Electronic Colloquium on Computational Complexity, Report No. 21 (2001)

Resolution Lower Bounds for the Weak Pigeonhole Principle

Ran Raz*
Weizmann Institute, and
The Institute for Advanced Study
ranraz@wisdom.weizmann.ac.il

Abstract

We prove that any Resolution proof for the weak pigeon hole principle, with n holes and any number of pigeons, is of length $\Omega(2^{n^\epsilon})$, (for some global constant $\epsilon > 0$).

1 Introduction

The *Pigeon Hole Principle* (PHP) is one of the most widely studied tautologies in propositional proof theory. The tautology PHP_n is a DNF encoding of the following statement: There is no one to one mapping from $n+1$ pigeons to n holes. The *Weak Pigeon Hole Principle* (WPHP) is a version of the pigeon hole principle that allows a larger number of pigeons. The tautology $WPHP_n^m$ (for $m \geq n+1$) is a DNF encoding of the following statement: There is no one to one mapping from m pigeons to n holes. For $m > n+1$, the weak pigeon hole principle is a weaker statement than the pigeon hole principle. Hence, it may have much shorter proofs in certain proof systems.

The weak pigeon hole principle is one of the most fundamental combinatorial principles. In particular, it is used in most probabilistic counting arguments and hence in many combinatorial proofs. Moreover, as observed by Razborov, there are certain connections between the weak pigeon hole principle and the problem of $P \neq NP$ [Razb]. Indeed, the weak pigeon hole principle (with a relatively large number of pigeons) can be interpreted as an encoding of the following statement: There are no small DNF formulas for SAT. Hence, (in most proof systems), a short proof for a certain formulation of the statement $NP \not\subset P/poly$ can be translated into a short proof for the weak pigeon hole principle. That is, a lower bound for the length of proofs for the weak pigeon hole principle implies a lower bound for the length of proofs for a certain formulation of the statement $NP \not\subset P/poly$.

*Research supported by US-Israel BSF grant 98-00349, and NSF grant CCR-9987077.

ISSN 1433-8092

Resolution is one of the most widely studied propositional proof systems. The *Resolution rule* says that if C and D are two clauses and x_i is a variable then any assignment that satisfies both of the clauses, $C \vee x_i$ and $D \vee \neg x_i$, also satisfies $C \vee D$. The clause $C \vee D$ is called the *resolvent* of the clauses $C \vee x_i$ and $D \vee \neg x_i$ on the variable x_i. A *Resolution refutation* for a CNF formula F is a sequence of clauses $C_1, C_2, \ldots, C_s$, such that: (1) Each clause C_j is either a clause of F or a resolvent of two previous clauses in the sequence. (2) The last clause, C_s, is the empty clause (and hence it has no satisfying assignments). We can represent a Resolution refutation as an acyclic directed graph on vertices $C_1, \ldots, C_s$, where each clause of F has out-degree 0, and any other clause has two edges pointing to the two clauses that were used to produce it. It is well known that Resolution is a sound and complete propositional proof system, that is, a formula F is unsatisfiable if and only if there exists a Resolution refutation for F. We think of a refutation for an unsatisfiable formula F also as a proof for the tautology $\neg F$. A well-known and widely studied restricted version of Resolution (that is still complete) is called *Regular Resolution*. In a Regular Resolution refutation, along any path in the directed acyclic graph, each variable is resolved upon at most once.

There are trivial Resolution proofs (and Regular Resolution proofs) of length $2^n \cdot poly(n)$ for the pigeon hole principle and for the weak pigeon hole principle. In a seminal paper, Haken proved that for the pigeon hole principle, the trivial proof is (almost) the best possible [Hak]. More specifically, Haken proved that any Resolution proof for the tautology PHP_n is of length $2^{\Omega(n)}$. Haken's argument was further developed in several other papers (e.g., [Urq, BeP, BSW]). In particular, it was shown that a similar argument gives lower bounds also for the weak pigeon hole principle, but only for small values of m. More specifically, super-polynomial lower bounds were proved for any Resolution proof for the tautology $WPHP_n^m$, for $m < c \cdot n^2/\log n$ (for some constant c) [BT]. The weak pigeon hole principle with larger values of m has attracted a lot of attention in recent years. However, the standard techniques for proving lower bounds for Resolution fail to give lower bounds for the weak pigeon hole principle. In particular, for $m \geq n^2$, no non-trivial lower bound was known.

For the weak pigeon hole principle with large values of m, there do exist Resolution proofs (and Regular Resolution proofs) which are much shorter than the trivial ones. In particular, it was proved by Buss and Pitassi that for $m > c^{\sqrt{n}\log n}$ (for some constant c), there are Resolution (and Regular Resolution) proofs of length $poly(m)$ for the tautology $WPHP_n^m$ [BuP]. Can this upper bound be further improved ? Can one prove a matching lower bound ? As mentioned above, for $m \geq n^2$, no non-trivial lower bound was known. A partial progress was made by Razborov, Wigderson and Yao, who proved exponential lower bounds for Regular Resolution proofs, but only when the Regular Resolution proof is of a certain restricted form [RWY]. An exponential lower bound for any Regular Resolution proof was proved in [PR]. In this paper, we prove an exponential lower bound for any Resolution proof.

More precisely, we prove that for any m, any Resolution proof for the weak pigeon hole principle, $WPHP_n^m$, is of length $\Omega(2^{n^\epsilon})$, (where $\epsilon > 0$ is some global constant).

2 Preliminaries

2.1 Resolution as a Search Problem

A *literal* is either an atom (i.e., a variable x_i) or the negation of an atom (i.e., $\neg x_i$). A *clause* is a disjunction of literals. If C and D are two clauses and x_i is a variable then any assignment that satisfies both of the clauses, $C \vee x_i$ and $D \vee \neg x_i$, obviously satisfies the clause $C \vee D$ as well. As mentioned in the introduction, the clause $C \vee D$ is called the resolvent of the clauses $C \vee x_i$ and $D \vee \neg x_i$ on the variable x_i. A Resolution refutation for a CNF formula F is a sequence of clauses $C_1, C_2, \ldots, C_s$, such that, each clause C_j is either a clause of F or a resolvent of two previous clauses in the sequence, and such that, the last clause, C_s, is the empty clause. We think of the empty clause as a clause that has no satisfying assignments. We think of a Resolution refutation for F also as a proof for $\neg F$. Without loss of generality, we assume that no clause in a Resolution proof contains both x_i and $\neg x_i$ (such a clause is always satisfied and hence it can be removed from the proof). The *length*, or *size*, of a Resolution proof is the number of clauses in it.

As mentioned in the introduction, we represent a Resolution proof as an acyclic directed graph G on the vertices $C_1, \ldots, C_s$. In this graph, each clause C_j which is an original clause of F has out-degree 0, and any other clause has two edges pointing to the two clauses that were used to produce it. We call the vertices of out-degree 0 (i.e., the clauses that are original clauses of F) the *leaves* of the graph. Without loss of generality, we can assume that the only clause with in-degree 0 is the last clause C_s (as we can just remove any other clause with in-degree 0). We call the vertex C_s the *root* of the graph, and we denote it also by *Root*.

