 1996*, pp. 274–282.
8. E. Ben-Sasson, N. Galesi. Space Complexity of Random Formulae in Resolution. *16th IEEE Annual Conference on Computational Complexity, CCC 2001*, pp. 42–51.
9. E. Ben-Sasson, R. Impagliazzo, A. Wigderson. Near optimal separation of tree-like and general Resolution. *Electronic Colloquium on Computational Complexity (ECCC) TR00-005, 2000*. To appear in *Combinatorica*.
10. E. Ben-Sasson, A. Wigderson. Short Proofs Are Narrow—Resolution Made Simple. *J. ACM* 48(2) pp. 149–168, 2001.
11. M.L. Bonet, N. Galesi. Optimality of Size-Width Tradeoffs for Resolution. *Computational Complexity*, Vol 10(4) 2001. pp. 261–276.
12. J.L. Esteban, N. Galesi, J. Messner. On the Complexity of Resolution with Bounded Conjunctions. *Theoretical Computer Science* 321(2-3) pp. 347–370, 2004.
13. J.L. Esteban, J. Torán. Space bounds for Resolution. *Inform. and Comput.* 171 (1) pp. 84–97, 2001.
14. N. Galesi, N. Thapen. The Complexity of Treelike Systems over λ Local Formulae *Proceedings of IEEE Conference on Computational Complexity* 2004.
15. N. Galesi, N. Thapen. Resolution and Pebbling Games *ECCC Technical Report TR04-112*. <http://www.eccc.uni-trier.de/eccc-reports/2004/TR04-112/index.html>
16. A. Haken. The Intractability of Resolution. *Theoret. Comp. Sci.* 39, pp. 297–308, 1985.
17. J. Krajíček. Bounded arithmetic, propositional logic, and complexity theory, *Encyclopedia of Mathematics and Its Applications*, Vol. **60**, *Cambridge University Press*,(1995),
18. J. Krajíček. On the weak pigeonhole principle. *Fund. Math.* 170(1-3) pp. 123–140, 2001.
J. Symbolic Logic 59(1) pp. 73–86, 1994.
19. P. Pudlak. Proofs as Games. *American Math. Monthly*, Vol. 2000-2001, pp.541-550
20. P. Pudlák, R. Impagliazzo. A lower bound for DLL algorithms for k -SAT”. *Conference Proceeding of Symposium on Distributed Algorithms* (2000), pp. 128-136.
21. S. Riis. A complexity gap for tree-resolution. *Computational Complexity* 10(3), pp. 179-209, 2001.
22. A. Urquhart. Hard examples for Resolution. *J. ACM* 34(1) pp. 209-219, 1987.