A New Proof of the Weak Pigeonhole Principle

Alexis Maciel*
Department of Mathematics
and Computer Science
Clarkson University
Potsdam, NY 13699-5815
U.S.A.
alexis@clarkson.edu

Toniann Pitassi†
Department of Computer
Science
University of Arizona
Tucson, AZ 85721
U.S.A.
toni@cs.arizona.edu

Alan R. Woods
Department of Mathematics
and Statistics
University of Western
Australia
Nedlands, WA 6907
Australia
woods@maths.uwa.edu.au

ABSTRACT
The exact complexity of the weak pigeonhole principle is an old and fundamental problem in proof complexity. Using a diagonalization argument, Paris, Wilkie and Woods [9] showed how to prove the weak pigeonhole principle with bounded-depth, quasipolynomial-size proofs. Their argument was further refined by Krajíček [5]. In this paper, we present a new proof: we show that the the weak pigeonhole principle has quasipolynomial-size proofs where every formula consists of a single AND/OR of polylog fan-in. Our proof is conceptually simpler than previous arguments, and is optimal with respect to depth.

1. INTRODUCTION
The pigeonhole principle is a fundamental axiom of mathematics, stating that there is no one-to-one mapping from m pigeons to n holes, $m > n$. It expresses a very basic fact about cardinalities of sets and is used ubiquitously in almost all areas of mathematics. As examples, the induction principle is simply a special case of the pigeonhole principle, and many combinatorial counting arguments reduce to the pigeonhole principle.

Perhaps not surprisingly, then, the inherent difficulty of proving the pigeonhole principle is tightly connected to important questions in proof theory and circuit complexity. It has served as the classic hard example for proof complexity, and versions of it have been used to obtain some of the strongest lower bounds and separations known to date. Examples include Resolution, bounded-depth Frege systems, Cutting Planes, and relativized bounded arithmetic.

There are several important open problems connected to the complexity of the weaker forms of the pigeonhole principle, which we define as those for which $n \le m/2$. First, the weak pigeonhole principle is connected to how much of elementary number theory, including the existence of infinitely many primes, can be proven in $I\Delta_0$, a weak system of arithmetic. Paris, Wilkie and Woods [9] show that a considerable part of elementary number theory, including the existence of infinitely many primes, is provable in $I\Delta_0$ with the weak pigeonhole principle for Δ_0-definable functions added as an axiom scheme. It is a longstanding open question whether or not one can dispense of the weak pigeonhole principle, by proving it within $I\Delta_0$.

Secondly, the complexity of the weak pigeonhole principle is related to the complexity of approximate counting. The problem of recognizing the approximate size of a set is in the polynomial-time hierarchy. However, all known proofs of this fact rely on the weak pigeonhole principle. These results translate downwards: there are bounded-depth, polynomial-size circuits that can approximately count the number of 1's in a 0/1 bit string. However, once again, all known proofs of correctness require much higher proof-theoretic complexity. This is a perplexing situation: is it possible to prove that small circuits exist for approximate counting, and also to prove that any correctness proof for these small circuits is inherently more complex than these circuits? A positive answer would follow if one could prove superpolynomial lower bounds on the size of bounded-depth Frege proofs of the weak pigeonhole principle for the case when $n = m/2$.

Lastly, the complexity of the weak pigeonhole principle is connected to the inherent complexity of *proving* circuit lower bounds. In the last decade, substantial effort has gone into understanding the metamathematics of the P versus NP question. In pioneering work, Razborov and Rudich [12] show that most circuit lower bounds are natural, and hence, under cryptographic certain assumptions, these methods cannot be extended to proving P $\neq$ NP. It would be a big breakthrough to extend this type of result to show that there can be no proof of $P \neq NP$ (formalized in a reasonable way) in bounded arithmetic. Razborov [11]

*Research supported by NSF grant CCR-9877150.

†Research supported by NSF Grant CCR-9457782, US-Israel BSF Grant 95-00238, and Grant INT-9600919/ME-103 from NSF and MŠMT (Czech Republic).

STOC 2000 Portland Oregon USA

has shown that this question is connected to the difficulty of proving the weak pigeonhole principle, since the circuit lower bound statement can encode the weak pigeonhole principle in a certain sense.

Resolving the above three questions amounts to understanding the exact complexity of proving the pigeonhole principle. This tautology is expressed propositionally by a formula of size polynomial in m, where the underlying variables are $P_{i,j}$, for $i \leq m$ and $j \leq n$. Three key complexity parameters are d, n and S: the first parameter, d, measures the depth of the Frege proof; the second parameter, n, is the number of holes; S is the size of the proof. Clearly, as d and S increase and n decreases, the pigeonhole principle becomes easier to prove. The ultimate goal is to obtain a precise and smooth characterization of the smallest value for S as we vary the other two parameters, d and n.

In [9], Paris, Wilkie and Woods use a very clever diagonalization argument to give constant-depth, quasipolynomial-size Frege proofs of the weak pigeonhole principle. This is surprising, especially since it has been shown that any constant-depth Frege proof of the weak pigeonhole principle requires exponential size whenever n is at least $m - c$, for c a constant [10; 7]. Their argument actually translates into depth-3.5, quasipolynomial-size proofs in the sequent calculus, and Krajíček [5] extends their argument to obtain depth-1.5, quasipolynomial-size proofs. (Depth $d + .5$, for any nonnegative integer d, means that each formula has depth at most $d + 1$, but the bottom level of gates are restricted to polylog fan-in.) Despite this breakthrough, there are still huge gaps in our overall understanding in terms of the three parameters mentioned above. In particular, are these results optimal in terms of depth? Is there a more constructive, constant-depth proof of the weak pigeonhole principle? Can the size be improved from quasipolynomial to polynomial?

The main result of this paper is a new proof of the weak pigeonhole principle. Our new proof is a step toward resolving the above-mentioned questions, and the exact complexity of the weak pigeonhole principle. We show that the weak pigeonhole principle has quasipolynomial-size proofs where every formula consists of a single AND/OR of polylog fan-in. In the above terminology, we obtain a depth-.5 proof. Translated to bounded arithmetic, it follows from our proof that the weak pigeonhole principle with respect to f can be proven in $T_2^2(f)$.

Our proof is optimal with respect to depth as exponential lower bounds are known for depth-0 sequent calculus proofs, i.e., Resolution proofs, of the weak pigeonhole principle [4]. Our upper bound is also tight in another sense: [6; 14] show that the proof cannot be made tree-like, unless the size becomes exponential. Moreover, our proof is conceptually simpler than the previous upper bound due to Paris, Wilkie and Woods: it is a simple divide and conquer, along the lines of the upper bounds for Resolution proofs of the weak pigeonhole principle [3], combined with an amplification phase which allows us to speed up the induction.

The outline for the remainder of the paper is as follows. In Section 2, we give precise definitions of the pigeonhole principle tautology and of the proof system that we will be working with. In Section 3, we give an overview and generalization of the Resolution upper bound of [3]. In Section 4, we present our main result. In Section 5, we optimize the argument given in Section 4. Finally, in Section 6, we put our new upper bound in perspective with the many previous results that are known in this area, and conclude with open problems.

