



A lower bound for the pigeonhole principle in tree-like Resolution by asymmetric Prover–Delayer games [☆]

Olaf Beyersdorff ^{a,*}, Nicola Galesi ^b, Massimo Lauria ^b

^a Institut für Theoretische Informatik, Leibniz Universität Hannover, Germany

^b Dipartimento di Informatica, Sapienza Università di Roma, Italy

ARTICLE INFO

Article history:

Received 1 July 2010

Accepted 17 September 2010

Available online 21 September 2010

Communicated by J. Torán

Keywords:

Computational complexity

Proof complexity

Prover–Delayer games

Resolution

ABSTRACT

In this note we show that the asymmetric Prover–Delayer game developed in Beyersdorff et al. (2010) [2] for Parameterized Resolution is also applicable to other tree-like proof systems. In particular, we use this asymmetric Prover–Delayer game to show a lower bound of the form $2^{\Omega(n \log n)}$ for the pigeonhole principle in tree-like Resolution. This gives a new and simpler proof of the same lower bound established by Iwama and Miyazaki (1999) [7] and Dantchev and Riis (2001) [5].

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Proving lower bounds by games is a very fruitful technique in proof complexity [9,11,10,1]. In particular, the Prover–Delayer game of Pudlák and Impagliazzo [10] is one of the canonical tools to study lower bounds in tree-like Resolution [10,3] and tree-like $Res(k)$ [6]. The Prover–Delayer game of Pudlák and Impagliazzo arises from the well-known fact [8] that a tree-like Resolution proof for a formula F can be viewed as a decision tree which solves the search problem of finding a clause of F falsified by a given assignment. In the game, Prover queries a variable and Delayer either gives it a value or leaves the decision to Prover and receives one point. The number of Delayer's points at the end of the game is then proportional to the height of the proof tree. It is easy to argue that showing lower bounds by this game only works if (the graph of)

every tree-like Resolution refutation contains a balanced sub-tree as a minor, and the height of that sub-tree then gives the size lower bound.

In [2] we developed a new asymmetric Prover–Delayer game which extends the game of Pudlák and Impagliazzo to make it applicable to obtain lower bounds to tree-like proofs when the proof trees are very unbalanced. In [2] we used the new asymmetric game to obtain lower bounds in tree-like Parameterized Resolution, a proof system in the context of parameterized proof complexity recently introduced by Dantchev, Martin, and Szeider [4]. The lower bounds we obtain in [2] for tree-like Parameterized Resolution are of the form $\Omega(n^k)$ (n is the formula size and k the parameter), but the tree-like Parameterized Resolution refutations of the formulas in question only contain balanced sub-trees of height k .

The aim of this note is to show that the asymmetric Prover–Delayer game is also applicable to other (non-parameterized) tree-like proof systems. One of the best studied principles is the pigeonhole principle. Iwama and Miyazaki [7] and independently Dantchev and Riis [5] show that the pigeonhole principle requires tree-like Resolution refutations of size roughly $n!$ while its tree-like Resolution proofs only contain balanced sub-trees of height n . Therefore the game of Pudlák and Impagliazzo only yields

[☆] This work was done while the first author was visiting Sapienza University of Rome under support of DFG grant KO 1053/5-2. The second author was supported by grant “Limiti di compressione in combinatoria e complessità computazionale” by Sapienza University Rome.

* Corresponding author.

E-mail addresses: beyersdorff@thi.uni-hannover.de (O. Beyersdorff), galesi@di.uniroma1.it (N. Galesi), lauria@di.uniroma1.it (M. Lauria).

a $2^{2^{\Omega(n)}}$ lower bound which is weaker than the optimal bound $2^{2^{\Omega(n \log n)}}$ established by Iwama and Miyazaki. Here we provide a new and easier proof of this lower bound by our asymmetric Prover–Delayer game.

