

Resolution width-size trade-offs for the Pigeon-Hole Principle

Stefan Dantchev

dantchev@dcs.qmul.ac.uk

BRICS¹, Dept. of Computer Science, University of Aarhus
Dept. of Computer Science, Queen Mary, University of London

Abstract

We prove the following two results:

(1) There is a resolution proof of the Weak Pigeon-Hole Principle, $WPHP_n^m$ of size $2^{O\left(\frac{n \log n}{\log m} + \log m\right)}$ for any number of pigeons m and any number of holes n .

(2) Any resolution proof of $WPHP_n^m$ of width $(\frac{1}{16} - \varepsilon) n^2$ has to be of size $2^{\Omega(n)}$, independently from m .

These results give not only a resolution size-width tradeoff for the Weak Pigeon-Hole Principle, but also almost optimal such trade-off for resolution in general. The upper bound (1) may be of independent interest, as it has been known for the two extreme values of m , $m = n + 1$ and $m = 2^{\sqrt{n \log n}}$, only.

1 Introduction

The Pigeon-Hole Principle is one of the simplest and, at the same time, the most important combinatorial principles. Its traditional formulation, denoted usually by PHP_n^m , states that there is no *injective* map from a finite m -element set into a finite n -element set if $m > n$. The Pigeon-Hole Principle also has the distinction to be the first and, since then, the most used combinatorial statement in propositional proof complexity. More specifically, it has been used quite a number of times in proving *lower bounds* for concrete *propositional proof systems*.

As the present paper is concerned with *resolution proofs* of PHP_n^m , we will briefly survey the history of proving resolution lower bounds for the Pigeon-Hole Principle. Everything started with the seminal Haken's result [5], that *any resolution proof* of PHP_n^{n+1} is of size $2^{\Omega(n)}$. The proof has been generalised and simplified in [4], [1], [2]. For quite a while, the best known result had been a $2^{\Omega(n^2/m)}$ lower bound from [4], thus having left the case $m = \Omega(n^2/\log n)$ as an important open problem in resolution proof complexity. Recently, a $2^{\Omega(n^\varepsilon)}$ lower bound on *any regular-resolution proof* of PHP_n^m has appeared in [6]. Shortly after that, the problem has finally been solved by Raz in [9], and further strengthened and improved by Razborov in [10], [11].

All the mentioned proofs use, explicitly or implicitly, the relation between width (or “pseudo-width”, as defined by Razborov) and size of the proof. This relation, used in almost all resolution lower bounds, not only for the Pigeon-Hole Principle,

¹Basic Research In Computer Science, Centre of the Danish National Research Foundation

was extracted and stated precisely in [2]. On the other hand, nobody has studied the question whether there is a width-size trade-off in resolution proofs of some concrete statements. That is, can one pay for decreasing the size by increasing the width? Our paper answers the question as far as the Pigeon-Hole Principle is concerned.

We first prove an upper bound, i.e. construct a resolution proof of $WPHP_n^m$ of size $2^{O(\frac{n \log n}{\log m} + \log m)}$ for any m and n . Such an upper bounds have been known so far only for the two extreme cases $m = n + 1$ and $m = 2^{\sqrt{n \log n}}$ (see [3]). Our result matches these, and, moreover, we believe that it is the exact resolution proof complexity of $WPHP_n^m$ for all m and n .

We also prove a $2^{\Omega(n)}$ lower bound on any resolution proof of the weak pigeon-hole principle, $WPHP_n^m$, when the width is bounded by $(\frac{1}{16} - \varepsilon) n^2$. Unlike the general lower bound in [10], it holds independently from the number of the pigeons m .

These two results not only give a resolution width-size trade-off for the Weak Pigeon-Hole Principle, but also have the following interesting consequence. Letting $m = 2^{\sqrt{n \log n}}$, we get a tautology on N variables which is provable in resolution within $\text{polylog}(N)$ width. Any such proof however is of super-polynomial in N size. At the same time, there is a proof of $\text{poly}(N)$ size and width N . This is asymptotically, i.e. modulo the degrees of the polynomials, the best resolution width-size trade-off, one could hope to prove.