2. DEFINITIONS

The propositional proof system that we will study in this paper is the sequent calculus, LK, modified to allow unbounded fan-in connectives. Formulas are built up using the connectives $\wedge$, $\vee$, and $\neg$. All connectives are assumed to have unbounded fan-in. The formula $\wedge(A_1, \ldots, A_n)$ denotes the logical AND of the multi-set consisting of $A_1, \ldots A_n$, and similarly for $\vee$. Thus commutativity and associativity of the connectives is implicit. Our proof system operates on sequents which are sets of formulas of the form $A_1, \ldots, A_i \rightarrow B_1, \ldots, B_j$. The intended meaning is that the conjunction of the A_i's implies the disjunction of the B_j's. A proof of a sequent S in LK is a sequence of sequents, $S_1, \ldots, S_q$, such that each sequent S_i is either an initial sequent, or follows from previous sequents by one of the rules of inference, and the final sequent, S_q, is S. The size of the proof is $\sum_{1 \leq i \leq q} \text{size}(S_i)$ and its depth is $\max_{1 \leq i \leq q}(\text{depth}(S_i))$.

The *initial sequents* are of the form: (1) $x \rightarrow x$ where x is a literal; (2) $\rightarrow \wedge()$; $\vee() \rightarrow$. The rules of inference are as follows. First we have simple structural rules such as weakening (formulas can always be added to the left or to the right), contraction (two copies of the same formula can be replaced by one), and permutation (formulas in a sequent can be reordered). The remaining rules are the cut rule, and logical rules which allow us to introduce each connective on both the left side and the right side. The cut rule allows the derivation of $\Gamma, \Gamma' \rightarrow \Delta, \Delta'$ from $\Gamma, A \rightarrow \Delta$, and $\Gamma' \rightarrow A, \Delta'$. The logical rules are as follows.

1. (Negation-left) From $\Gamma \rightarrow A, \Delta$, we can derive $\neg A, \Gamma \rightarrow \Delta$.
2. (Negation-right) From $A, \Gamma \rightarrow \Delta$, derive $\Gamma \rightarrow \neg A, \Delta$.
3. (And-left) From $A_1, \wedge(A_2, \ldots, A_n), \Gamma \rightarrow \Delta$ derive $\wedge(A_1, \ldots, A_n), \Gamma \rightarrow \Delta$.
4. (And-right) From $\Gamma \rightarrow A_1, \Delta$ and $\Gamma \rightarrow \wedge(A_2, \ldots, A_n), \Delta$ derive $\Gamma \rightarrow \wedge(A_1, \ldots, A_n), \Delta$
5. (Or-left) From $A_1, \Gamma \rightarrow \Delta$ and $\vee(A_2, \ldots, A_n), \Gamma \rightarrow \Delta$ derive $\vee(A_1, \ldots, A_n), \Gamma \rightarrow \Delta$
6. (Or-right) From $\Gamma \rightarrow A_1, \vee(A_2, \ldots, A_n), \Delta$ derive $\Gamma \rightarrow \vee(A_1, \ldots, A_n), \Delta$.

DEFINITION 1. *Let d be a nonnegative integer. A formula is of depth $d + .5$ if it is of depth d or of depth $d + 1$ but with the arity of the level 1 connectives restricted to polylogarithmic in the size of the formula.*

A sequent calculus proof is of depth $d+.5$ if all the formulas that appear in it are either of depth d or of depth $d+1$ but with the arity of the level 1 connectives restricted to polylogarithmic in the size of the final sequent.

There is a well-known translation between propositional proofs of certain types of statements, and proofs of the corresponding first order principle in systems of bounded arithmetic. In particular, it is well-known that uniform, quasipolynomial-size, bounded-depth proofs of $\Sigma_1^b(f)$ statements (such as the pigeonhole principle) can be translated into $S_2(f)$ proofs. Also, uniform, polynomial-size, bounded-depth proofs of the same type of statements can be translated into $I\Delta_0(f)$ proofs. The upper bounds that we will be presenting are all sufficiently uniform that they will also carry through in the uniform setting. In particular, our proofs can be straightforwardly translated to show that the weak pigeonhole principle with respect to f has a proof in $S_2^3(f)$, and then by conservativity of S_2^3 over T_2^2, the proof can also be carried out in $T_2^2(f)$. (Details of this translation will be given in the full version of the paper.)

The pigeonhole principle on m pigeons and n holes says that there is no one-to-one function from a set of size m to a set of size n. Formally, this can be stated as follows:

$$\mathrm{PHP}_n^m : \ldots, \bigvee_{y\in[n]} P_{xy}, \ldots \rightarrow \ldots, P_{x_1 y}P_{x_2 y}, \ldots$$

where, on the left, x ranges over $[m]$ and, on the right, $x_1 \neq x_2$ range over $[m]$ and y ranges over $[n]$. Note that PHP_n^m is actually more general than the informal statement above since it asserts the nonexistence of any injective, many-valued function from $[m]$ to $[n]$.

Clearly as n decreases, the principle becomes weaker and weaker. When $n = m-1$, it is usually referred to as just the pigeonhole principle, and when $n \leq m/2$ it is referred to as the *weak* pigeonhole principle. The *onto* pigeonhole principle is a weaker version stating that there is no one-to-one, onto, many-valued function from m pigeons to n holes.

3. THE RESOLUTION UPPER BOUND

As mentioned in the introduction, the new proof of the weak pigeonhole principle presented in this paper uses some of the same ideas as the Resolution upper bound of Buss and Pitassi [3]. More precisely, they show that PHP_n^m has polynomial-size Resolution proofs whenever $n \leq (\log m)^2/\log\log m$. In this section, we provide an overview and generalization of this result.

First note that when $n = O(\log m)$ there are trivially polynomial-size Resolution proofs, by ignoring all but $n+1$ pigeons, and performing a brute-force refutation on these pigeons and holes.

Now assume for sake of contradiction that there is a mapping from m to n (for appropriately chosen n). Divide the m pigeons up into blocks, each of size $\log m + 1$. The first case is that some block of pigeons maps in a one-to-one way into the first $\log m$ holes, and in this case we get a direct contradiction by brute force. The other case is where no block of pigeons all map to the first $\log m$ holes. But in this case, each block of pigeons can be viewed as a metapigeon, and now we have a one-to-one map from $m/(\log m+1)$ metapigeons to the last $n - \log m$ holes, and we can proceed inductively. This argument can be translated into a Resolution proof because each inductive instance of the pigeonhole principle is still a conjunction of a set of clauses.

We can use this idea more generally to prove PHP_n^m with a size-S Resolution refutation, where $n \leq \log m \log S/\log\log S$. Let the block size be b, where $b = \log S$. Dividing up the m pigeons into m/b blocks, each of size b, either some block maps one-to-one into the first b holes, or not. In the first case, we can use brute-force to get a size $O(S)$ refutation, and in the second case, we have $m/(\log S)$ metapigeons, and $n - \log S$ holes. Continuing for $k = n/(\log S)$ iterations, as long as $n \leq \log m \log S/(\log\log S)$, we reach the desired contradiction.

Thus, we obtain polynomial-size Resolution refutations of PHP_n^m for $n = O((\log m)^2/\log\log m)$, quasipolynomial-size Resolution refutations for $n = O(\log m)^c$, etc.

Our new upper bound gives small proofs of PHP_n^m for much larger n, but the depth increases slightly, from 0 to .5.

4. OUR NEW UPPER BOUND

Our goal is to show that PHP_n^{2n} has a quasipolynomial-size, tree-like proof of depth 1.5. We start by presenting the argument that we will then formalize as a sequent calculus proof.

The proof is in two parts: first we prove $\mathrm{PHP}_n^{n^2}$ and then we prove PHP_n^{2n}. Let us start with $\mathrm{PHP}_n^{n^2}$. By contradiction, suppose that there is an injective, many-valued function from $A = [n^2]$ to $B = [n]$. (For the remainder of this section, we will simply speak of functions even though we really mean many-valued functions.) Let $A_1, \ldots, A_n$ be the partition of A into sets of size n. Let B_1, B_2 be the partition of B into sets of size $n/2$. Then either

1. all the pigeons of some block A_i are sent to holes in the first block B_1, or
2. in every block there is at least one pigeon that is sent to a hole in the second block B_2.