We label each vertex C_j in the graph by the variable x_i that was used to derive it (i.e., the variable x_i that was resolved upon), unless the clause C_j is an original clause of F (and then C_j is not labelled). If a clause C_j is labelled by a variable x_i we label the two edges going out from C_j by 0 and 1, where the edge pointing to the clause that contains x_i is labelled by 0, and the edge pointing to the clause that contains $\neg x_i$ is labelled by 1. That is, if the clause $C \vee D$ was derived from the two clauses $C \vee x_i$ and $D \vee \neg x_i$ then the vertex $C \vee D$ is labelled by x_i, the edge from the vertex $C \vee D$ to the vertex $C \vee x_i$ is labelled by 0 and the edge from the vertex $C \vee D$ to the vertex $D \vee \neg x_i$ is labelled by 1. For a non-leaf node u of the graph G, define,

$$\mathbf{Label}(\mathbf{u}) = \text{the variable labelling } u.$$

We think of $Label(u)$ as a variable queried at the node u.

Let p be a path on G, starting from the root. Note that along a path p, a variable x_i may appear (as a label of a node u) more than once. We say that the path p evaluates x_i to 0 if $x_i = Label(u)$ for some node u on the path p, and after the last appearance of x_i as $Label(u)$ (of a node u on the path) the path p continues on the edge labelled by 0 (i.e., if u is the last node on p such that $x_i = Label(u)$ then p contains the edge labelled by 0 that goes out from u). In the same way, we say that the path p evaluates x_i to 1 if $x_i = Label(u)$ for some node u on the path p, and after the last appearance of x_i as $Label(u)$ (of a node u on the path) the path p continues on the edge labelled by 1 (i.e., if u is the last node on p such that $x_i = Label(u)$ then p contains the edge labelled by 1 that goes out from u).

For any node u of the graph G, we define $Zeros(u)$ to be the set of variables that the node u "remembers" to be 0, and $Ones(u)$ to be the set of variables that the node u "remembers" to be 1, that is,

> **Zeros**(**u**) = the set of variables that are evaluated to 0 by every path p from the root to u.
>
> **Ones**(**u**) = the set of variables that are evaluated to 1 by every path p from the root to u.

Note that for any u, the two sets $Zeros(u)$ and $Ones(u)$ are disjoint.

The following proposition gives the connection between the sets $Zeros(u)$, $Ones(u)$ and the literals appearing in the clause u. The proposition is particularly interesting when u is a leaf of the graph.

Proposition 2.1 *Let F be an unsatisfiable CNF formula and let G be (the graph representation of) a Resolution refutation for F. Then, for any node u of G and for any x_i, if the literal x_i appears in the clause u then $x_i \in Zeros(u)$, and if the literal $\neg x_i$ appears in the clause u then $x_i \in Ones(u)$.*

Proof:
Assume that the literal x_i appears in the clause u. (The claim for the literal $\neg x_i$ is proved in the same way). Let p be a path from the root C_s to u. We will show that the path p evaluates x_i to 0. If no node $v < u$ on the path p satisfies $Label(v) = x_i$ then the literal x_i appears in the clause C_s, in contradiction to the fact that C_s is the empty clause. Hence, there exists a node $v < u$ on the path p, such that, $Label(v) = x_i$. Let v be the last (i.e., the largest) such node. Let w be the next node on p (i.e., the successor of v on the path p). Thus, the edge (v, w) is contained in the path p. Since v is the last node on p such that $Label(v) = x_i$, no node z on the path p from w to u satisfies $Label(z) = x_i$. Hence, since the literal x_i appears in u, it appears in w as well. Thus, (v, w) is the edge labelled by 0. That is, p evaluates x_i to 0. □

2.2 The Weak Pigeonhole Principle

The propositional weak pigeon hole principle, $WPHP^m_n$, states that there is no one-to-one mapping from m pigeons to n holes. The underlying Boolean variables, $x_{i,j}$, for $1 \leq i \leq m$ and $1 \leq j \leq n$, represent whether or not pigeon i is mapped to hole j. The negation of the pigeonhole principle, $\neg WPHP^m_n$, is expressed in conjunctive normal form (CNF) as the conjunction of m *pigeon clauses* and $\binom{m}{2} \cdot n$ *hole clauses*. For every $1 \leq i \leq m$, we have a pigeon clauses, $(x_{i,1} \vee \ldots \vee x_{i,n})$, stating that pigeon i maps to some hole. For every $1 \leq i_1 < i_2 \leq m$ and every $1 \leq j \leq n$, we have a hole clauses, $(\neg x_{i_1,j} \vee \neg x_{i_2,j})$, stating that pigeons i_1 and i_2

do not both map to hole j. We refer to the pigeon clauses and the hole clauses also as pigeon axioms and hole axioms.

Let G be (the graph representation of) a Resolution refutation for $\neg WPHP_n^m$. Then, by Proposition 2.1, for any leaf u of the graph G, one of the following is satisfied:

1. u is a pigeon axiom, and then for some $1 \leq i \leq m$, the variables $x_{i,1}, \ldots, x_{i,n}$ are all contained in $Zeros(u)$.
2. u is a hole axiom, and then for some $1 \leq j \leq n$, there exist two different variables $x_{i_1,j}, x_{i_2,j}$ in $Ones(u)$.

2.3 Basic Notations

We denote by n the number of holes, and by m the number of pigeons. We denote by $Holes$ the set of holes, and by $Pigeons$ the set of pigeons. That is,

Holes $= \{1, ..., n\}$.

Pigeons $= \{1, ..., m\}$.

We will usually denote a hole by j, and a pigeon by i. By $x_{i,j}$ we denote the variable corresponding to pigeon i and hole j. We denote by $Variables$ the set of all these variables, and by $Variables_i$ the set of variables corresponding to the i^{th} pigeon. That is,

Variables $= \{x_{i,j} | i \in Pigeons, j \in Holes\}$.

$\mathbf{Variables_i} = \{x_{i,j} | j \in Holes\}$.

We will consider (the graph representation of) Resolution proofs for the weak pigeon hole principle. We denote such a graph by G. By u we will usually denote a node in the graph. We say that $u' < u$ if there is a path in the graph from u' to u. By p we will usually denote a path in the graph, starting from the root. Note that for any non-leaf node u of the graph, $Label(u)$ is a variable $x_{i,j}$. The sets $Zeros(u)$ and $Ones(u)$ are subsets of $Variables$.

We denote by ϵ a small fixed constant (say $\epsilon = 1/100$). For simplicity, we assume that n is large enough (say $n^\epsilon \geq 1000$). For simplicity, we assume that expressions like $n^\epsilon, n^{1-\epsilon}, n^{1-8\epsilon}/2$, etc', are all integers.

3 The Lower Bound

In this section, we prove our lower bound on the size of Resolution proofs for the weak pigeon hole principle. Fix $\epsilon = 1/100$, and assume for simplicity that $n^\epsilon \geq 1000$. (We do not attempt here to optimize the value of ϵ).