The rest of the paper is organised as follows. In the section 2 we introduce the concepts, denotations and techniques. In the section 3 we show two trivial results, for the sake of the completeness only. The upper bound for $WPHP_n^m$ is shown in the section 4. The lower bound, when the width is restricted, is proven in the section 5 Finally, in the section 6, we discuss some open questions.

2 Preliminaries

We first introduce some denotations and give some definitions. $[n]$ denotes the set $\{1, 2, \dots, n\}$, and $[n : m]$ denotes $\{n + 1, n + 2, \dots, m\}$.

A *literal* is either a propositional variable or the negation of a propositional variable. A *clause* is a set of literals. It is satisfied by a truth assignment if at least one of its literals is true under this assignment. A set of clauses is *satisfiable* if there exists a truth assignment satisfying all the clauses.

PHP_n^m denotes the claim that there is no *injective map* from a set of size m to a set of size n , where $m > n$. We encode its negation as the following set of clauses

1. $\{p_{i,j} \mid j \in [n]\}$ for every pigeon $i \in [m]$
2. $\{\bar{p}_{i,k}, \bar{p}_{j,k}\}$ for every hole $k \in [n]$ and pigeons $i, j \in [m], i \neq j$

Although we consider the *injective PHP*, all the results and proofs from the paper remain valid for the *functional* (where the map is required to be a function) and *bijective PHP*, too.

Resolution is a proof system designed to *refute* given set of clauses i.e. to prove that it is unsatisfiable. This is done by means of the resolution rule

$$\frac{C_1 \cup \{v\} \quad C_2 \cup \{\bar{v}\}}{C_1 \cup C_2}.$$

Thus, we can derive a new clause from two other clauses that contain a variable and its negation respectively. The goal is to derive the empty clause from the initial ones. Anywhere we say we *prove* some proposition, we mean that first we take

its negation in a clausal form and then resolution is used to refute these clauses. Sometimes, for technical reasons only, we could also use the *weakening* rule

$$\frac{C_1}{C_1 \cup C_2}.$$

There is an obvious way to represent every resolution refutation as a directed acyclic graph whose nodes are labelled by clauses. The sources, i.e. the vertices with no incoming edges, are the initial clauses. The only sink, i.e. the vertex with no outgoing edges, is the empty clause. We then define *the size* of a proof to be the number of internal vertices, i.e. non-leaves, which is equal to the number of applications of the resolution rule. We do not however count the number of weakenings if any, as they are not essential and can be removed from the proof.

A very important technique, we use to prove lower bounds on proofs, is considering a proof as a *Prover-Adversary game*. It is first introduced in [8] and developed further in [7] especially for resolution.

There are two players, named *Prover* and *Adversary*. An unsatisfiable set of clauses is given. Adversary claims wrongly that there is a satisfying assignment. Prover's task is to convict him in lying. A *position* in the game is a partial assignment of the propositional variables. The game start from the empty position. Prover has two kind of moves:

1. She queries a variable, whose value is unknown in the current position. Adversary answers, and the position then is extended with the answer.
2. She forgets a value of a variable, which is known. The current position is then reduced, i.e., the variable value becomes unknown.

The game is over, when the current partial assignment falsifies one of the clauses. Prover then wins, having shown a contradiction.

The link to resolution proofs is easily seen. We take such a proof, reverse all the edges in its graph, and label all the vertices by the negation of the corresponding clauses. What we get is a description of Prover's *winning strategy*. Any node in the new graph corresponds to a position in the game. The negation of the corresponding clause is a conjunction of literals which defines the current partial assignment. The variable to be queried by Prover at this position is the variable resolved in the proof. The leaves in the new graph are negation of the initial clauses, i.e. arriving there means that an assignment falsifying such a clause is found, and therefore Prover wins the game.

A *deterministic* Adversary's strategy corresponds to a *single path* in the proof's graph. Therefore, he has to use a *randomised strategy* (called "super-strategy" in Pudlak's paper) in order to enforce a *big enough subgraph*. Any such Adversary's strategy can be used to prove a lower bound on any Prover's strategy description, and therefore on any resolution proof.