If the first case occurs, then we have an injective function from a set of n pigeons to a set of $n/2$ holes. The function is injective because the original function is.

We now claim that the second case also gives an injective function from a set of n pigeons to a set of $n/2$ holes. View each block as a new superpigeon. Send each superpigeon to all the holes where its member pigeons are sent. We are guaranteed that each superpigeon is sent to at least one hole in the second block. The induced function from these n superpigeons to the $n/2$ holes in B_2 is injective again because of the injectivity of the original function.

This is the first step of the proof. In this step, the number of pigeons was reduced to n and the number of holes was reduced by half. In the second step, we will *amplify* the number of pigeons back up to n^2. Let f be the original function from $[n^2]$ to $[n]$ and let g be the new function from $[n]$ to $[n/2]$. Define a function h from $[n^2]$ to $[n/2]$ by setting $h(i) = k$ iff there is $j \in [n]$ such that $f(i) = j$ and $g(j) = k$. This new function h is injective because of the injectivity of both f and g.

We now repeat these two steps to obtain a sequence of injective functions from $[n]$ to $[n/4]$, from $[n^2]$ to $[n/4]$, from $[n]$ to $[n/8]$, from $[n^2]$ to $[n/8]$, ... , until an injective function from $[n]$ to $[1]$ is obtained. This is the desired contradiction, which proves $\mathrm{PHP}_n^{n^2}$.

Now we prove PHP_n^{2n}. Again by contradiction, suppose that there is an injective function f from $[2n]$ to $[n]$. We define a function g from $[4n]$ to $[2n]$ as follows. Partition $[4n]$ into two blocks A_1, A_2 of size $2n$ and partition $[2n]$ into two blocks B_1, B_2 of size n. The function g is defined by using f to map A_1 to B_1 and a *translated* version of f to map A_2 to B_2. Now compose g and f as was done above to obtain a function h from $[4n]$ to $[n]$. Both g and h are injective because of the injectivity of f. This process can be generalized and repeated to obtain a sequence of injective functions with increasingly larger domain. Eventually, we get an injective function from $[n^2]$ to $[n]$, which contradicts $\mathrm{PHP}_n^{n^2}$ and completes the proof of PHP_n^{2n}.

We now turn to the formalization of this argument as a quasipolynomial-size, tree-like sequent calculus proof of depth 1.5. The proof will consist of a sequence alternations between the two steps mentioned above. Since pigeons will eventually be not just simple pigeons but superpigeons, as a result of the reduction step, and since the function from pigeons to holes will eventually be the composition of earlier functions, as a result of the amplification step, we generalize the statement of the pigeonhole principle as follows. Let A and B be any two sets.

$$\mathrm{PHP}_B^A(Q) : \dots , \bigvee_{y \in B} Q_{xy}, \dots \to \dots , Q_{x_1 y} Q_{x_2 y}, \dots$$

where, on the left, x ranges over A and, on the right, $x_1 \neq x_2$ range over A and y ranges over B. Here, the Q_{xy} can be arbitrary formulas and not just propositional variables.

In fact, in our proof, each Q_{xy} will be a OR of small AND's, say $\bigvee_k Q_{xy}^{(k)}$. Since our goal is to obtain a proof of depth 1.5, we have to be able to state $\mathrm{PHP}_B^A(Q)$ in depth 1.5. To achieve this, we say that $\bigvee_{y \in B} Q_{xy}$ actually stands for $\bigvee_{y \in B} \bigvee_k Q_{xy}^{(k)}$, and that $Q_{x_1 y} Q_{x_2 y}$ stands for $\bigvee_{k_1, k_2} Q_{x_1 y}^{(k_1)} Q_{x_2 y}^{(k_2)}$.

The following two lemmas establish that the reduction and amplification steps in the above argument can be carried out by a quasipolynomial-size, tree-like sequent calculus proof of depth 1.5.

LEMMA 2. *Let A be any set of size n^2 and let B be any set of size $m \leq n$. Let $A_1, \dots , A_n$ be the partition of A into sets of size n and let B_1, B_2 be the partition of B into sets of size $m/2$. For every set of size-s, depth-1.5 formulas $(Q_{xy})_{x \in A, y \in B}$ of the form OR of small AND's, there is a set of size-(ns), depth-1.5 formulas $(R_{iy})_{i \in [n], y \in B_2}$ of the form OR of small AND's such that $\mathrm{PHP}_B^A(Q)$ has a size-$(ns)^{O(1)}$, tree-like, depth-1.5 sequent calculus proof from $\mathrm{PHP}_{B_1}^{A_1}(Q)$, ... , $\mathrm{PHP}_{B_1}^{A_n}(Q)$ and $\mathrm{PHP}_{B_2}^{n}(R)$.*

PROOF. For the moment, ignore the fact that the Q_{xy}'s are formulas and pretend that they are simple propositional variables. $\mathrm{PHP}_B^A(Q)$ can be written as follows:

$$\mathrm{PHP}_B^A(Q) : \dots , \bigvee_{y \in B} Q_{xy}, \dots \to \dots , Q_{x_1 y} Q_{x_2 y}, \dots$$

where, on the left, x ranges over A and, on the right, $x_1 \neq x_2$ range over A and y ranges over B. For any $i \in \{1, \dots , n\}$, $\mathrm{PHP}_{B_1}^{A_i}(Q)$ can be written as

$$\mathrm{PHP}_{B_1}^{A_i}(Q) : \dots , \bigvee_{y \in B_1} Q_{xy}, \dots \to \dots , Q_{x_1 y} Q_{x_2 y}, \dots$$

where, on the left, x ranges over A_i and, on the right, $x_1 \neq x_2$ range over A_i and y ranges over B_1.

The idea behind the set of R formulas is the following: R_{iy} will say that some pigeon from the ith block A_i is sent hole y. This is formalized as

$$R_{iy} = \bigvee_{x \in A_i} Q_{xy}.$$

$\mathrm{PHP}_{B_2}^{n}(R)$ can then be written as

$$\mathrm{PHP}_{B_2}^{n}(Q) : \dots , \bigvee_{y \in B_2} R_{iy}, \dots \to \dots , R_{iy} R_{jy}, \dots$$

where, on the left, i ranges over $[n]$ and, on the right, $i \neq j$ range over $[n]$ and y ranges over B_2. Of course, it is understood that $R_{iy} R_{jy}$ actually stands for

$$\bigvee_{x_1 \in A_i} \bigvee_{x_2 \in A_j} Q_{x_1 y} Q_{x_2 y}.$$

The proof of $\mathrm{PHP}_B^A(Q)$ from $\mathrm{PHP}_{B_1}^{A_1}(Q)$, ... , $\mathrm{PHP}_{B_1}^{A_n}(Q)$ and $\mathrm{PHP}_{B_2}^{n}(R)$ starts with the following sequents:

$$\bigvee_{y \in B} Q_{xy} \to \bigvee_{y \in B_1} Q_{xy}, \bigvee_{y \in B_2} Q_{xy} \quad (x \in A). \tag{1}$$

For every i, cut $\mathrm{PHP}_{B_1}^{A_i}(Q)$ with the corresponding sequents in (1). This gives

$$\begin{aligned} &\dots , \bigvee_{y \in B} Q_{xy}, \dots \\ &\qquad \to \dots , Q_{x_1 y} Q_{x_2 y}, \dots , \dots , \bigvee_{y \in B_2} Q_{xy}, \dots \\ &\qquad\qquad (i \in [n]) \end{aligned} \tag{2}$$

where, on the left, x ranges over A_i and, on the right, $x_1 \neq x_2$ range over A_i, y ranges over B_1 and x ranges over A_i.