Theorem 3.1 *For any $m \geq n+1$, any Resolution proof for the tautology $WPHP_n^m$ is of length larger than $2^{n^\epsilon/100}$.*

In the rest of the section, we give the proof of Theorem 3.1. Let G be the graph representation of a Resolution proof for $WPHP_n^m$, and assume for a contradiction that the size of G (i.e., the number of vertices in G) is at most $2^{n^\epsilon/100}$. Note that since the size of G is at most $2^{n^\epsilon/100}$, we can assume w.l.o.g. that $m < 2^{n^\epsilon/100}$.

3.1 Adding Axioms

First, define for any integer $0 \leq k \leq n^\epsilon$,

$\mathbf{n_k} = k \cdot n^{1-\epsilon}$,

$\mathbf{m_k} = 2^{n^\epsilon - k}$.

(Recall that we assume that $n^\epsilon, n^{1-\epsilon}$ are integers). For any node u of the graph G and for any pigeon i, we define $Zeros_i(u)$ to be the set of variables $x_{i,j}$ that the node u "remembers" to be 0, that is, the set of variables $x_{i,j}$ that are evaluated to 0 by every path from the root to u. In other words,

$\mathbf{Zeros_i(u)} = Zeros(u) \cap Variables_i$.

We use the size of $Zeros_i(u)$ as a measure for the progress made on pigeon i along paths from the root. We compare this measure with the "mileposts" $\{n_k\}$. We define $Over^k(u)$ to be the set of pigeons that passed the k^{th} milepost (for the node u), and we say that a node u is an axiom of order k if for the node u at least m_k pigeons passed the k^{th} milepost. That is, for any integer $0 \leq k \leq n^\epsilon$ and any node u of G, we define,

$\mathbf{Over^k(u)} =$ the set of pigeons i such that $|Zeros_i(u)| \geq n_k$.

For $1 \leq k \leq n^\epsilon$, we say that a node u is a **pigeon-axiom of order k** if $|Over^k(u)| \geq m_k$.

Note that a pigeon-axiom of order $k = n^\epsilon$ is just a standard pigeon-axiom of the weak pigeon hole principle. We say that a node u is a hole-axiom if there exists a hole j and two different pigeons i_1, i_2, such that $x_{i_1,j}, x_{i_2,j} \in Ones(u)$. Note that this is just a standard hole-axiom of the pigeon hole principle.

In our lower bound proof, we allow the leaves of the graph G to be pigeon-axioms of any order k (and not only pigeon-axioms of order $k = n^\epsilon$ as in the usual weak pigeon hole principle). That is, we say that G is a Resolution proof for the weak pigeon hole principle if all its leaves are axioms (i.e., all the leaves of G are either hole-axioms or pigeon-axioms of some order). We assume w.l.o.g. that in G, a non-leaf node u is never an axiom (otherwise, we can just disconnect the two edges going out from u and hence convert u to a leaf). In particular, no non-leaf node is a pigeon-axiom of order k, for any k.

One consequence of the assumption that no non-leaf node is a pigeon-axiom of order k is that if a node u is a pigeon-axiom of order k then $|Over^k(u)| = m_k$. This is true because if for some u we had $|Over^k(u)| > m_k$ then any node v, such that there is an edge from v to u, would satisfy $|Over^k(v)| \geq m_k$, and hence the non-leaf node v would be a pigeon-axiom of order k. Therefore, we can assume that for any node u in the graph,

$$|Over^k(u)| \leq m_k.$$

For our lower bound proof, we need to refine the scale $\{n_k\}$. For any two integers $0 \leq k < n^\epsilon$ and $0 \leq l \leq n^\epsilon$ and any node u, define,

$\mathbf{n_{k,l}} = k \cdot n^{1-\epsilon} + l \cdot n^{1-2\epsilon}.$

$\mathbf{Over^{k,l}(u)}$ = the set of pigeons i such that $|Zeros_i(u)| \geq n_{k,l}$.

Note that $n_{k,0} = n_k$, and $n_{k,n^\epsilon} = n_{k+1}$.

3.2 The Random Assignment

We will define a probabilistic assignment $A_{i,j}$ to the variables $x_{i,j}$. Unlike in previous lower bound proofs, one should not interpret the assignment $A_{i,j}$ as a "random restriction" of the Resolution proof. The assignment $A_{i,j}$ will be used in a different way. The assignment $A_{i,j}$ is chosen at random according to some specific probability distribution, defined below. First, define,

$\left\{\mathbf{Holes^k}\right\}_{\mathbf{k=1}}^{\mathbf{n^\epsilon}}$ = a random partition of $Holes$ into n^ϵ sets of size $n^{1-\epsilon}$ each.

That is, we partition $Holes$ into n^ϵ sets of size $n^{1-\epsilon}$ each. The intuition is that the set of holes $Holes^k$ will be used "against" pigeon-axioms of order k. For each $1 \leq k \leq n^\epsilon$, define,

$\left\{\mathbf{Holes^{k,l}}\right\}_{\mathbf{l=1}}^{\mathbf{n^\epsilon}}$ = a random partition of $Holes^k$ into n^ϵ sets of size $n^{1-2\epsilon}$ each.

That is, we further partition each set $Holes^k$ into n^ϵ sets of size $n^{1-2\epsilon}$ each. Altogether, the set $Holes$ was partitioned into $n^{2\epsilon}$ sets of size $n^{1-2\epsilon}$ each. We denote by $Variables_i^{k,l}$ the set of variables corresponding to the i^{th} pigeon and holes in $Holes^{k,l}$. That is,

$\mathbf{Variables_i^{k,l}} = \{x_{i,j} | j \in Holes^{k,l}\}.$

Next, we would like to define for every $1 \leq k \leq n^\epsilon$, a set of pigeons $Pigeons^k$. For $m_k \leq n^\epsilon$, we would like the set $Pigeons^k$ to contain all pigeons. For larger values of m_k, we would like each pigeon to be chosen (independently, at random) with a certain probability. For every $1 \leq k \leq n^\epsilon$, define,

$\mathbf{p_k} = \min\left[1, \frac{n^\epsilon}{m_k}\right]$.

$\mathbf{Pigeons^k}$ = a random subset of $Pigeons$, such that each pigeon is chosen (independently, at random) with probability p_k.

For every pigeon i, and every $1 \leq k \leq n^\epsilon$ and $1 \leq l \leq n^\epsilon$, define the subset $AOnes_i^{k,l}$ of the set $Variables_i^{k,l}$, in the following way.

$$\mathbf{AOnes_i^{k,l}} = \begin{cases} \text{a random subset of size } n^{1-6\epsilon} \text{ of } Variables_i^{k,l} & \text{if } i \in Pigeons^k \\ \emptyset & \text{if } i \notin Pigeons^k \end{cases}$$

The set $AOnes$ is now defined to be the union of all the sets $AOnes_i^{k,l}$, and the set $AZeros$ is defined to be the complement of $AOnes$. The assignment $A_{i,j}$ is defined by, $A_{i,j} = 1$ iff $x_{i,j} \in AOnes$. That is,

$$\mathbf{AOnes} = \bigcup_{i,k,l} AOnes_i^{k,l}.$$

$$\mathbf{AZeros} = Variables \setminus AOnes.$$

$$\mathbf{A_{i,j}} = \begin{cases} 1 & \text{if } x_{i,j} \in AOnes \\ 0 & \text{if } x_{i,j} \in AZeros \end{cases}$$

3.3 Properties of the Assignment

For our lower bound proof, we do not need the assignment $A_{i,j}$, and the sets that were involved in defining it, to be probabilistic. We just need them to satisfy certain properties. These properties are satisfied (with high probability) by the probabilistic construction that we defined, but we will only need one assignment (and sets) that satisfy the properties. The properties that we will need are summarized in the following claim.