3 Two simple proofs of the Pigeon-Hole Principle.

In this section we construct two resolution proofs. They are stated and proven for the ordinary Pigeon-Hole Principle, PHP_n^{n+1} as one can always prove the Weak Pigeon-Hole Principle by restricting it to the first $n + 1$ pigeons. The first proof shows that $WPHP_n^m$ can be proven in very small width which is also optimal. The second one is the optimal-size proof for the ordinary pigeon-hole principle. We will use it as a basis case in constructing a smaller proof in the case $m \gg n$.

Proposition 3.1 *There is a resolution proof of PHP_n^{n+1} of (optimal) width n and size $2^{O(n \log n)}$.*

Proof The proof contains n stages, numbered from n to 0. For each k , the k -th stage encodes the statement “there is no injective mapping of the first k pigeons into the holes”. The corresponding clauses are

$$\left\{ \bar{p}_{1,\pi(1)}, \bar{p}_{2,\pi(2)}, \dots, \bar{p}_{k,\pi(k)} \right\},$$

where π is any injective function from $[k]$ into $[n]$. Thus the stage 0 consists of the empty clause only. The n -th stage clauses can be derived by consecutively resolving the axiom $\{p_{n+1,1}, p_{n+1,2}, \dots, p_{n+1,n}\}$ with the axioms $\{\bar{p}_{n+1,\pi(1)}, \bar{p}_{1,\pi(1)}\}$, $\{\bar{p}_{n+1,\pi(2)}, \bar{p}_{2,\pi(2)}\}, \dots$ and $\{\bar{p}_{n+1,\pi(n)}, \bar{p}_{n,\pi(n)}\}$ for every injective function $\pi : [n] \rightarrow [n]$.

What remains is to show how to go from the stage $k+1$ to the stage k . Given an injection $\pi : [k] \rightarrow [k]$, we shall derive the clause

$$\left\{ \bar{p}_{1,\pi(1)}, \bar{p}_{2,\pi(2)}, \dots, \bar{p}_{k,\pi(k)} \right\} \quad (1)$$

from axioms and the $k+1$ -th stage clauses

$$\left\{ \bar{p}_{1,\pi(1)}, \bar{p}_{2,\pi(2)}, \dots, \bar{p}_{k,\pi(k)}, \bar{p}_{k+1,j} \right\}, \quad (2)$$

for all $j \in [n] \setminus \text{Range}(\pi)$. The derivation is as follows. We first use the axioms $\{p_{k+1,1}, p_{k+1,2}, \dots, p_{k+1,n}\}$ and $\{\bar{p}_{k+1,\pi(1)}, \bar{p}_{1,\pi(1)}\}, \{\bar{p}_{k+1,\pi(2)}, \bar{p}_{2,\pi(2)}\}, \dots, \{\bar{p}_{k+1,\pi(k)}, \bar{p}_{k,\pi(k)}\}$ to obtain

$$\left\{ \bar{p}_{1,\pi(1)}, \bar{p}_{2,\pi(2)}, \dots, \bar{p}_{k,\pi(k)}, p_{k+1,j_1}, p_{k+1,j_2}, p_{k+1,j_{n-k}} \right\},$$

where $\{j_1, j_2, \dots, j_{n-k}\} = [n] \setminus \text{Range}(\pi)$. We then consecutively resolve it with the clauses (2) to “kill” the indices j_1, j_2, \dots, j_{n-k} and finally get the desired clause (1).