Now consider the following sequents:

$$Q_{xy} \to R_{iy} \quad (i \in [n], x \in A_i, y \in B_2). \tag{3}$$

By using the OR-left rule and then the OR-right rule, combine the sequents in (3) that correspond to the various values of y:

$$\bigvee_{y \in B_2} Q_{xy} \to \bigvee_{y \in B_2} R_{iy} \quad (i \in [n], x \in A_i). \tag{4}$$

Cut each of the sequents in (2) with the corresponding sequents in (4) to obtain

$$\ldots, \bigvee_{y \in B} Q_{xy}, \ldots \to \ldots, Q_{x_1 y} Q_{x_2 y}, \ldots, \bigvee_{y \in B_2} R_{iy} \quad (i \in [n]) \tag{5}$$

where, on the left, x ranges over A_i and, on the right, $x_1 \neq x_2$ range over A_i and y ranges over B_1. Cut each of these sequents with $\mathrm{PHP}^n_{B_2}(R)$ to obtain

$$\ldots, \bigvee_{y \in B} Q_{xy}, \ldots \to \ldots, Q_{x_1 y} Q_{x_2 y}, \ldots, \ldots, R_{iy} R_{jy}, \ldots \tag{6}$$

On the left of this sequent, x ranges over A. On the right, we have one $Q_{x_1 y} Q_{x_2 y}$ for every $x_1 \neq x_2 \in A_i$, every $i \in [n]$ and every $y \in B_1$. On the right, we also have one $R_{iy} R_{jy}$ for every $i \neq j \in [n]$ and every $y \in B_2$.

Finally, recall that $R_{iy} R_{jy}$ actually stands for

$$\bigvee_{x_1 \in A_i} \bigvee_{x_2 \in A_j} Q_{x_1 y} Q_{x_2 y}.$$

Consider the following sequents:

$$R_{iy} R_{jy} \to \ldots, Q_{x_1 y} Q_{x_2 y}, \ldots \quad (i \neq j \in [n], y \in B_2) \tag{7}$$

where, on the right, x_1 ranges over A_i and x_2 ranges over A_j. Cut each of these sequents with (6) to obtain

$$\ldots, \bigvee_{y \in B} Q_{xy}, \ldots \to \ldots, Q_{x_1 y} Q_{x_2 y}, \ldots, \tag{8}$$

where, on the left, x still ranges over A, but, on the right, we now have one $Q_{x_1 y} Q_{x_2 y}$ for every $i \in [n]$, every $x_1 \neq x_2 \in A_i$ and every $y \in B_1$, and another $Q_{x_1 y} Q_{x_2 y}$ for every $i \neq j \in [n]$, every $x_1 \in A_i$, every $x_2 \in A_j$ and every $y \in B_2$. $\mathrm{PHP}^A_B(Q)$ can now be easily obtained by weakening, which completes the proof.

This was all done under the assumption that the Q_{xy}'s are simple propositional variables. Generalizing to OR's of small AND's is fairly easy since it requires only minor modifications of the proof. To illustrate, suppose that $Q_{xy} = \bigvee_k Q^{(k)}_{xy}$. Then the sequents in (7) become

$$\bigvee_{x_1 \in A_i} \bigvee_{k_1} \bigvee_{k_2} \bigvee_{x_2 \in A_j} Q^{(k_1)}_{x_1 y} Q^{(k_2)}_{x_2 y} \to \ldots, \bigvee_{k_1, k_2} Q^{(k_1)}_{x_1 y} Q^{(k_2)}_{x_2 y}, \ldots \quad (i \neq j \in [n], y \in B_2) \tag{9}$$

where, on the right, x_1 ranges over A_i and x_2 ranges over A_j. These sequents are proved in essentially the same way as the sequents in (7). We leave the remaining details to the reader as well as the straightforward task of verifying that the proof is tree-like and of size $(ns)^{O(1)}$. □

Lemma 3. *For every set C of size n, for every set D, for every set of size-s, depth-1.5 formulas $(Q_{xy})_{x \in C, y \in D}$ of the form OR of small AND's, and for every set of size-t, depth-1.5 formulas $(P_{wx})_{w \in [n^2], x \in [n]}$ of the form OR of small AND's, there is a set of size-$O(nst)$, depth-1.5 formulas $(R_{wy})_{w \in [n^2], y \in D}$ of the form OR of small AND's such that $\mathrm{PHP}^C_D(Q)$ weakened by the cedents of $\mathrm{PHP}^{n^2}_n(P)$ has a size-$(nst)^{O(1)}$, tree-like, depth-1.5 sequent calculus proof from $\mathrm{PHP}^{n^2}_D(R)$.*

Proof. Suppose that $Q_{xy} = \bigvee_k Q^{(k)}_{xy}$ and that $P_{wx} = \bigvee_j P^{(j)}_{wx}$. $\mathrm{PHP}^C_D(Q)$ weakened by the cedents of $\mathrm{PHP}^{n^2}_n(P)$ can be written as follows:

$$\ldots, \bigvee_{x \in [n]} P_{wx}, \ldots, \ldots, \bigvee_{y \in D} Q_{xy}, \ldots \to \ldots, P_{w_1 x} P_{w_2 x}, \ldots, \ldots, Q_{x_1 y} Q_{x_2 y}, \ldots \tag{10}$$

where, on the left, w ranges over $[n^2]$ and x ranges over C, and, on the right, x ranges over $[n]$, $x_1 \neq x_2$ range over C and y ranges over D. As mentioned earlier, it is understood that $\bigvee_{y \in D} Q_{xy}$ stands for

$$\bigvee_{y \in D} \bigvee_{k} Q^{(k)}_{xy},$$

that $Q_{x_1 y} Q_{x_2 y}$ stands for

$$\bigvee_{k_1, k_2} Q^{(k_1)}_{x_1 y} Q^{(k_2)}_{x_2 y},$$

and similarly for P.

We now want to define a set of R formulas that will allow us to prove the above sequent from $\mathrm{PHP}^{n^2}_D(R)$. The P formulas describe a function between a set of size n^2 and a set of size n, while the Q formulas describe a function between a set C of size n and a D set of size m. The idea is that the R formulas will describe the composition of those two functions. First, in what follows, we will identify C with $[n]$. More precisely, let f be any one-to-one, onto function from $[n]$ to C. Whenever x is in $[n]$ and we write Q_{xy}, we will actually mean $Q_{f(x)y}$. Now R_{wy} will be defined as follows:

$$R_{wy} = \bigvee_{x \in [n]} P_{wx} Q_{xy}.$$

Once again, this last formula actually stands for

$$\bigvee_{x \in [n]} \bigvee_{j} \bigvee_{k} P^{(j)}_{wx} Q^{(k)}_{xy}.$$

The sequent $\mathrm{PHP}^{n^2}_D(R)$ can be written as follows:

$$\mathrm{PHP}^{n^2}_D(R) : \ldots, \bigvee_{y \in D} R_{wy}, \ldots \to \ldots, R_{w_1 y} R_{w_2 y}, \ldots \tag{11}$$

where, on the left, w ranges over $[n^2]$ and, on the right, $w_1 \neq w_2$ range over $[n^2]$ and y ranges over D. In other words, this sequent says that if every w is sent to some y, then at least two w's will be sent to the same y.