Claim 3.1 *With exponentially high probability, all the following are satisfied, for every pigeon i, every hole j, every nodes u, v, and every $1 \leq k \leq n^\epsilon$ and $1 \leq l \leq n^\epsilon$.*

1. *If $j \in Holes^k$ and $i \notin Pigeons^k$ then*
$$A_{i,j} = 0.$$

2. *If $i \in Pigeons^k$ and $|Zeros_i(u)| - |Zeros_i(v)| \geq n^{1-2\epsilon}$ then*
$$\left|[Zeros_i(u) \setminus Zeros_i(v)] \cap AOnes_i^{k,l}\right| > n^{1-8\epsilon}/2.$$

3. *If i_1 and i_2 are two different pigeons then*
$$\left|\left\{j' \in Holes^{k,l} \mid A_{i_1,j'} = 1 \text{ and } A_{i_2,j'} = 1\right\}\right| < 2n^{1-10\epsilon}.$$

4. If $|Ones(u)| \geq n^{\epsilon}$ *then*

$$Ones(u) \cap AZeros \neq \emptyset.$$

5. If u *is a pigeon-axiom of order* k *then*

$$Pigeons^{k} \cap Over^{k}(u) \neq \emptyset.$$

6. For any u*,*

$$\left|Pigeons^{k} \cap Over^{k-1}(u)\right| < 10n^{\epsilon}.$$

Proof:
Recall that the number of pigeons and the number of nodes are both bounded by $2^{n^{\epsilon}/100}$. The number of holes is n. Recall that we assume that $\epsilon = 1/100$, and $n^{\epsilon} \geq 1000$. For the proof of the claim, we just have to verify that (for specific objects, $i, j, k, l, u, v, i_1, i_2$), the requirement in each one of the properties is falsified with exponentially small probability (say, with probability smaller than $2^{-n^{\epsilon}/25}$). This will usually follow by the standard Chernoff-Hoeffding bounds or by other simple probabilistic arguments. Let us analyze the properties one by one.

Property 1:
By the definition of $AOnes_i^{k,l}$, the requirement in this property is always satisfied.

Property 2:
$Zeros_i(u) \setminus Zeros_i(v)$ is a fixed subset of $Variables_i$ of size $\geq n^{1-2\epsilon}$. Assume w.l.o.g. that the size of $Zeros_i(u) \setminus Zeros_i(v)$ is exactly $n^{1-2\epsilon}$. Since $i \in Pigeons^{k}$, the set $AOnes_i^{k,l}$ is a random subset of $Variables_i$ of size exactly $n^{1-6\epsilon}$. Hence, the intersection

$$[Zeros_i(u) \setminus Zeros_i(v)] \cap AOnes_i^{k,l}$$

is of expected size $n^{1-8\epsilon}$, and by the standard Chernoff-Hoeffding bounds the actual size of the intersection is very close to $n^{1-8\epsilon}$, with high probability. In particular, the probability that the size of the intersection is $\leq n^{1-8\epsilon}/2$ is exponentially small (and in particular, smaller than $2^{-n^{\epsilon}/25}$).

Property 3:
Denote,

$$H_{i_1} = \left\{ j' \in Holes^{k,l} \mid A_{i_1,j'} = 1 \right\},$$

and,

$$H_{i_2} = \left\{ j' \in Holes^{k,l} \mid A_{i_2,j'} = 1 \right\}.$$

Then,

$$\left\{ j' \in Holes^{k,l} \mid A_{i_1,j'} = 1 \text{ and } A_{i_2,j'} = 1 \right\} = H_{i_1} \cap H_{i_2}.$$

If either $i_1 \notin Pigeons^{k}$ or $i_2 \notin Pigeons^{k}$ then $H_{i_1} \cap H_{i_2}$ is empty. If both $i_1, i_2 \in Pigeons^{k}$ then H_{i_1} and H_{i_2} are both random subsets of $Holes^{k,l}$ of size $n^{1-6\epsilon}$ each. Recall that $Holes^{k,l}$ is a set of size $n^{1-2\epsilon}$. Hence, the intersection $H_{i_1} \cap H_{i_2}$ is of expected size $n^{1-10\epsilon}$, and by the standard Chernoff-Hoeffding bounds the actual size of the intersection is very close to

$n^{1-10\epsilon}$, with high probability. In particular, the probability that the size of the intersection is $\geq 2n^{1-10\epsilon}$ is exponentially small (and in particular, smaller than $2^{-n^{\epsilon}/25}$).

Property 4:
Denote $s = |Ones(u)|$, and assume w.l.o.g. that s is exactly n^{ϵ}. Let $x_{i_1,j_1}, x_{i_2,j_2}, ..., x_{i_s,j_s}$ be the s variables in $Ones(u)$. It is easy to verify that for any $1 \leq t \leq s$, the probability for $A_{i_t,j_t} = 1$ is smaler than $1/2$, even under the condition that $A_{i_1,j_1}, ..., A_{i_{t-1},j_{t-1}}$ are all 1. Hence, the probability that $A_{i_1,j_1}, ..., A_{i_s,j_s}$ are all 1 is smaller than $2^{-n^{\epsilon}}$.

Property 5:
$Over^k(u)$ is a set of m_k pigeons. If $m_k \leq n^{\epsilon}$ then each one of these pigeons is in $Pigeons^k$ with probability 1. Otherwise, the probability for each one of these pigeons to be in $Pigeons^k$ is n^{ϵ}/m_k, and hence the probability that none of them is in $Pigeons^k$ is

$$\left(1 - \frac{n^{\epsilon}}{m_k}\right)^{m_k} < 2^{-n^{\epsilon}}.$$

Property 6:
As we have seen, $Over^{k-1}(u)$ is a set of at most $m_{k-1} = 2m_k$ pigeons. Assume w.l.o.g. that $Over^{k-1}(u)$ is a set of exactly $2m_k$ pigeons. If $m_k \leq n^{\epsilon}$ then $2m_k \leq 2n^{\epsilon}$ and the requirement is obviously satisfied. Otherwise, each one of these $2m_k$ pigeons is in $Pigeons^k$ with probability n^{ϵ}/m_k. Hence, the intersection $Pigeons^k \cap Over^{k-1}(u)$ is of expected size $2n^{\epsilon}$, and by the standard Chernoff-Hoeffding bounds the actual size of the intersection is very close to $2n^{\epsilon}$, with high probability. In particular, the probability that the size of the intersection is $\geq 10n^{\epsilon}$ is exponentially small (and in particular, smaller than $2^{-n^{\epsilon}/25}$).

□

3.4 The Adversary Strategy

In this subsection, we give the proof of Theorem 3.1, given one lemma (the main lemma).