All the clauses used in the proof are of width less than or equal to n . The size can be estimated as follows. For every k , the k -th stage contains $\frac{n!}{(n-k)!}$ clauses, and in deriving each k -stage clause we have used n resolution steps. The overall number of resolution steps therefore is

$$n \sum_{k=0}^n \frac{n!}{(n-k)!} \sim en n! = 2^{O(n \log n)}.$$

□

Proposition 3.2 *There is a resolution proof of PHP_n^{n+1} of (optimal) size $2^{n+O(\log n)}$ and width $\frac{1}{4}n^2 + O(n)$*

Proof The proof contains n stages, the k -th one encoding the statement “there is no injective mapping of the first $k+1$ pigeons into any k -element subset of $[n]$ ”. The corresponding set of clauses is

$$P_{1,S} \cup P_{2,S} \cup \dots \cup P_{k+1,S},$$

where S is any $n-k$ -element subset of $[n]$, and $P_{i,S}$ is defined as $P_{i,S} := \{p_{i,j} \mid j \in S\}$. Thus the stage 0 consists of the single axiom $\{p_{1,1}, p_{1,2}, \dots, p_{1,n}\}$, and the last stage, the n -th one, consists of the empty clause only.

What remains is to show how to go from the stage $k-1$ to the stage k . Given any $n-k$ -element set $S \subseteq [n]$ we shall derive the clause

$$P_{1,S} \cup P_{2,S} \cup \dots \cup P_{k+1,S} \quad (3)$$

from axioms and the $k - 1$ -th stage clauses

$$P_{1,S \cup \{j\}} \cup P_{2,S \cup \{j\}} \cup \dots \cup P_{k,S \cup \{j\}}, \quad (4)$$

for all $j \in [n] \setminus S$. The derivation is as follows. For each $j \in [n] \setminus S$, we consecutively resolve the clause (4) with the clauses $\{\bar{p}_{1,j}, \bar{p}_{k+1,j}\}, \{\bar{p}_{2,j}, \bar{p}_{k+1,j}\}, \dots, \{\bar{p}_{k,j}, \bar{p}_{k+1,j}\}$ to obtain

$$P_{1,S} \cup P_{2,S} \cup \dots \cup P_{k+1,S} \cup \bar{p}_{k+1,j}. \quad (5)$$

We then consecutively resolve these with the axiom $\{p_{k+1,1}, p_{k+1,2}, \dots, p_{k+1,n}\}$ to get the desired clause (3).

We shall estimate the size and the width of the above resolution proof. The k -th stage consists of 2^{n-k} clauses, and in deriving each of these we have used $k^2 + k$ resolution steps. Thus the size of the proof is

$$\sum_{k=1}^n (k^2 + k) 2^{n-k} = O(n^2 2^n) = 2^{n+O(\log n)}.$$

The maximum width at the k -th stage is achieved when resolving the first clause of (5) with the axiom, and it is $(k+1)(n-k) + n - 2$. It is easy to see that the maximum over k is achieved near to $k = \frac{n-1}{2}$, and it is $\frac{1}{4}n^2 + O(n)$. \square

4 The upper bound

In this section we construct the following resolution proof of the Weak Pigeon-Hole Principle.

Theorem 4.1 *There is a resolution proof of $WPHP_n^m$ of size $2^{O(\frac{n \log n}{\log m} + \log m)}$ for every m and n .*

Proof We construct the proof recursively as follows. We divide the pigeons into $\frac{m}{i}$ blocks of i pigeons in each block. We divide the holes into two groups of size j and $n - j$, respectively (the figure 1). We first take each block of i pigeons against the first group of j holes, i.e. we consider PHP_j^i (the grey area on the figure 1), provided that $i > j$. As there are additional $n - j$ holes, we can derive a clause saying that at least one pigeon from the block goes to the second group of holes. We now consider each block as a single pigeon against the second group of holes (shown by arrows on the figure 1), i.e. we have $PHP_{n-j}^{\frac{m}{i}}$, assuming $\frac{m}{i} > n - j$. Strictly speaking, this is not an instance of the Pigeon-Hole Principle. The propositional variables p_{ij} are now replaced by big disjunctions, and therefore the negations \bar{p}_{ij} are now big conjunctions. This is a potential problem, because a resolution proof of the Pigeon-Hole principle can contain clauses having many negative literals. If we try to transform directly such a proof into a proof of the new version of PHP , where we have blocks instead of single pigeons, the expansion of these clauses would cause exponential blow-up. However, the difficulty can be overcome by considering only proofs that contains small (constant) number of negative literals in each clause. Indeed, the proof from proposition 3.2, we shall use as a basis case in our construction, is of such kind. By expanding the clause and simulating each resolution step, we multiply the size of the proof by only a polynomial (in this particular case quadratic) factor in the block-size. Therefore, denoting by $S(m, n)$ the size of the resolution proof of PHP_n^m , obtained in this way we get

