# Bounded-depth Frege lower bounds for weaker pigeonhole principles


Josh Buresh-Oppenheim
Computer Science Department
University of Toronto
Toronto, ON M5S 3G4
bureshop@cs.toronto.edu

Paul Beame*
Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350
beame@cs.washington.edu

Toniann Pitassi†
Computer Science Department
University of Toronto
Toronto, ON M5S 3G4
toni@cs.toronto.edu

Ran Raz‡
Faculty of Math and Computer Science
Weizmann Institute of Science
Rehovot 76100, Israel
ranraz@wisdom.weizmann.ac.il

Ashish Sabharwal*
Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350
ashish@cs.washington.edu



## Abstract

*We prove a quasi-polynomial lower bound on the size of bounded-depth Frege proofs of the pigeonhole principle $PHP^m_n$ where $m = (1 + 1/\mathit{polylog}\, n)n$. This lower bound qualitatively matches the known quasi-polynomial-size bounded-depth Frege proofs for these principles. Our technique, which uses a switching lemma argument like other lower bounds for bounded-depth Frege proofs, is novel in that the tautology to which this switching lemma is applied remains random throughout the argument.*


## 1 Introduction

The propositional pigeonhole principle asserts that $m$ pigeons cannot be placed in $n$ holes with at most one pigeon per hole whenever $m$ is larger than $n$. It is an exceptionally simple fact that underlies many theorems in mathematics, and is the most extensively studied combinatorial principle in proof complexity. (See [17] for an excellent survey on the proof complexity of pigeonhole principles.) It can be formalized as a propositional formula, denoted $PHP^m_n$, in a standard way which we shall describe later in the paper.

Clearly for $m > n$ the pigeonhole principle is valid and proving super-polynomial lower bounds on the length of propositional proofs of the pigeonhole principle when $m = n + 1$ has been a major achievement in proof complexity. There are several ways in which the principle can be made weaker (and hence easier to prove). One way is to increase the number of pigeons with respect to the number of holes. Another is to put more constraints on how the pigeons can be placed in the holes. The onto-PHP and the functional-PHP are two well-studied weakenings, which assert, respectively, that the mapping from pigeons to holes must be surjective and a well-defined function. In this paper, we will prove lower bounds which apply to all of these variations of the basic PHP.

For all $m > n$, Buss has given polynomial-size Frege proofs of $PHP^m_n$. The basic idea is to formulate polynomial-size circuits that count the number of 1's in an $N$-bit string. Then using basic axioms about the counting circuit plus the pigeon clauses, derive that the counting circuit applied to the vector of variables $P_{i,j}$ outputs a number that is larger than $n$, while on the other hand, using the basic axioms about the circuit plus the hole clauses, derive that the counting circuit applied to the same vector of variables, outputs a number no larger than $n$.

In weaker proof systems, however, the proof complexity of the pigeonhole principle depends in a crucial way on the number of pigeons, $m$, as a function of the number of holes, $n$. As $m$ increases, the principle becomes weaker (easier to prove) and in turn the proof complexity question becomes more difficult. We review the basics of what is known for Resolution and bounded-depth Frege systems below. Generally, the weak pigeonhole principle has been used to refer to $PHP^m_n$ whenever $m$ is at least a constant factor larger than $n$. We will be


*Research supported by NSF grant CCR-0098066
†Research supported by US-Israel BSF grant 98-00349
‡Research supported by US-Israel BSF grant 98-00349

primarily concerned with forms of the pigeonhole principle that are significantly weaker than the usual pigeonhole principle but somewhat stronger than these typical weak forms.

For the Resolution proof system, the complexity of the pigeonhole principle is essentially resolved. In 1985, Haken proved the first super-polynomial lower bounds for unrestricted Resolution proofs of $PHP_n^m$, for $m = n+1$ [8]. This lower bound was generalized by Buss and Turan [7] for $m < n^2$. For the next 10 years, the resolution complexity of $PHP_n^m$ for $m \geq n^2$ was completely open. A recent result due to Raz gave exponential Resolution lower bounds for the weak pigeonhole principle, and subsequently Razborov has resolved the problem for most interesting variants of the PHP. (See [15, 18].)