With high probability, all the properties in Claim 3.1 are satisfied. Hence, we can fix the assignment $A_{i,j}$ (and all the sets involved in defining it, such as, $Pigeons^k$, $Holes^{k,l}$, etc') to some fixed values that satisfy all these properties. Thus, from now on, we assume that the assignment $A_{i,j}$ (and all the sets involved in defining it) are fixed (and are not probabilistic any more), and that all the properties in Claim 3.1 are satisfied.

For every non-leaf node u of the graph G, we define a value $Answer(u) \in \{0, 1\}$. We think of $Answer(u)$ as an adversary "answer" for the "query" $Label(u)$. The answer $Answer(u)$ depends on the assignment $A_{i,j}$ and the sets $Holes^{k,l}$.

> Assume that $Label(u) = x_{i,j}$, and $j \in Holes^{k,l}$. We define $Answer(u)$ in the following way:

1)	If $i \notin Over^{k-1,l-1}(u)$	**Answer(u)** =	**0**
2)	If $\exists i' \neq i$ s.t. $x_{i',j} \in Ones(u)$	**Answer(u)** =	**0**
3)	Otherwise,	**Answer(u)** =	$\mathbf{A_{i,j}}$

That is, the answer is automatically 0 if $i \notin Over^{k-1,l-1}(u)$, or if there exists $i' \neq i$ such that $x_{i',j} \in Ones(u)$. Otherwise, the answer is the value of $A_{i,j}$. Given the values $Answer(u)$ (for every non-leaf node u), we define a path (called $Path$) on the graph G. The path starts from the root of G and in each step it follows the edge labelled by $Answer(u)$, where u is the current node. We denote by $Leaf$ the leaf reached by the path $Path$. That is,

Path = the path that starts from *Root*, and that satisfies that for every (non-leaf) node u on the path, the path contains the edge that goes out from u and is labelled by $Answer(u)$.

Leaf = the leaf reached by $Path$.

Lemma 3.1 (Main Lemma) *For any* $1 \leq k \leq n^{\epsilon}$*, and any node* u *on the path* $Path$*,*

$$Pigeons^k \cap Over^k(u) = \emptyset.$$

Lemma 3.1 is proved in the next subsection. Let us show how the proof of Theorem 3.1 follows from Lemma 3.1.

Proof of Theorem 3.1:
By Lemma 3.1 and by Property 5 of Claim 3.1, no node u on $Path$ is a pigeon-axiom (of any order k). By the definition of $Answer(u)$, if there exists $i' \neq i$ such that $x_{i',j} \in Ones(u)$ then $Answer(u) = 0$. Hence, for no node u on $Path$ we will have that both $x_{i,j}$ and $x_{i',j}$ are in $Ones(u)$. That is, no node u on $Path$ is a hole-axiom. In particular, $Leaf$ is neither a pigeon-axiom (of any order k) nor a hole-axiom, in contradiction to the fact that all leaves of the graph G must be axioms. □

3.5 Pigeon-Sections

In this subsection, we give the proof of Lemma 3.1, given one claim (the main claim). For any node u on $Path$, define,

$\mathbf{u^+}$ = the successor of u on $Path$.

$\mathbf{u^-}$ = the predecessor of u on $Path$.

(u^+ is undefined for $u = Leaf$, and u^- is undefined for $u = Root$). For two nodes $v \leq w$ on $Path$, denote by $[v, w]$ the section of nodes (on $Path$) between them. That is,

$[\mathbf{v}, \mathbf{w}]$ = the set of nodes u on $Path$, such that, $v \leq u \leq w$.

For a pigeon $i \in Pigeons^k$, we will be interested in maximal sections on $Path$, such that, for every node u in the section, $i \in Over^{k-1}(u)$. For $1 \leq k \leq n^\epsilon$, we define a *pigeon-section of type* k, and the set $PigSec^k$ (of all these pigeon-sections), in the following way.

$(\mathbf{i}, [\mathbf{v}, \mathbf{w}])$ is a **pigeon-section of type k** if all the following are satisfied:

1. $i \in Pigeons^k$, and $v \leq w$ are nodes on $Path$.
2. For any node $u \in [v, w]$, we have $i \in Over^{k-1}(u)$.
3. The section $[v, w]$ is maximal with this property. That is, if $v \neq Root$ then $i \notin Over^{k-1}(v^-)$ and if $w \neq Leaf$ then $i \notin Over^{k-1}(w^+)$.

$\mathbf{PigSec^k}$ = the set of all pigeon-sections of type k.

We will further refine the categorization of pigeon-sections into types. We say that a pigeon-section of type k is of type (k, l) if the section $[v, w]$ contains a node u such that $i \in Over^{k-1,l-1}(u)$, and we define the set $PigSec^{k,l}$ to be the set of all these pigeon-sections. That is, for $1 \leq k \leq n^\epsilon$ and $1 \leq l \leq n^\epsilon + 1$,

$(\mathbf{i}, [\mathbf{v}, \mathbf{w}])$ is a **pigeon-section of type** $(\mathbf{k}, \mathbf{l})$ if all the following are satisfied:

1. $(i, [v, w])$ is a pigeon-section of type k.
2. For some node $u \in [v, w]$, we have $i \in Over^{k-1,l-1}(u)$.

$\mathbf{PigSec^{k,l}}$ = the set of all pigeon-sections of type (k, l).

Note the asymmetric role of k and l in the definition of pigeon-section of type (k, l). Note also that $PigSec^{k,1} = PigSec^k$.

Claim 3.2 (Main Claim) *For every* $1 \leq k \leq n^\epsilon$ *and* $1 \leq l \leq n^\epsilon$,

$$\left|PigSec^{k,l+1}\right| \leq \frac{1}{2} \cdot \left|PigSec^{k,l}\right|.$$

Claim 3.2 is proved in the next subsections. Let us show how the proof of Lemma 3.1 follows from Claim 3.2.

Proof of Lemma 3.1:
Since the number of pigeons and the number of nodes in the graph are both bounded by $2^{n^\epsilon/100}$, the number of pigeon-sections of type k is bounded by $2^{n^\epsilon/50}$. That is,

$$\left|PigSec^{k,1}\right| = \left|PigSec^k\right| \leq 2^{n^\epsilon/50}.$$

Hence, by n^ϵ applications of Claim 3.2,

$$\left|PigSec^{k,n^\epsilon+1}\right| \le 2^{-n^\epsilon} \cdot \left|PigSec^{k,1}\right| \le 2^{-n^\epsilon} \cdot 2^{n^\epsilon/50} < 1,$$

and since $\left|PigSec^{k,n^\epsilon+1}\right|$ is integer,

$$\left|PigSec^{k,n^\epsilon+1}\right| = 0.$$

That is, there are no pigeon-sections of type $(k, n^\epsilon + 1)$.