2. Preliminaries

A *literal* is a positive or negated propositional variable and a *clause* is a set of literals. A clause is interpreted as the disjunctions of its literals and a set of clauses as the conjunction of the clauses. Hence clause sets correspond to formulas in CNF. The *Resolution system* is a refutation system for the set of all unsatisfiable CNF. Resolution uses as its only rule the *Resolution rule*

$$\frac{\{x\} \cup C \quad \{\neg x\} \cup D}{C \cup D}$$

for clauses C, D and a variable x . The aim in Resolution is to demonstrate unsatisfiability of a clause set by deriving the empty clause. If in a derivation every derived clause is used at most once as a prerequisite of the Resolution rule, then the derivation is called *tree-like*, otherwise it is *dag-like*. The size of a Resolution proof is the number of its clauses. Undoubtedly, Resolution is the most studied and best-understood propositional proof system (cf. [12]).

It is well known (cf. [8]) that a tree-like refutation of F can equivalently be described as a *boolean decision tree*. A boolean decision tree for F is a binary tree where inner nodes are labeled with variables from F and leaves are labeled with clauses from F . Each path in the tree corresponds to a partial assignment where a variable x gets value 0 or 1 according to whether the path branches left or right at the node labeled with x . The condition on the decision tree is that each path α must lead to a clause which is falsified by the assignment corresponding to α . Therefore, a boolean decision tree solves the *search problem* for F which, given an assignment α , asks for a clause from F falsified by α . It is easy to verify that each tree-like Resolution refutation of F yields a boolean decision tree for F and vice versa, where the size of the Resolution proof equals the number of nodes in the decision tree. In the sequel, we will therefore concentrate on boolean decision trees to prove our lower bound to tree-like Resolution.

3. Tree-like lower bounds via asymmetric Prover–Delayer games

We review the asymmetric Prover–Delayer game from [2]. Let F be a set of clauses in n variables x_1, \dots, x_n . In the asymmetric game, Prover and Delayer build a (partial) assignment to x_1, \dots, x_n . The game is over as soon as the partial assignment falsifies a clause from F . The game proceeds in rounds. In each round, Prover suggests a variable x_i , and Delayer either chooses a value 0 or 1 for x_i or leaves the choice to the Prover. In this last case, if the Prover sets the value, then the Delayer gets some points. The number of points Delayer earns depends on the variable x_i , the assignment α constructed so far in the game, and two functions $c_0(x_i, \alpha)$ and $c_1(x_i, \alpha)$. More precisely, the number of points that Delayer will get is

$$\begin{aligned} & 0 && \text{if Delayer chooses the value,} \\ & \log c_0(x_i, \alpha) && \text{if Prover sets } x_i \text{ to 0, and} \\ & \log c_1(x_i, \alpha) && \text{if Prover sets } x_i \text{ to 1.} \end{aligned}$$

Moreover, the functions $c_0(x, \alpha)$ and $c_1(x, \alpha)$ are chosen in such a way that for each variable x and assignment α

$$\frac{1}{c_0(x, \alpha)} + \frac{1}{c_1(x, \alpha)} = 1 \tag{1}$$

holds. Let us call this game the (c_0, c_1) -game on F .

The connection of this game to size of proofs in tree-like Resolution is given by Theorem 1. The theorem is essentially contained in [2], but for completeness we include the full proof.

Theorem 1. (See [2].) *Let F be unsatisfiable formula in CNF and let c_0 and c_1 be two functions satisfying (1) for all partial assignments α to the variables of F . If F has a tree-like Resolution refutation of size at most S , then the Delayer gets at most $\log S$ points in each (c_0, c_1) -game played on F .*

Proof. Let F be an unsatisfiable CNF in variables x_1, \dots, x_n and let Π be a tree-like Resolution refutation of F . Assume now that Prover and Delayer play a game on F where they successively construct an assignment α . Let α_i be the partial assignment constructed after i rounds of the game, i.e., α_i assigns i variables a value 0 or 1. By p_i we denote the number of points that Delayer has earned after i rounds, and by Π_{α_i} we denote the sub-tree of the decision tree of Π which has as its root the node reached in Π along the path specified by α_i .