The proof of (10) from this sequent consists of two main steps. First, we show that if two w's go to the same y, then either two w's go to the same x or two x's go to the same y. This can be written as

$$R_{w_1y}R_{w_2y} \rightarrow \ldots, P_{w_1x}P_{w_2x}, \ldots, \ldots, Q_{x_1y}Q_{x_2y}, \ldots \quad (w_1 \neq w_2 \in [n^2], y \in D) \quad (12)$$

where, on the right, x and $x_1 \neq x_2$ range over $[n]$.

Second, we show that if w goes to some x and every x goes to some y, then w goes to some y. That is,

$$\bigvee_{x\in[n]} P_{wx}, \ldots, \bigvee_{y\in[D]} Q_{xy}, \ldots \rightarrow \bigvee_{y\in D} R_{wy} \quad (w \in [n^2]) \quad (13)$$

where, on the left, x ranges over $[n]$. Applying the cut rule to (11) and all the sequents in (12) and (13) produces the desired result, i.e., sequent (10).

We now examine in more detail the proofs of the sequents in (12) and (13). For the sequents in (12), consider arbitrary values of $w_1 \neq w_2 \in [n^2]$ and $y \in D$. First note that $R_{w_1y}R_{w_2y}$ stands for

$$\bigvee_{x_1\in[n]} \bigvee_{j_1} \bigvee_{k_1} \bigvee_{x_2\in[n]} \bigvee_{j_2} \bigvee_{k_2} P^{(j_1)}_{w_1x_1} Q^{(k_1)}_{x_1y} P^{(j_2)}_{w_2x_2} Q^{(k_2)}_{x_2y}.$$

Now start with the following sequents:

$$P^{(j_1)}_{w_1x_1} Q^{(k_1)}_{x_1y} P^{(j_2)}_{w_2x_2} Q^{(k_2)}_{x_2y} \rightarrow P^{(j_1)}_{w_1x_1} P^{(j_2)}_{w_2x_2} \quad (x_1, x_2 \in [n], j_1, j_2, k_1, k_2) \quad (14)$$

and

$$P^{(j_1)}_{w_1x_1} Q^{(k_1)}_{x_1y} P^{(j_2)}_{w_2x_2} Q^{(k_2)}_{x_2y} \rightarrow Q^{(k_1)}_{x_1y} Q^{(k_2)}_{x_2y} \quad (x_1, x_2 \in [n], j_1, j_2, k_1, k_2) \quad (15)$$

By using the OR-left rule, combine all the sequents in (14) with $x_1 = x_2$ and all the sequents in (15) with $x_1 \neq x_2$. This gives

$$R_{w_1y}R_{w_2y} \rightarrow \ldots, P^{(j_1)}_{w_1x} P^{(j_2)}_{w_2x}, \ldots, \ldots, Q^{(k_1)}_{x_1y} Q^{(k_2)}_{x_2y}, \ldots$$

where, on the right, x and $x_1 \neq x_2$ range over $[n]$ and j_1, j_2, k_1 and k_2 range over all possible values. Several applications of the OR-right rule now yield the desired sequent in (12).

Let us now turn to the proof of the sequents in (13). Let $w \in [n^2]$ be arbitrary. Again, first note that $\bigvee_{y\in D} R_{wy}$ stands for

$$\bigvee_{y\in D} \bigvee_{x\in[n]} \bigvee_{j} \bigvee_{k} P^{(j)}_{wx} Q^{(k)}_{xy}.$$

Start with the following sequents:

$$P^{(j)}_{wx}, Q^{(k)}_{xy} \rightarrow P^{(j)}_{wx} Q^{(k)}_{xy} \quad (x \in [n], y \in D, j, k) \quad (16)$$

By using the OR-left rule and then the OR-right rule, combine the sequents in (16) that correspond to the various values of k:

$$P^{(j)}_{wx}, Q_{xy} \rightarrow \bigvee_{k} P^{(j)}_{wx} Q^{(k)}_{xy} \quad (x \in [n], y \in D, j) \quad (17)$$

Again by using the OR-left rule and then the OR-right rule, combine the sequents in (17) that correspond to the various values of j:

$$P_{wx}, Q_{xy} \rightarrow \bigvee_{j} \bigvee_{k} P^{(j)}_{wx} Q^{(k)}_{xy} \quad (x \in [n], y \in D) \quad (18)$$

Once more, by using the OR-left rule and then the OR-right rule, combine the sequents in (18) that correspond to the various values of y:

$$P_{wx}, \bigvee_{y\in D} Q_{xy} \rightarrow \bigvee_{y\in D} \bigvee_{j} \bigvee_{k} P^{(j)}_{wx} Q^{(k)}_{xy} \quad (x \in [n]) \quad (19)$$

Finally, in a similar way, combine all the sequents in (19) to obtain the desired sequent in (13).

There only remains to say that it is easy to verify that the proof is tree-like and of size $(nst)^{O(1)}$. □

THEOREM 4. *For every set of size-t, depth-1.5 formulas $(P_{xy})_{x\in[n^2],y\in[n]}$ of the form OR of small AND's, $\mathrm{PHP}^{n^2}_n(P)$ has a size-$(nt)^{O(\log n)}$, tree-like, depth-1.5 sequent calculus proof. In particular, if the P_{xy}'s are simple propositional variable, then the size of the proof is $n^{O(\log n)}$.*

PROOF. As mentioned earlier, the proof consists in a sequence of alternations between the reduction and amplification steps formalized in the preceding lemmas. Before describing the proof, first note that these two lemmas also hold when all the sequents involved are weakened by the cedents of $\mathrm{PHP}^{n^2}_n(P)$. This is simply because in every application of any of the inference rules, both the hypotheses and the conclusion can be weakened in this way. In what follows, we assume that all sequents are weakened by the cedents of $\mathrm{PHP}^{n^2}_n(P)$.

We describe the proof in a top-down fashion. Let c be the maximum of all the hidden constants in the statements of Lemmas 2 and 3. First, by Lemma 2, we prove $\mathrm{PHP}^{n^2}_n(P)$ from $\mathrm{PHP}^{A_1}_{B_1}(P)$, ..., $\mathrm{PHP}^{A_n}_{B_1}(P)$ and $\mathrm{PHP}^{n}_{B_2}(R)$, where $A_1, \ldots, A_n$ is the partition of $[n^2]$ into sets of size n, B_1, B_2 is the partition of $[n]$ into sets of size $n/2$, and R is a set of size-(nt), depth-1.5 formulas. In other words, we prove $\mathrm{PHP}^{n^2}_n(P)$ from $n+1$ sequents of the form $\mathrm{PHP}^{C}_{D}(Q)$ where $|C| = n$, $|D| = n/2$ and the Q's are sets of size-(nt), depth-1.5 formulas.

Second, by Lemma 3, we prove each of these sequents from a sequent of the form $\mathrm{PHP}^{n^2}_{D}(R)$ where the R's are sets of size-$c(nt)^2$, depth-1.5 formulas.

We continue using the two lemmas in alternation. In general, it is easy to verify that after k reductions and amplifications, we will be left with proving $(n+1)^k$ sequents of the form $\mathrm{PHP}^{n^2}_{n/2^k}(R)$ where the R's are sets of size-$(cnt)^{2k}$, depth-1.5 formulas.

After $\log n$ steps, we are left with only sequents of the form $\mathrm{PHP}^{n^2}_1(R)$, and these are very easy to prove.