Assume for a contradiction to the statement of the lemma that for some node u on $Path$,

$$Pigeons^k \cap Over^k(u) \neq \emptyset.$$

Then, since

$$Over^k(u) = Over^{k-1,n^\epsilon}(u),$$

there exists $i \in Pigeons^k$, such that,

$$i \in Over^{k-1,n^\epsilon}(u).$$

Denote by $[v, w]$ the largest section (on $Path$) that contains u, and such that for every $u' \in [v, w]$ we have $i \in Over^{k-1}(u')$ (such a section exists because $i \in Over^{k-1}(u)$). Then, $(i, [v, w])$ is a pigeon-section of type $(k, n^\epsilon + 1)$, in contradiction to the fact that there are no such pigeon-sections. □

3.6 Forcing

Let u be a node such that $Label(u) = x_{i,j}$, and such that $i \in Pigeons^k$ and $j \in Holes^{k,l}$ (for some $1 \le k \le n^\epsilon$ and $1 \le l \le n^\epsilon$). Recall that $Answer(u)$ is 0 if there exists $i' \neq i$ such that $x_{i',j} \in Ones(u)$. If, in addition, $i \in Over^{k-1,l-1}(u)$ and $A_{i,j} = 1$ we say that $x_{i,j}$ is forced to 0 at the node u by $x_{i',j}$. (Recall that if $i \notin Over^{k-1,l-1}(u)$ or $A_{i,j} = 0$ then $Answer(u)$ would be 0 anyways, so we do not consider it as "forcing"). That is,

> Assume that $Label(u) = x_{i,j}$, and $j \in Holes^{k,l}$. We say that $\mathbf{x_{i,j}}$ **is forced to 0 at the node u by** $\mathbf{x_{i',j}}$ if all the following are satisfied:
>
> 1. $i \in Pigeons^k$ and $A_{i,j} = 1$.
> 2. $i \in Over^{k-1,l-1}(u)$.
> 3. $x_{i',j} \in Ones(u)$.

Assume that $x_{i,j}$ is forced to 0 by $x_{i',j}$ at a node u on $Path$. Then, since $i \in Pigeons^k$ and $i \in Over^{k-1,l-1}(u)$, there exists a (unique) pigeon-section $(i, [v, w])$ of type (k, l) such that $u \in [v, w]$. (To see this, just denote by $[v, w]$ the largest section on $Path$ that contains u, and such that for every $\hat{u} \in [v, w]$ we have $i \in Over^{k-1}(\hat{u})$, such a section exists because $i \in Over^{k-1}(u)$. Then, $(i, [v, w])$ is a pigeon-section of type (k, l)).

Consider the nodes on $Path$ from the root to u, that is, the nodes in $[Root, u]$. Denote by u' the last node in $[Root, u]$, such that, $Label(u') = x_{i',j}$. Since $x_{i',j} \in Ones(u)$, we know that $Answer(u')$ is 1. Therefore, by the definition of $Answer(u')$, we know that $i' \in Over^{k-1,l-1}(u')$, and by Property 1 of Claim 3.1 we know that $i' \in Pigeons^k$ (otherwise, $A_{i',j}$ would be 0, and hence $Answer(u')$ would be 0 as well). By the same argument as before, there exists a (unique) pigeon-section $(i', [v', w'])$ of type (k, l) such that $u' \in [v', w']$. We categorize the "forcing" to types according to the relations between the nodes u', v, w, w', as follows.

Let u be a node on $Path$. Assume that $Label(u) = x_{i,j}$, and $j \in Holes^{k,l}$. Assume that $x_{i,j}$ is forced to 0 by $x_{i',j}$ at the node u. Let u' be the last node in $[Root, u]$, such that $Label(u') = x_{i',j}$. Let $(i, [v, w])$ be the pigeon-section of type (k, l) such that $u \in [v, w]$, and let $(i', [v', w'])$ be the pigeon-section of type (k, l) such that $u' \in [v', w']$.

1. We say that the forcing is **a forcing of type 1** if $u' < v$.
2. We say that the forcing is **a forcing of type 2** if $u' \in [v, w]$ and $w' \geq w$.
3. We say that the forcing is **a forcing of type 3** if $u' \in [v, w]$ and $w' < w$.

Note that since $u' < u$ and $u \in [v, w]$, any forcing is a forcing of one of these three types. For every k, l, we would like to count the number of variables forced to 0 at pigeon-sections of type (k, l). In our counting, we would like to count a variable more than once if it is forced to 0 at more than one pigeon-section. However, we would like to count a variable only once for each pigeon-section, that is, if the variable is forced to 0 many times at the same pigeon-section we count it only once. For every, $1 \leq k \leq n^\epsilon$ and $1 \leq l \leq n^\epsilon$, define,

$\mathbf{Forced^{k,l}}$ = the set of all pairs $(x_{i,j}, [v, w])$, such that all the following are satisfied:

1. $(i, [v, w])$ is a pigeon-section of type (k, l).
2. $j \in Holes^{k,l}$.
3. $x_{i,j}$ is forced to 0 at some node $u \in [v, w]$.

$\mathbf{Forced_1^{k,l}}$ = the set of all pairs $(x_{i,j}, [v, w])$, such that all the following are satisfied:

1. $(i, [v, w])$ is a pigeon-section of type (k, l).
2. $j \in Holes^{k,l}$.
3. $x_{i,j}$ is forced to 0 at some node $u \in [v, w]$, and the forcing is type 1.

$\mathbf{Forced_2^{k,l}}$ = the set of all pairs $(x_{i,j}, [v, w])$, such that all the following are satisfied:

1. $(i, [v, w])$ is a pigeon-section of type (k, l).

2. $j \in Holes^{k,l}$.
3. $x_{i,j}$ is forced to 0 at some node $u \in [v, w]$, and the forcing is type 2.

$\mathbf{Forced_3^{k,l}}$ = the set of all pairs $(x_{i,j}, [v, w])$, such that all the following are satisfied:

1. $(i, [v, w])$ is a pigeon-section of type (k, l).
2. $j \in Holes^{k,l}$.
3. $x_{i,j}$ is forced to 0 at some node $u \in [v, w]$, and the forcing is type 3.

3.7 Bounding the Number of Forced Variables

In this subsection, we give the proof of Claim 3.2. The proof will follow easily by the following four claims.

Claim 3.3 *For every* $1 \leq k \leq n^\epsilon$ *and* $1 \leq l \leq n^\epsilon$,

$$\left|Forced_1^{k,l}\right| \leq \left|PigSec^{k,l}\right| \cdot n^\epsilon.$$

Claim 3.4 *For every* $1 \leq k \leq n^\epsilon$ *and* $1 \leq l \leq n^\epsilon$,

$$\left|Forced_2^{k,l}\right| \leq \left|PigSec^{k,l}\right| \cdot 20n^{1-9\epsilon}.$$

Claim 3.5 *For every* $1 \leq k \leq n^\epsilon$ *and* $1 \leq l \leq n^\epsilon$,

$$\left|Forced_3^{k,l}\right| \leq \left|PigSec^{k,l}\right| \cdot 20n^{1-9\epsilon}.$$

Claim 3.6 *For every* $1 \leq k \leq n^\epsilon$ *and* $1 \leq l \leq n^\epsilon$,

$$\left|Forced^{k,l}\right| \geq \left|PigSec^{k,l+1}\right| \cdot n^{1-8\epsilon}/2.$$

Proof of Claim 3.2:
Since any forcing is a forcing of type 1 or type 2 or type 3,

$$\left|Forced^{k,l}\right| \leq \left|Forced_1^{k,l}\right| + \left|Forced_2^{k,l}\right| + \left|Forced_3^{k,l}\right|.$$

Hence, the proof follows immediately from Claims 3.3, 3.4, 3.5, 3.6, using the assumptions that $\epsilon = 1/100$ and $n^\epsilon \geq 1000$. □

Proof of Claim 3.3:
Let $(i, [v, w])$ be a pigeon-section of type (k, l). Denote,

$$F^1_{(i,[v,w])} = \left\{(x_{i,j}, [v, w]) \in Forced_1^{k,l}\right\}.$$

We will show that for every such $(i, [v, w])$,

$$\left|F^1_{(i,[v,w])}\right| \leq n^\epsilon,$$

(and hence the claim follows).