We use induction on the number of rounds in the game to prove the following claim:

$$|\Pi_{\alpha_i}| \leq \frac{|\Pi|}{2^{p_i}} \quad \text{for any round } i.$$

To see that the theorem follows from this claim, let α be an assignment constructed during the game yielding p_α points to the Delayer. As a contradiction has been reached in the game, the size of Π_α is 1, and therefore by the inductive claim

$$1 \leq \frac{|\Pi|}{2^{p_\alpha}},$$

yielding $p_\alpha \leq \log |\Pi|$ as desired.

In the beginning of the game, Π_{α_0} is the full tree and the Delayer has 0 points. Therefore the claim holds.

For the inductive step, assume that the claim holds after i rounds and Prover asks for a value of the variable x in round $i + 1$. If the Delayer chooses the value, then $p_{i+1} = p_i$ and hence

$$|\Pi_{\alpha_{i+1}}| \leq |\Pi_{\alpha_i}| \leq \frac{|\Pi|}{2^{p_i}} = \frac{|\Pi|}{2^{p_{i+1}}}.$$

If the Delayer defers the choice to the Prover, then the Prover uses the following strategy to set the value of x . Let $\alpha_i^{x=0}$ be the assignment extending α_i by setting x to 0, and let $\alpha_i^{x=1}$ be the assignment extending α_i by setting x to 1. Now, Prover sets $x = 0$ if $|\Pi_{\alpha_i^{x=0}}| \leq \frac{1}{c_0(x, \alpha_i)} |\Pi_{\alpha_i}|$, otherwise he sets $x = 1$. Because $\frac{1}{c_0(x, \alpha_i)} + \frac{1}{c_1(x, \alpha_i)} = 1$, we know

that if Prover sets $x = 1$, then $|\Pi_{\alpha_i^{x=1}}| \leq \frac{1}{c_1(x, \alpha_i)} |\Pi_{\alpha_i}|$. Thus, if Prover's choice is $x = j$ with $j \in \{0, 1\}$, then we get

$$\begin{aligned} |\Pi_{\alpha_{i+1}}| &= |\Pi_{\alpha_i^{x=j}}| \leq \frac{|\Pi_{\alpha_i}|}{c_j(x, \alpha_i)} \leq \frac{|\Pi|}{c_j(x, \alpha_i) 2^{p_i}} \\ &= \frac{|\Pi|}{2^{p_i + \log c_j(x, \alpha_i)}} = \frac{|\Pi|}{2^{p_{i+1}}}. \end{aligned}$$

This completes the proof of the induction. \square

As remarked in [2] we get the game of Pudlák and Impagliazzo [10] by setting $c_0(x, \alpha) = c_1(x, \alpha) = 2$ for all variables x and partial assignments α .

4. Tree-like resolution lower bounds for the pigeonhole principle

The *weak pigeonhole principle* PHP_n^m with $m > n$ uses variables $x_{i,j}$ with $i \in [m]$ and $j \in [n]$, indicating that pigeon i goes into hole j . PHP_n^m consists of the clauses

$$\bigvee_{j \in [n]} x_{i,j} \quad \text{for all pigeons } i \in [m]$$

and $\neg x_{i_1,j} \vee \neg x_{i_2,j}$ for all choices of distinct pigeons $i_1, i_2 \in [m]$ and holes $j \in [n]$. We prove that PHP_n^m is hard for tree-like Resolution. Showing the lower bound by the asymmetric game from the last section, requires a suitable choice of the functions c_0 and c_1 and then the definition of the Delayer-strategy for the (c_0, c_1) -game.

Theorem 2. Any tree-like Resolution refutation of PHP_n^m has size $2^{\Omega(n \log n)}$.