It is easy to see that the entire proof is tree-like. To calculate its size, note that the largest subproofs occur in the last amplification step. There, we have $(n+1)^{\log n}$ proofs of size at most $(cnt)^{2c\log n}$. The total size of the proof is therefore $(nt)^{O(\log n)}$. $\square$

The next lemma formalizes the proof of PHP_n^{2n} from $\mathrm{PHP}_n^{n^2}$ that was outlined at the beginning of this section. The overall structure of the proof is similar to the amplification step that was formalized in Lemma 3.

LEMMA 5. *For every set of size-s, depth-1.5 formulas $(Q_{xy})_{x\in[2n],y\in[n]}$ of the form OR of small AND's, there is a set of size-$(ns)^{O(\log n)}$, depth-1.5 formulas $(R_{wy})_{w\in[n^2],y\in[n]}$ of the form OR of small AND's such that $\mathrm{PHP}_n^{2n}(Q)$ has a size-$(ns)^{O(\log n)}$, tree-like, depth-1.5 sequent calculus proof from $\mathrm{PHP}_n^{n^2}(R)$.*

The main result of this section now follows directly from Lemma 5 and Theorem 4.

THEOREM 6. *For every set of size-t, depth-1.5 formulas $(P_{xy})_{x\in[2n],y\in[n]}$ of the form OR of small AND's, $\mathrm{PHP}_n^{2n}(P)$ has a size-$(nt)^{O(\log n)^2}$, tree-like, depth-1.5 sequent calculus proof. In particular, if the P_{xy}'s are simple propositional variable, then the size of the proof is $(n)^{O(\log n)^2}$.*

5. OPTIMAL DEPTH

In this section we will show how to prove PHP_n^{2n} in depth .5. Note that the statement of PHP_n^{2n} itself has depth 1, so in order for the theorem to make sense, we will need to convert the proof into refutation form. Let Clauses(PHP_n^{2n}) denote the set of depth-0 sequents that underly the pigeonhole principle. That is, Clauses(PHP_n^{2n}) consists of the sequents $\rightarrow P_{i1},\dots,P_{in}$ for each $i\in[2n]$, and the sequents $P_{ik},P_{jk}\rightarrow$ for each $i\neq j\in[2n]$, $k\in[n]$. The following lemma shows that it is easy to convert a proof of PHP_n^{2n} into a refutation of Clauses(PHP_n^{2n}) with no significant change in size or depth.

LEMMA 7. *Let PHP_n^{2n} have a size-s, tree-like, depth-1.5 sequent calculus proof. Then there is a size-$O(s^2)$, tree-like, depth-1.5 refutation of Clauses(PHP_n^{2n}).*

PROOF. Recall that PHP_n^{2n} is the following sequent:

$$\mathrm{PHP}_n^{2n}:\dots,\bigvee_{k\in[n]}P_{ik},\dots\rightarrow\dots,P_{ik}P_{jk},\dots$$

where, on the left, i ranges over $[2n]$ and, on the right, $i\neq j$ range over $[2n]$ and k ranges over $[n]$. Start with the sequents

$$\rightarrow P_{i1},\dots,P_{in}\quad(i\in[2n]).$$

By several applications of the OR-right rule, we get

$$\rightarrow\bigvee_{k\in[n]}P_{ik}\quad(i\in[2n]).$$

Now cut each of these sequents as well as each of the sequents $P_{ik}P_{jk}\rightarrow$ with PHP_n^{2n} to obtain the desired contradiction. The bound on the size of the refutation is easy to verify. $\square$

We will now show how to convert a tree-like refutation of Clauses(PHP_n^{2n}) of depth-1.5 into a (dag-like) refutation of Clauses(PHP_n^{2n}) of depth-.5. The following result is due to Krajíček.

THEOREM 8 ([5]). *Let Q be a set of sequents of depth 0. That is, each sequent in Q is of the form $\Gamma\rightarrow\Delta$ where Γ and Δ are sets of literals. Let d be a nonnegative integer. Suppose that there is a tree-like, depth-$(d+1.5)$ LK refutation of Q of size S. Then Q has a depth-$(d+.5)$ LK refutation of size polynomial in S.*

For completeness, we include the proof for the case of reducing the depth from 1.5 to .5.

PROOF. Consider an arbitrary sequent in the depth-1.5 LK refutation, of the form:

$$\Gamma,\bigvee_i A_i^1,\dots,\bigvee_i A_i^m,\bigwedge_i C_i^1,\dots,\bigwedge_i C_i^m$$
$$\rightarrow\bigvee_i B_i^1,\dots\bigvee_i B_i^m,\bigwedge_i D_i^1,\dots\bigwedge_i D_i^m,\Delta$$

where Γ and Δ are sets of formulas of depth at most .5, and A_i^j, B_i^j, C_i^j and D_i^j are formulas of depth .5.

Let P be the tree-like, depth-1.5 LK refutation, and let P_k denote the first k lines of P. Assume that s_k is the sequent at line k, and assume without loss of generality that it has the above form. We will prove by induction on k that P_k can be efficiently converted into a dag-like, depth-.5 proof of

$$\Gamma,C_1^1,\dots C_q^1,\dots C_1^m,\dots C_q^m$$
$$\rightarrow B_1^1,\dots,B_q^1,\dots B_1^m,\dots,B_q^m,\Delta$$

from axioms:

$$\begin{array}{ll}
(1a) & \rightarrow A_1^1,\dots,A_q^1\\
(2a) & \rightarrow A_1^2,\dots,A_q^2\\
 & \dots\\
(ma) & \rightarrow A_1^m,\dots,A_q^m\\
 & \text{and}\\
(1b) & D_1^1,\dots,D_q^1\rightarrow\\
(2b) & D_1^2,\dots,B_q^2\rightarrow\\
 & \dots\\
(mb) & D_1^m,\dots,D_q^m\rightarrow.
\end{array}$$

This suffices to prove the theorem since the final line has depth 0.

When $k=1$, s_k is an axiom, of the form $x\rightarrow x$, so the inductive statement holds. Now suppose that the k^{th} line follows from two previous lines by a rule. The two rules requiring work are the cut-rule, and the $\vee$-right rule. First, suppose that the two previous lines have the form:

$$\Gamma, \bigvee_i A_i^1, \ldots, \bigvee_i A_i^m, \bigwedge_i C_i^1, \ldots, \bigwedge_i C_i^m, \bigvee_i E_i$$
$$\rightarrow \Delta, \bigvee_i B_i^1, \ldots, \bigvee_i B_i^m, \bigwedge_i D_i^1, \ldots, \bigwedge_i D_i^m$$

and

$$\Gamma, \bigvee_i A_i^1, \ldots, \bigvee_i A_i^m, \bigwedge_i C_i^1, \ldots, \bigwedge_i C_i^m$$
$$\rightarrow \Delta, \bigvee_i B_i^1, \ldots, \bigvee_i B_i^m, \bigwedge_i D_i^1, \ldots, \bigwedge_i D_i^m, \bigvee_i E_i$$

In the above notation, the formulas in Γ, Δ, and A_i^j, B_i^j, C_i^j, D_i^j and E_i have depth .5. And the k^{th} line, obtained by cutting on $\bigvee_i E_i$, has the form:

$$\Gamma, \bigvee_i A_i^1, \ldots, \bigvee_i A_i^m, \bigwedge_i C_i^1, \ldots, \bigwedge_i C_i^m$$
$$\rightarrow \Delta, \bigvee_i B_i^1, \ldots, \bigvee_i B_i^m, \bigwedge_i D_i^1, \ldots, \bigwedge_i D_i^m$$