Fix $(i,[v,w])$ to be a pigeon-section of type (k,l). For every $(x_{i,j},[v,w]) \in F^1_{(i,[v,w])}$, we know that $x_{i,j}$ is forced to 0 at some node $u \in [v,w]$ by some $x_{i',j}$, and the forcing is type 1. Hence, the last node $u' \in [Root, u]$, such that $Label(u') = x_{i',j}$, satisfies $u' < v$. That is, $x_{i',j}$ does not appear as $Label(\hat{u})$ for any $\hat{u} \in [v,u]$, and since $x_{i',j} \in Ones(u)$ we conclude that $x_{i',j} \in Ones(v)$. Thus, for every $(x_{i,j},[v,w]) \in F^1_{(i,[v,w])}$, there is (at least one) corresponding $x_{i',j} \in Ones(v)$. Hence,

$$\left|F^1_{(i,[v,w])}\right| \leq |Ones(v)|.$$

To finish the proof of the claim, it is enough to show that for every node v on $Path$,

$$|Ones(v)| \leq n^{\epsilon}.$$

Let v be a node such that $|Ones(v)| > n^{\epsilon}$. We will show that v is not on $Path$. By Property 4 of Claim 3.1,

$$Ones(v) \cap AZeros \neq \emptyset.$$

Hence, there exists $x_{\tilde{i},\tilde{j}} \in Ones(v)$, such that $A_{\tilde{i},\tilde{j}} = 0$. Hence, for any node $\tilde{u}$ such that $Label(\tilde{u}) = x_{\tilde{i},\tilde{j}}$, we have $Answer(\tilde{u}) = 0$. Since $Path$ always follows the edge $Answer(\tilde{u})$ (when $\tilde{u}$ is the current node), it will never evaluate $x_{\tilde{i},\tilde{j}}$ to 1. Since every path to v evaluates $x_{\tilde{i},\tilde{j}}$ to 1, we conclude that v is not on $Path$. □

Proof of Claim 3.4:
Let $(i,[v,w])$ be a pigeon-section of type (k,l). Denote,

$$F^2_{(i,[v,w])} = \left\{(x_{i,j},[v,w]) \in Forced^{k,l}_2\right\}.$$

We will show that for every such $(i,[v,w])$,

$$\left|F^2_{(i,[v,w])}\right| \leq 20n^{1-9\epsilon},$$

(and hence the claim follows).

Fix $(i,[v,w])$ to be a pigeon-section of type (k,l). We will count the number of possibilities for $(x_{i,j},[v,w]) \in F^2_{(i,[v,w])}$. For every $(x_{i,j},[v,w]) \in F^2_{(i,[v,w])}$, we know that $x_{i,j}$ is forced to 0 at some node $u \in [v,w]$ by some $x_{i',j}$, and the forcing is type 2. Therefore, there exists a pigeon-section $(i',[v',w'])$ of type (k,l) such that $v' < u$ and $w' \geq w$. Thus, $w \in [v',w']$. Hence, since $(i',[v',w'])$ is a pigeon-section of type k, we know that $i' \in Pigeons^k$ and $i' \in Over^{k-1}(w)$. Thus, for every $(x_{i,j},[v,w]) \in F^2_{(i,[v,w])}$, each corresponding $x_{i',j}$ satisfies that i' is in

$$Pigeons^k \cap Over^{k-1}(w).$$

By Property 6 of Claim 3.1,

$$\left|Pigeons^k \cap Over^{k-1}(w)\right| < 10n^{\epsilon},$$

and hence for the pigeon-section $(i,[v,w])$, the number of possibilities for i' is bounded by $10n^{\epsilon}$.

Since $x_{i,j}$ is forced to 0 at u by $x_{i',j}$, we know that $A_{i,j}=1$ (by the definition of forcing), and $A_{i',j}=1$ (since $x_{i',j}\in Ones(u)$ and u is on $Path$, and as in the proof of Claim 3.3 $Path$ cannot evaluate $x_{i',j}$ to 1 if $A_{i',j}=0$). Hence, j is in

$$\left\{j \in Holes^{k,l} \mid A_{i,j}=1 \text{ and } A_{i',j}=1\right\}.$$

By Property 3 of Claim 3.1,

$$\left|\left\{j \in Holes^{k,l} \mid A_{i,j}=1 \text{ and } A_{i',j}=1\right\}\right| < 2n^{1-10\epsilon},$$

and hence for every i', the number of possibilities for j is bounded by $2n^{1-10\epsilon}$.

Altogether, for the pigeon-section $(i,[v,w])$, the number of possibilities for i' is bounded by $10n^{\epsilon}$, and for every i' the number of possibilities for j is bounded by $2n^{1-10\epsilon}$. Hence,

$$\left|F^2_{(i,[v,w])}\right| \le 10n^{\epsilon}\cdot 2n^{1-10\epsilon} = 20n^{1-9\epsilon}.$$

□

Proof of Claim 3.5:
For every $(x_{i,j},[v,w])\in Forced_3^{k,l}$, we know that $x_{i,j}$ is forced to 0 at some node $u\in[v,w]$ by some $x_{i',j}$, and the forcing is type 3. Let u' be the last node in $[Root,u]$, such that $Label(u')=x_{i',j}$, and let $(i',[v',w'])$ be the pigeon-section of type (k,l) such that $u'\in[v',w']$. We will say, in this case, that the pigeon-section $(i',[v',w'])$ is **responsible** for $(x_{i,j},[v,w])\in Forced_3^{k,l}$. Thus, for every $(x_{i,j},[v,w])\in Forced_3^{k,l}$, there is (at least one) pigeon-section $(i',[v',w'])$ of type (k,l) responsible for it.

Let $(i',[v',w'])$ be a pigeon-section of type (k,l). Denote by $F^3_{(i',[v',w'])}$ the set of all $(x_{i,j},[v,w])\in Forced_3^{k,l}$ that $(i',[v',w'])$ is responsible for. We will show that for every such $(i',[v',w'])$,

$$\left|F^3_{(i',[v',w'])}\right| \le 20n^{1-9\epsilon},$$

and hence, obviously,

$$\left|Forced_3^{k,l}\right| \le \left|PigSec^{k,l}\right|\cdot 20n^{1-9\epsilon}.$$

The bound for $|F^3_{(i',[v',w'])}|$ is proved in a similar way to the proof of the bound for $|F^2_{(i,[v,w])}|$, in Claim 3.4.