Proof. Let α be a partial assignment to the variables $\{x_{i,j} \mid i \in [m], j \in [n]\}$. Let

$$\begin{aligned} p_i(\alpha) &= |\{j \in [n] \mid \alpha(x_{i,j}) = 0 \text{ and} \\ &\alpha(x_{i',j}) \neq 1 \text{ for all } i' \in [m]\}|. \end{aligned}$$

Intuitively, $p_i(\alpha)$ corresponds to the number of holes which are still free but are explicitly excluded for pigeon i by α (we do not count the holes which are excluded because some other pigeon is sitting there). We define

$$\begin{aligned} c_0(x_{i,j}, \alpha) &= \frac{\frac{n}{2} + 1 - p_i(\alpha)}{\frac{n}{2} - p_i(\alpha)} \quad \text{and} \\ c_1(x_{i,j}, \alpha) &= \frac{n}{2} + 1 - p_i(\alpha). \end{aligned}$$

For simplicity we assume that n is divisible by 2. During the game it will never be the case that Prover gets the choice when $p_i(\alpha) \geq \frac{n}{2}$. Therefore the functions c_0 and c_1 are always greater than zero when the Delayer gets points, thus the score function is always well defined. Furthermore notice that this definition satisfies (1).

We now describe Delayer's strategy in a (c_0, c_1) -game played on PHP_n^m . If Prover asks for a value of $x_{i,j}$, then Delayer decides as follows:

set $\alpha(x_{i,j}) = 0$ if there exists $i' \in [m] \setminus \{i\}$ such that $\alpha(x_{i',j}) = 1$ or if there exists $j' \in [n] \setminus \{j\}$ such that $\alpha(x_{i,j'}) = 1$;
 set $\alpha(x_{i,j}) = 1$ if $p_i(\alpha) \geq \frac{n}{2}$ and there is no $i' \in [m]$ with $\alpha(x_{i',j}) = 1$, and
 let Prover decide otherwise.

Intuitively, Delayer leaves the choice to Prover as long as pigeon i does not already sit in a hole, hole j is still free, and there are at most $\frac{n}{2}$ excluded free holes for pigeon i .

Let us pause to give an intuitive explanation of why we choose the functions c_0 and c_1 and thus the points for Delayer as above. As a first observation, Delayer always earns more when Prover is setting a variable $x_{i,j}$ to 1 instead of setting it to 0. This is intuitively correct as the amount of freedom for Delayer to continue the game is by far more diminished by sending pigeon i to some hole j than by just excluding that hole j for pigeon i . In fact, our choice of scores can be completely explained by the following information-theoretic interpretation: When Prover sends a pigeon to a hole, Delayer should always get about $\log n$ points on that pigeon. For our Delayer strategy, sending pigeon i to a hole either means that Prover excluded $\frac{n}{2}$ holes for pigeon i or was setting pigeon i directly to a hole. When we play the game, in each round Delayer should get some number of points proportional to the progress Prover made towards fixing pigeon i to a hole. For instance, if Prover fixes i to a hole in the very beginning by answering 1 to $x_{i,j}$, Delayer should get the $\log n$ points immediately. On the other extreme, if Prover has already excluded $\frac{n}{2} - 1$ holes for pigeon i , then it does not matter whether Prover sets $x_{i,j}$ to 0 or 1 because after both answers pigeon i will be forced to a hole. Consequently, in the latter case, Delayer gets just 1 point regardless of whether Prover answers 0 or 1. This is exactly what our score function provides.

If Delayer uses the above strategy, then the small clauses $\neg x_{i_1,j} \vee \neg x_{i_2,j}$ from PHP_n^m will not be violated in the game. Therefore, a contradiction will always be reached on one of the big clauses $\bigvee_{j \in [n]} x_{i,j}$. Let us assume now that the game ends by violating $\bigvee_{j \in [n]} x_{i,j}$, i.e. for pigeon i all variables $x_{i,j}$ with $j \in [n]$ have been set to 0. As soon as the number $p_i(\alpha)$ of excluded free holes for pigeon i reaches the threshold $\frac{n}{2}$, Delayer will not leave the choice to Prover. Instead, Delayer will try to place pigeon i into some hole. If Delayer still answers 0 to $x_{i,j}$ even after $p_i(\alpha) > \frac{n}{2}$, it must be the case that some other pigeon already sits in hole j , i.e. for some $i' \neq i$, $\alpha(x_{i',j}) = 1$. Therefore, at the end of the game at least $\frac{n}{2}$ variables have been set to 1. W.l.o.g. we assume that these are the variables x_{i,j_i} for $i = 1, \dots, \frac{n}{2}$.