By induction, there is a dag-like, depth-.5 proof, Q_1, of

$$\Gamma, C_1^1, \ldots, C_q^1, \ldots, C_1^m, \ldots, C_q^m$$
$$\rightarrow \Delta, B_1^1, \ldots, B_q^1, \ldots, B_1^m, \ldots, B_q^m$$

from axioms (1a) through (ma) and (1b) through (mb), and $\rightarrow E_1, \ldots E^q$, and a dag-like, depth-.5 proof, Q_2, of

$$\Gamma, C_1^1, \ldots C_q^1, \ldots, C_1^m, \ldots, C_q^m$$
$$\rightarrow \Delta, B_1^1, \ldots, B_q^1, \ldots, B_1^m, \ldots, B_q^m, E^1, \ldots, E_q$$

from axioms (1a) through (ma) and (1b) through (mb). We want to combine Q_1 and Q_2 to obtain a dag-like, depth-.5 proof, Q, of

$$\Gamma, C_1^1, \ldots C_q^1, \ldots C_1^m, \ldots, C_q^m$$
$$\rightarrow \Delta, B_1^1, \ldots, B_q^1, \ldots, B_1^m, \ldots, B_q^m$$

from axioms (1a) through (ma) and (1b) through (mb). Replacing each axiom of the form $\rightarrow E_1, \ldots E_q$ in Q_1 by the entire Q_2 proof gives the desired proof Q.

The other case, where the cut rule is applied to $\bigwedge_i E_i$, and the $\vee$-right rule is proven similarly; the other rules require little or no modifications. $\square$

This result, combined with Theorem 6 and Lemma 7, gives the main theorem of this section.

THEOREM 9. *The propositional weak pigeonhole principle,* PHP_n^{2n}, *has size-*$n^{O(\log n)^2}$, *depth-*.5 *LK proofs. Also, the first-order version of the weak pigeonhole principle,* $\mathrm{PHP}_n^{2n}(R)$, *has* $T_2^2(R)$ *proofs.*

Our upper bound is optimal with respect to depth since it is known that depth-0 proofs, i.e., Resolution proofs, of PHP_n^{2n} require exponential size [4]. In addition, our upper bound is tight in another sense: the proof cannot be made tree-like, unless the size becomes exponential, as the following theorem shows.

THEOREM 10 ([6; 14]). *For sufficiently large* n, *if* P *is a tree-like LK refutation of* PHP_n^m, *where each formula in* P *involves at most* k *variables, then* P *has size at least* $2^{\lfloor n/2k \rfloor}$.

The results of [6; 14] are very elegant and apply to a large class of formulas. However, the exact form of the lower bound for the weak pigeonhole principle is not made explicit and their proof is more complicated than needed for the particular case that concerns us. Therefore, we will give here a simpler proof of the theorem, one that extends the lower bound for tree-like Resolution given in [3].

PROOF. The proof will consist of two stages.

1. Show that if there is a small tree-like, depth-.5 LK refutation of PHP_n^m, then there is a decision tree of the same structure, with nodes queried by decisions of the form $f(X) = 0/1$, where f is a function, and X is a set of at most k variables upon which f depends, with the property that each leaf is labeled by some clause of PHP_n^m that is falsified.
2. Show that any such decision tree for PHP_n^m has to be large.

We will prove the first step by induction on the size of the proof. The only rules that really matter are the ones that take two sequents to one sequent: these are AND-right, OR-left and cut.

First, suppose we derive $C = \Gamma \rightarrow \Delta$ from $A = \Gamma, g \rightarrow \Delta$ and $B = \Gamma \rightarrow g, \Delta$ by an application of the cut rule. Consider an assignment α that makes C false. Then if $g(\alpha)$ is false, then B is false. Otherwise, if $g(\alpha)$ is true, then A is false. So we label this node with g. Since the proof has depth .5, g is a function involving at most k variables, and so satisfies the conditions required of the decision tree.

Now suppose we derive $C = \Gamma \rightarrow \wedge(A_1 \ldots A_n), \Delta$ from $A = \Gamma \rightarrow A_1, \Delta$ and $B = \Gamma \rightarrow \wedge(A_2, \ldots A_k), \Delta$ by an application of the AND-right rule. Consider an assignment α that makes C false. This implies that $\wedge(A_1, \ldots, A_n)(\alpha)$ is false. Now if $A_1(\alpha)$ is false, then A is false. On the other hand, if $A_1(\alpha)$ is true, then $\wedge(A_2, \ldots A_n)$ is false and thus B is false. So we can label this node with A_1. The OR-left rule is handled in a similar way.

We will now prove step two. We want to show that any decision tree for solving the search problem associated with PHP_n^m, where the queries made are of the form $f(X)$, where each f depends on at most k variables, must have size at least $2^{\lfloor n/2k \rfloor}$.

Consider the critical truth assignments (cta's) where n pigeons are mapped to n holes, and the remaining $m - n$

pigeons are unassigned. Consider the restricted tree T, where we only care about paths that are followed by at least one critical truth assignment. Now we want to claim that T must be large.

We want to prove that along any path in T, the number of branching nodes must be at least $\lfloor n/2k \rfloor$, and hence the total size of T is at least $2^{\lfloor n/2k \rfloor}$. We will prove it by induction on n. When $n = 0$, any of the m pigeons is a valid answer, and the size is therefore 1.

Now suppose $n > 0$, and assume that Q is a decision tree for PHP^m_n. Let $f(X)$ be the first query in Q, and suppose that the the left subtree of Q is labeled by $f(X) = 0$ and the right subtree of Q is labeled by $f(X) = 1$. If all cta's are such that $f(X) = 0$, then proceed on the left subtree. Similarly if all cta's are such that $f(X) = 1$, then proceed on the right subtree.

Otherwise, $f(X)$ splits up the problem in two pieces in a nontrivial way. First consider the left subtree, the one labeled by $f(X) = 0$. In this case, we want to find a restriction ρ_0 so that: (1) $f(X)$ is forced to 0 by ρ_0, and (2) ρ_0 is a partial one-to-one map from at most $2k$ pigeons to holes. To obtain ρ_0, since $f(X)$ is forced to 0 by some cta, select an assignment to the variables of X consistent with one of these cta's. Then minimally extend the assignment so that we are left with a partial assignment ρ_0 that leaves m' unassigned pigeons and n' unassigned holes, and the remaining pigeons are mapped in a one-to-one way onto the remaining holes. Since $|X| \leq k$, at most k pigeons and at most k holes are mentioned by ρ_0, and therefore the extended assignment leaves $m' \geq m - 2k$ and $n' \geq n - 2k$. Now applying ρ_0, it follows that the left subtree, Q_0, solves the decision problem for $\mathrm{PHP}^{m'}_{n'}$, where $m' = m - 2k$, $n' = n - 2k$. By the inductive hypothesis it follows that any path of Q_0 must have at least $\lfloor (n-2k)/2k \rfloor$ branching nodes.

Similarly, for the right subtree (labeled $f(X) = 1$), we can find a restriction so that $f(X)$ is forced to 1 by ρ_1 and ρ_1 is a partial map from at most $2k$ pigeons to holes. Applying ρ_1 it follows that the right subtree Q_1 solves the decision problem for $\mathrm{PHP}^{m'}_{n'}$, and again by the inductive hypothesis, any path in Q_1 must have at least $\lfloor (n-2k)/2k \rfloor$ branching nodes.

Thus, in total, it follows that any path in Q has at least $\lfloor n/2k \rfloor$ branching nodes, and thus the size of Q is at least $2^{\lfloor n/2k \rfloor}$. □

6. DISCUSSION AND RELATED RESULTS

We summarize what is currently known in Table 1. The symbol * in the References column indicates the current paper. All of the lower bounds are exponential in n. (Some of these are actually proven generally, as a function of n and m.)