Fix $(i',[v',w'])$ to be a pigeon-section of type (k,l). We will count the number of possibilities for $(x_{i,j},[v,w])\in F^3_{(i',[v',w'])}$. For every $(x_{i,j},[v,w])\in F^3_{(i',[v',w'])}$, we know that $x_{i,j}$ is forced to 0 at some node $u\in[v,w]$ by $x_{i',j}$, and the forcing is type 3. We also know that if u' is the last node in $[Root,u]$ such that $Label(u')=x_{i',j}$ then $u'\in[v',w']$ (by the definition of $F^3_{(i',[v',w'])}$). Since the forcing is type 3, we know that $u'\in[v,w]$ and $w'<w$. Thus,

$w' \in [v, w]$. Hence, since $(i, [v, w])$ is a pigeon-section of type k, we know that $i \in Pigeons^k$ and $i \in Over^{k-1}(w')$. Thus, for every $(x_{i,j}, [v, w]) \in F^3_{(i',[v',w'])}$, we know that i is in

$$Pigeons^k \cap Over^{k-1}(w').$$

By Property 6 of Claim 3.1,

$$\left| Pigeons^k \cap Over^{k-1}(w') \right| < 10n^\epsilon,$$

and hence for the pigeon-section $(i', [v', w'])$, the number of possibilities for i is bounded by $10n^\epsilon$.

Since $x_{i,j}$ is forced to 0 at u by $x_{i',j}$, we know that $A_{i,j} = 1$ (by the definition of forcing), and $A_{i',j} = 1$ (as in the proof of Claim 3.4). Hence, j is in

$$\left\{ j \in Holes^{k,l} \mid A_{i,j} = 1 \text{ and } A_{i',j} = 1 \right\}.$$

By Property 3 of Claim 3.1,

$$\left| \left\{ j \in Holes^{k,l} \mid A_{i,j} = 1 \text{ and } A_{i',j} = 1 \right\} \right| < 2n^{1-10\epsilon},$$

and hence for every i, the number of possibilities for j is bounded by $2n^{1-10\epsilon}$.

Altogether, for the pigeon-section $(i', [v', w'])$, the number of possibilities for i is bounded by $10n^\epsilon$, and for every i the number of possibilities for j is bounded by $2n^{1-10\epsilon}$. For every i, the number of possibilities for $[v, w]$ is (at most) one, because there is (at most) one pigeon-section $(i, [v, w])$ of type (k, l) such that $w' \in [v, w]$. Hence,

$$\left| F^3_{(i',[v',w'])} \right| \leq 10n^\epsilon \cdot 2n^{1-10\epsilon} = 20n^{1-9\epsilon}.$$

□

Proof of Claim 3.6:
Let $(i, [v, w])$ be a pigeon-section of type $(k, l + 1)$. Then, obviously, $(i, [v, w])$ is a pigeon-section of type (k, l) as well. Denote,

$$F_{(i,[v,w])} = \left\{ (x_{i,j}, [v, w]) \in Forced^{k,l} \right\}.$$

We will show that for every such $(i, [v, w])$,

$$\left| F_{(i,[v,w])} \right| \geq n^{1-8\epsilon}/2,$$

(and hence the claim follows).

Fix $(i, [v, w])$ to be a pigeon-section of type $(k, l + 1)$. Then, for some node $u \in [v, w]$, we have

$$|Zeros_i(u)| \geq n_{k-1,l}.$$

For simplicity of the notations, assume that v is not the root, and hence v^- exists. Let s be the last node in $[v^-, u]$, such that, $i \notin Over^{k-1,l-1}(s)$ (such an s exists because $i \notin Over^{k-1,l-1}(v^-)$). Denote $t = s^+$. Then,

$$|Zeros_i(t)| = n_{k-1,l-1}.$$

(This is true because by the definition of s we know that $i \in Over^{k-1,l-1}(t)$, and if we had $|Zeros_i(t)| > n_{k-1,l-1}$ then we would have had $|Zeros_i(s)| \geq n_{k-1,l-1}$, in contradiction to the definition of s). Thus,

$$|Zeros_i(u)| - |Zeros_i(t)| \geq n^{1-2\epsilon},$$

and hence, by Property 2 of Claim 3.1,

$$\left|[Zeros_i(u) \setminus Zeros_i(t)] \cap AOnes_i^{k,l}\right| > n^{1-8\epsilon}/2.$$

To finish the proof of the claim, it is enough to show that

$$x_{i,j} \in [Zeros_i(u) \setminus Zeros_i(t)] \cap AOnes_i^{k,l} \quad \Longrightarrow \quad (x_{i,j}, [v,w]) \in F_{(i,[v,w])}.$$

Let $x_{i,j} \in [Zeros_i(u) \setminus Zeros_i(t)] \cap AOnes_i^{k,l}$. First note that since $x_{i,j} \in AOnes_i^{k,l}$, we know that $i \in Pigeons^k$ and $j \in Holes^{k,l}$. Since $x_{i,j} \in [Zeros_i(u) \setminus Zeros_i(t)]$, there is a node $z \in [t,u]$, such that, $Label(z) = x_{i,j}$ and $Answer(z) = 0$. Since $x_{i,j} \in AOnes_i^{k,l}$, we know that $A_{i,j} = 1$, and since $z \in [t,u]$, we know that $i \in Over^{k-1,l-1}(z)$. Hence, $Answer(z)$ is 0 only if $x_{i,j}$ is forced to 0 at the node z. Thus, $x_{i,j}$ is forced to 0 at the node z, and since $(i,[v,w])$ is a pigeon-section of type (k,l), we conclude that $(x_{i,j}, [v,w]) \in F_{(i,[v,w])}$. □

Acknowledgment

I would like to thank Toni Pitassi for many helpful conversations and for years of collaboration that led to this work. I would also like to thank the Small World Café, where most of this research was done. Sometimes, the small things make the difference.

References

[BeP] Beame, P., and Pitassi, T., "Simplified and improved resolution lower bounds," *Foundations of Computer Science*, 1996, pp. 274-282.

[BuP] Buss, S., and Pitassi, T., "Resolution and the weak pigeonhole principle," *Springer-Verlag Lecture Notes in Computer Science*, Publications of selected papers presented at *Proceedings from Computer Science Logic 1997.*

[BSW] Ben-Sasson, E., and Wigderson, A., "Short proofs are narrow–resolution made simple," *Symposium on Theory of Computing*, 1999, pp. 517-526.

[BT] Buss, S., and Turan, G., "Resolution proofs of generalized pigeonhole principles," *Theoretical Computer Science*, 62, 1988, pp. 311-317.

[Hak] Haken, A. "The intractability of resolution," *Theoretical Computer Science*, 39, 1985, pp. 297-308.

[PR] Pitassi, T., and Raz, R., "Regular resolution lower bounds for the weak pigeonhole principle," to appear in *Symposium on Theory of Computing*, 2001.

[Razb] Razborov, A., "Lower bounds for the polynomial calculus," *Computational Complexity*, 7, 1998, pp. 291-324.

[RWY] Razborov, A., Wigderson, A., and Yao, A., "Read-once branching programs, rectangular proofs of the pigeonhole principle, and the transversal calculus," *Symposium on Theory of Computing*, 1997, pp. 739-748.

[Urq] Urquhart, A., "Hard examples for resolution," *Journal of ACM*, 34, 1987, pp. 209-219.

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