Let us check how many points Delayer earns in this game. We calculate the points separately for each pigeon $i = 1, \dots, \frac{n}{2}$ and distinguish two cases: whether x_{i,j_i} was set to 1 by Delayer or Prover. Let us first assume that Delayer sets the variable x_{i,j_i} to 1. Then pigeon i was not assigned to a hole yet and, moreover, there must be $\frac{n}{2}$ unoccupied holes which are already excluded for pigeon i by α , i.e., there is some $J \subseteq [n]$ with $|J| = \frac{n}{2}$, $\alpha(x_{i',j'}) \neq 1$

for $i' \in [m]$, $j' \in J$, and $\alpha(x_{i,j'}) = 0$ for all $j' \in J$. All of these 0's have been assigned by Prover, as Delayer has only assigned a 0 to $x_{i,j'}$ when some other pigeon was already sitting in hole j' , and this is not the case for the holes from J (at the moment when Delayer assigns the 1 to x_{i,j_i}). Thus, before Delayer sets $\alpha(x_{i,j_i}) = 1$, she has already earned points for all $\frac{n}{2}$ variables $x_{i,j'}$, $j' \in J$, yielding

$$\sum_{p=0}^{\frac{n}{2}-1} \log \frac{\frac{n}{2} + 1 - p}{\frac{n}{2} - p} = \log \prod_{p=0}^{\frac{n}{2}-1} \frac{\frac{n}{2} + 1 - p}{\frac{n}{2} - p} = \log \left(\frac{n}{2} + 1 \right)$$

points for the Delayer. Let us note that because Delayer never allows a pigeon to go into more than one hole, she will really get the number of points calculated above for every of the variables which she set to 1.

If, conversely, Prover sets variable x_{i,j_i} to 1, then Delayer gets $\log(\frac{n}{2} + 1 - p_i(\alpha))$ points for this, but she also received points for the $p_i(\alpha)$ variables set to 0 before by Prover. Thus, in this case Delayer earns on pigeon i

$$\begin{aligned} & \log \left(\frac{n}{2} + 1 - p_i(\alpha) \right) + \sum_{p=0}^{p_i(\alpha)-1} \log \frac{\frac{n}{2} + 1 - p}{\frac{n}{2} - p} \\ &= \log \left(\frac{n}{2} + 1 - p_i(\alpha) \right) + \log \frac{\frac{n}{2} + 1}{\frac{n}{2} - p_i(\alpha) + 1} \\ &= \log \left(\frac{n}{2} + 1 \right) \end{aligned}$$

points. In total, Delayer gets at least

$$\frac{n}{2} \log \left(\frac{n}{2} + 1 \right)$$

points in the game. Applying Theorem 1, we obtain $2^{\frac{n}{2} \log(\frac{n}{2} + 1)}$ as a lower bound to the size of each tree-like Resolution refutation of PHP_n^m . \square

By inspection of the above Delayer strategy it becomes clear that the lower bound from Theorem 2 also holds for the *functional pigeonhole principle* where in addition to the clauses from PHP_n^m we also include $\neg x_{i,j_1} \vee \neg x_{i,j_2}$ for all pigeons $i \in [m]$ and distinct holes $j_1, j_2 \in [n]$.