$$S(m, n) \leq \frac{m}{i} S(i, j) + i^2 S\left(\frac{m}{i}, n - j\right).$$




Figure 1: The recursion

We now let $i = \sqrt{m}$ and $i = \frac{n}{2}$ and get

$$S(m, n) \leq \sqrt{m}S\left(\sqrt{m}, \frac{n}{2}\right) + mS\left(\sqrt{m}, \frac{n}{2}\right) \leq m^2S\left(\sqrt{m}, \frac{n}{2}\right).$$

After k iterations of the above, we get

$$S(m, n) \leq m^{2(1+\frac{1}{2}+\dots+\frac{1}{2^{k-1}})}S\left(m^{\frac{1}{2^k}}, \frac{n}{2^k}\right) \leq m^4S\left(m^{\frac{1}{2^k}}, \frac{n}{2^k}\right).$$

Let us now substitute $\log \frac{\log m}{\log n}$ for k :

$$S(m, n) \leq m^4S\left(n, \frac{n \log n}{\log m}\right).$$

By Proposition 3.2, $S\left(n, \frac{n \log n}{\log m}\right)$, which is the basis case of the recursion, can be bounded by $2^{O\left(\frac{n \log n}{\log m}\right)}$. Therefore we have

$$S(m, n) \leq 2^{O\left(\frac{n \log n}{\log m} + \log m\right)},$$

as claimed. \square

5 The lower bound

We shall prove the following

Theorem 5.1 *For any positive constant ε , every resolution proof of $WPHP_n^m$ of width $(\frac{1}{16} - \varepsilon)n^2$ has to be of size $2^{\Omega(n)}$*

Proof We first describe Adversary’s strategy. He divides the set of holes into two equally big sets H_1 and H_2 of size $\frac{n}{2}$ each (this does not need to be at random), and then assign each pigeon to either H_1 or H_2 , independently at random with probability $\frac{1}{2}$. At this stage, Adversary does not assign a particular hole to any pigeon, so we say that all the pigeons are *floating*. During the game, a pigeon becomes *fixed* iff either an “yes” answer or at least $\frac{n}{16}$ “no” answer about it are present in the current position. “Fixed” means that a hole is assigned to the pigeon. We can now explain Adversary replies to Prover’s moves. There are several cases to be considered

1. Prover “forgets” some information. This is the easiest case. Adversary may need only to change the status of some pigeons from “fixed” to “floating”.
2. Prover makes a query about a pigeon P . Assuming w.l.o.g. that P has been assigned to H_1 , Adversary answers as follows:
 - (a) P is fixed: consistently with the assigned hole.
 - (b) The query is about a hole in H_2 : “no”.
 - (c) P is floating and the query is about a hole in H_1 : if the number of “no” answers, involving P and holes in H_1 , in the current position, is less than $\frac{n}{4} - 1$, Adversary answers “no”. Otherwise, he assigns the first empty hole in H_1 to P , and then answers consistently with the assignment. P becomes fixed.

It is easy to see that Adversary can play until there are $\frac{n}{4}$ fixed pigeons in either part. Indeed, the only danger is when a pigeon P changes the status from “floating” to “fixed”. Exactly $\frac{n}{4}$ holes are explicitly forbidden by “no” answers at this stage, while less than $\frac{n}{4}$ holes are implicitly forbidden, because they are assigned to the fixed pigeons. Thus there is an empty hole to be assigned to P .