For depth 0 (Resolution proofs), the best known upper bound are polynomial-size proofs of PHP^m_n, where $n \leq (\log m)^2/\log\log m$ [3]. As mentioned in the introduction, prior to the result of this paper, the only nontrivial constant-depth LK proof of the weak pigeonhole principle was that of [9], and the optimization with respect to depth of [5]. Krajíček also shows that the proof of [9] can also be modified to give depth-.5 LK proofs of the *onto* pigeonhole principle.

When $m - n = O(1)$, it is known that any bounded-depth LK proof of PHP^m_n requires exponential size. Moreover, it is known that even if one adds the onto pigeonhole principle as an axiom scheme, there is still no subexponential, bounded-depth LK proof of PHP^m_n.

In this paper, we showed how to prove PHP^{2n}_n with depth-.5, quasipolynomial-size LK proofs. It is not known whether or not there are constant-depth, *polynomial*-size LK proofs of the weak pigeonhole principle. If we restrict attention to polynomial-size proofs, then all that is known is that one can prove PHP^m_n in constant depth, where $n = \text{polylog } m$. Moreover, the depth of the proof is dependent on the the exponent in the polylog m.

Lastly, by formalizing circuits that count, one can prove PHP^m_n for any $n < m$ with polynomial-size Frege proofs [2].

There are many interesting open problems that are raised by this work. Most importantly, are there polynomial-size, constant-depth proofs of either the weak pigeonhole principle, or the onto weak pigeonhole principle? As mentioned in the introduction, a sufficiently uniform positive answer would answer a longstanding open question of [15] about the provability of infinitely many primes in $I\Delta_0$.

The original proof of [9] actually shows that $\mathrm{PHP}^{n^2}_n$ has depth-d, size-$n^{\log^{\Omega(d)} n}$ proofs. That is, as d increases, the size is reduced. We do not know how to extend our new proof to this more general situation.

In the introduction, we mentioned a close connection between the weak pigeonhole principle and approximate counting. Here we elaborate further on this connection and a related open problem. Buss's Frege proof [2] of PHP^{n+1}_n views the pigeonhole variables as a bipartite graph with pigeons on the left and holes on the right. If every pigeon maps to at least one hole, then the number of edges out of the left side of the graph is at least m. To say this, we construct a polynomial-size circuit that counts the number of 1's in a binary string with one index for each of the edges of the graph, and prove inductively (using the pigeon axioms) that this circuit outputs a number of size at least m. Similarly, if each hole has at most one pigeon mapped to it, then the number of edges into the right side of the graph is at most n, and again we say this by proving inductively (using the hole axioms) that the counting circuit outputs a number of size at most n. Finally, if $m > n$, this gives us the desired contradiction.

Using a pairwise independent collection of hash functions, approximating the number of 1's in a binary string is computable with bounded-depth, polynomial-size circuits. It is tempting to try to use such circuits to prove the weak pigeonhole principle, in a similar manner to the above argument of Buss. However, the proofs of correctness of all known constructions involve probablistic counting

	Upper Bounds			Lower Bounds	
	Ref.	$n \leq$	Size	Ref.	$n \geq$
Resolution	[3]	$\frac{(\log m)^2}{\log\log m}$	poly(m)	[3; 13]	$\frac{(\log m)^2}{\log\log m}$ (tree-like)
(depth-0 LK)	*	$\frac{\log m \log S}{\log\log S}$	poly(S)	[4]	$m^{1/2+\epsilon}$
Depth-.5 LK	*	$\sqrt{m}$	$m^{O(\log m)}$	[6; 14]	$\sqrt{m}$ (tree-like)
	*	$m/2$	$m^{O(\log m)^2}$		
Depth-1.5 LK	[9; 5]	$m/2$	$m^{O(\log m)}$	[10; 7; 1]	$m - m^{1/480}$
Depth-c LK	[8]	polylog m	poly(m)	[10; 7; 1]	$m - m^{o(1)}$
	[9; 5]	$m/2$	$m^{\log^{O(c)} m}$		

Table 1: Summary of related results.

and hence rely on the weak pigeonhole principle to prove correctness. It is not clear whether this can be avoided. We conjecture that it is not possible to prove the weak pigeonhole principle with polynomial-size, small-depth (say depth 2 or 3) Frege proofs. Such a result would be very striking, as it would be the first instance where there are known explicit constructions of circuits computing a function (in this case approximate counting), but where *any* proof of correctness of the function cannot be carried out in an equally feasible way.

Lastly, a question left open in [4], are there polynomial-size Resolution proofs of PHP_n^m, when $m \geq n^2/\log n$? We conjecture that the answer is no.

7. ACKNOWLEDGEMENTS

We would like to thank Jan Krajíček for stimulating discussions that led to the writing of this paper. We also thank him, Paul Beame and Sam Buss for useful comments on an earlier version of the paper.

8. REFERENCES

[1] P. Beame and S. Riis. More on the relative strength of counting principles. In P. W. Beame and S. R. Buss, editors, *Proof Complexity and Feasible Arithmetics*, volume 39 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 13–35. American Mathematical Society, 1998.

[2] S. R. Buss. Polynomial size proofs of the propositional pigeonhole principle. *J. Symbolic Logic*, 52:916–927, 1987.

[3] S. R. Buss and T. Pitassi. Resolution and the weak pigeonhole principle. In *Computer Science Logic*, pages 149–156, 1997.

[4] S. R. Buss and G. Turán. Resolution proofs of generalized pigeonhole principles. *Theoret. Comput. Sci.*, 62:311–317, 1988.

[5] J. Krajíček. Lower bounds to the size of constant-depth propositional proofs. *J. Symbolic Logic*, 59:73–86, 1994.

[6] J. Krajíček. On the weak pigeonhole principle. Manuscript, August 1999.

[7] J. Krajíček, P. Pudlák, and A. Woods. Exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Structures and Algorithms*, 7:15–39, 1995.

[8] J. B. Paris and A. J. Wilkie. Counting problems in bounded arithmetic. In *Methods in Mathematical Logic: Proceedings of the 6th Latin American Symposium on Mathematical Logic 1983*, volume 1130 of *Lectures Notes in Mathematics*, pages 317–340, 1985.

[9] J. B. Paris, A. J. Wilkie, and A. R. Woods. Provability of the pigeonhole principle and the existence of infinitely many primes. *J. Symbolic Logic*, 53:1235–1244, 1988.

[10] T. Pitassi, P. Beame, and R. Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Comput. Complexity*, pages 97–140, 1993.

[11] A. A. Razborov. Unprovability of lower bounds on the circuit size in certain fragments of bounded arithmetic. *Izvestiya of the R.A.N.*, 59(1):201–224, 1995.

[12] A. A. Razborov and S. Rudich. Natural proofs. *J. Comput. System Sci.*, 55:24–35, 1997.

[13] A. A. Razborov, A. Wigderson, and A. C. Yao. Read-once branching programs, rectangular proofs of the pigeonhole principle and the transversal calculus. In *Proceedings of the 29th ACM Symposium on Theory of Computing*, pages 739–748, 1997.

[14] S. Riis. A complexity gap for tree-resolution. Manuscript, September 1999.

[15] A. J. Wilkie. Some results and problems on weak systems of arithmetic. In A. Macintyre, L. Pacholski, and J. Paris, editors, *Logic Colloquium '77*, pages 237–248. North-Holland, 1978.