We remark that the choice of the score functions c_0 and c_1 in the proof of Theorem 2 is by no means unique. It is even possible to obtain the same asymptotic lower bound $2^{\Omega(n \log n)}$ by choosing simpler score functions c_0, c_1 which do not depend on the game played so far, i.e. c_0 and c_1 just depend on n , but are independent of the assignment α and the queried variable x . Namely, setting

$$c_1 = \frac{n}{\log n} \quad \text{and}$$

$$c_0 = \frac{c_1}{c_1 - 1} = 1 + \frac{1}{c_1 - 1} = \Omega \left(e^{\frac{1}{c_1 - 1}} \right) = 2^{\Omega \left(\frac{\log n}{n} \right)}$$

we obtain score functions which satisfy (1) and lead to the following modified analysis in the proof of Theorem 2: if the Prover sets $x_{i,j}$ to 1, then Delayer earns at least $\log c_1 = \Omega(\log n)$ points. Otherwise, she still earns at least $\frac{n}{2} \log c_0 = \Omega(\log n)$ points on pigeon i . Thus, in total Delayer earns $\frac{n}{2} \cdot \Omega(\log n)$ points during the game, yielding the lower bound.

Our first proof of Theorem 2 has the advantage that it yields more precise and better bounds, namely exactly $2^{\frac{n}{2} \log(\frac{n}{2} + 1)}$ which is the same lower bound obtained by Dantchev and Riis [5]. There might also be scenarios where the adaptive definition of points according to our above information-theoretic interpretation indeed yields better asymptotic bounds.

Acknowledgements

We thank Stasys Jukna for suggesting the simplified choice of the score function explained after the proof of Theorem 2 and Jan Johannsen for pointing out to us the Ref. [7].

References

- [1] Albert Atserias, Víctor Dalmau, A combinatorial characterization of resolution width, *Journal of Computer and System Sciences* 74 (3) (2008) 323–334.
- [2] Olaf Beyersdorff, Nicola Galesi, Massimo Lauria, Hardness of parameterized resolution, Technical Report TR10-059, Electronic Colloquium on Computational Complexity, 2010.
- [3] Eli Ben-Sasson, Russell Impagliazzo, Avi Wigderson, Near optimal separation of tree-like and general resolution, *Combinatorica* 24 (4) (2004) 585–603.
- [4] Stefan S. Dantchev, Barnaby Martin, Stefan Szeider, Parameterized proof complexity, In: *Proc. 48th IEEE Symposium on the Foundations of Computer Science*, 2007, pp. 150–160.
- [5] Stefan S. Dantchev, Søren Riis, Tree resolution proofs of the weak pigeon-hole principle, In: *Proc. 16th Annual IEEE Conference on Computational Complexity*, 2001, pp. 69–75.
- [6] Juan Luis Esteban, Nicola Galesi, Jochen Messner, On the complexity of resolution with bounded conjunctions, *Theoretical Computer Science* 321 (2–3) (2004) 347–370.
- [7] Kazuo Iwama, Shuichi Miyazaki, Tree-like resolution is superpolynomially slower than dag-like resolution for the pigeonhole principle, in: *Proc. 10th International Symposium on Algorithms and Computation*, in: *Lecture Notes in Computer Science*, vol. 1741, Springer-Verlag, Berlin, Heidelberg, 1999, pp. 133–142.
- [8] Jan Krajíček, Bounded Arithmetic, Propositional Logic, and Complexity Theory, *Encyclopedia of Mathematics and Its Applications*, vol. 60, Cambridge University Press, Cambridge, 1995.
- [9] Pavel Pudlák, Samuel R. Buss, How to lie without being (easily) convicted and the length of proofs in propositional calculus, in: *Proc. 8th Workshop on Computer Science Logic*, in: *Lecture Notes in Computer Science*, vol. 933, Springer-Verlag, Berlin, Heidelberg, 1994, pp. 151–162.
- [10] Pavel Pudlák, Russell Impagliazzo, A lower bound for DLL algorithms for SAT, In: *Proc. 11th Symposium on Discrete Algorithms*, 2000, pp. 128–136.
- [11] Pavel Pudlák, On the complexity of propositional calculus, in: *Sets and Proofs, Invited papers from Logic Colloquium'97*, Cambridge University Press, 1999, pp. 197–218.
- [12] Nathan Segerlind, The complexity of propositional proofs, *Bulletin of symbolic Logic* 13 (4) (2007) 482–537.