Let us now consider the situation when Adversary gives up, i.e. there are exactly $\frac{n}{4}$ in one of the parts, either H_1 or H_2 . There have to be at least $4\epsilon n$ “yes” fixed pigeons, as otherwise there would be more than $(\frac{1}{4} - 4\epsilon)n$ busy “no” pigeons and therefore the size of the position would be bigger than $(\frac{1}{16} - \epsilon)n^2$. As the parts H_1 and H_2 have been assigned to the pigeons independently at random with probability $\frac{1}{2}$, it follows that the probability that the $4\epsilon n$ “yes” busy pigeons are consistent with the initial assignment, is $\frac{1}{2^{4\epsilon n}}$, and therefore the description of Prover’s winning strategy has to be at least $2^{4\epsilon n}$. \square

6 Conclusion

We have proven a resolution width-size trade-offs for the Weak Pigeon-Hole Principle. There are, however, several open questions left.

The bound on the width, $(\frac{1}{16} - \epsilon)n^2$, is too strong. At least for the ordinary Pigeon-Hole Principle, one should be able to prove a stronger result. We conjecture the following sharp-threshold.

Conjecture 6.1 *For any positive constant ϵ , every resolution proof of PHP_n^{n+1} of width $(\frac{1}{4} - \epsilon)n^2$ has to be of size $2^{\Omega(n \log n)}$*

It is in agreement with the upper bounds we have proven, the propositions 3.1 and 3.2. In a preliminary version of this work, we claimed the above conjecture as a proven result. Unfortunately, it seems that one cannot get this result with an easy adaptation of the current lower-bound techniques.

Another interesting open problem is to tighten the gap between the known upper and lower bounds on the resolution proofs of the weak pigeon hole principle, $WPHP_n^m$. The best known lower bound, $2^{\Omega(\frac{n}{(\log m)^2})}$, appears in [12]. We conjecture that our upper bound is optimal.

Conjecture 6.2 *Any optimal proof of the $WPHP_n^m$ is of size $2^{\Omega(\frac{n \log n}{\log m})}$ for $n < m \leq 2^{\theta(\sqrt{n \log n})}$.*

References

- [1] P. Beame and T. Pitassi. Simplified and improved resolution lower bounds. In *Proceedings of the 37th annual IEEE symposium on Foundation Of Computer Science*, pages 274–282, 1996.
- [2] E. Ben-Sasson and A. Wigderson. Short proofs are narrow - resolution made simple. *Journal of the ACM*, 48(2), March 2001. A preliminary version appears at the 31st STOC, 1999.
- [3] S. Buss and T. Pitassi. Resolution and the weak pigeonhole principle. In *Computer science logic (Aarhus, 1997)*, pages 149–156. Springer, 1998.
- [4] S.R. Buss and G. Turán. Resolution proofs of generalized pigeonhole principles. *Theoretical Computer Science*, 62:311–317, 1988.
- [5] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.
- [6] T. Pitassi and R. Raz. Regular resolution lower bounds for the weak pigeonhole principle. In *Proceedings of the 33rd annual ACM Symposium on Theory Of Computing*, pages 347–355, 2001.
- [7] P. Pudlák. Proofs as games. *American Mathematical Monthly*, pages 541–550, June-July 2000.
- [8] P. Pudlák and S.R. Buss. How to lie without being (easily) convicted and the lengths of proofs in propositional calculus. In *Computer Science Logic'94*, volume 993 of *Lecture Notes in Computer Science*, pages 151–162, 1995.
- [9] R. Raz. Resolution lower bounds for the weak pigeonhole principle. Technical Report 21, Electronic Colloquium on Computational Complexity, 2001. Available at <http://www.eccc.uni-trier.de/eccc/>.
- [10] A. Razborov. Improved resolution lower bounds for the weak pigeonhole principle. Technical Report 55, Electronic Colloquium on Computational Complexity, 2001. Available at <http://www.eccc.uni-trier.de/eccc/>.
- [11] A. Razborov. Resolution lower bounds for the weak functional pigeonhole principle. Technical Report 75, Electronic Colloquium on Computational Complexity, 2001. Available at <http://www.eccc.uni-trier.de/eccc/>.
- [12] A. Razborov. Resolution lower bounds for perfect matching principles. In *Proceedings of the 17th annual IEEE Conference on Computational Complexity*. IEEE, May 2002. Available at <http://www.mi.ras.ru/razborov/>.