Substantially less is known about the complexity of the pigeonhole principle in bounded-depth Frege systems, although strong lower bounds are known when the number of pigeons $m$ is close to the number of holes $n$. Ajtai proved super-polynomial lower bounds for $PHP_n^{n+1}$ with an ingenious blend of combinatorics and nonstandard model theory [1]. This result was improved to exponential lower bounds in [5]. It was observed in [3] that the above lower bounds can in fact be applied to $PHP_n^m$ for $m \leq n + n^{\varepsilon}$, for some $\varepsilon$ that falls off exponentially in the depth of the formulas involved in the proof.

For the case of larger $m$ (the topic of this paper), the complexity of bounded-depth Frege proofs of $PHP_n^m$ is slowly emerging, with surprising and interconnected results. There are several deep connections between the complexity of the weak pigeonhole principle and other important problems. First, lower bounds for bounded-depth Frege proofs of the weak pigeonhole principles suffice to show unprovability results for the $P$ versus $NP$ statement. Secondly, the long-standing question of whether or not the existence of infinitely many primes has an $I\Delta_0$ proof is closely related to the complexity of WPHP in bounded-depth Frege systems [14]. Thirdly, the question is closely related to the complexity of approximate counting [13].

There are two significant prior results concerning the proof complexity of weak pigeonhole principles in bounded-depth Frege systems more powerful than resolution. There are bounded-depth Frege proofs of $PHP_n^m$ for $m$ as small as $n + n/\text{polylog}\, n$ of quasi-polynomial size [14, 10, 12]; thus exponential lower bounds for the weak pigeonhole principle are out of the question. In fact, this upper bound is actually provable in a very restricted form of bounded-depth Frege where all lines in the proof are disjunctions of polylog $n$-sized conjunctions, a proof system known as $Res(\text{polylog}\, n)$. On the other hand, [2] shows exponential lower bounds for weak pigeonhole principles in $Res(2)$, a proof system which allows lines to be disjunctions of size-2 conjunctions.

In this paper we prove quasi-polynomial lower bounds for the weak pigeonhole principle whenever $m \leq n + n/\text{polylog}\, n$. More precisely, we show that given integers $c$ and $h$ such that $c$ is sufficiently large compared to $h$, there exists an integer $a > 1$ such that any depth-$h$ proof of $PHP_n^m$, where $m \leq n + n/\log^c n$, requires size $2^{\log^a n}$. This is a substantial improvement over previous lower bounds. Our proof technique applies a switching lemma to a weaker tautology based on certain bipartite graphs. This type of tautology was introduced by [6]. Although we rely heavily on the simplified arguments presented in [4] and [19], in a major difference from previous switching-lemma-based proofs, both the tautologies themselves and the restrictions we consider remain random throughout most of the argument.

In the next section we give an overview of the key issues involved in extending bounded-depth Frege lower bounds to the weak pigeonhole principle and how we deal with them. We follow this with fairly standard definitions of bounded-depth Frege proof systems and the precise formulation of the tautologies we use in section 3, definitions for restrictions and matching decision trees in section 4. We present the lower bound argument itself in section 5.

# 2 Overview

The high-level schema of our proof is not new. Ignoring parameters for a minute, we start with an alleged proof of $PHP_n^m$ of small size. We then show that assigning values to some of the variables in the proof leaves us with a sequence of formulas, each of which can be represented as a particular type of decision tree of small height. This part of the argument is generally referred to as the switching lemma. We then prove that the leaves of any such short tree corresponding to a formula in the proof must all be labelled 1 if the proof is to be sound. Finally, we show that the tree corresponding to $PHP_n^m$ has leaves labelled 0, which is a contradiction since it must appear as a formula in the alleged proof. We now overview the components of lower bound in more detail and context.

The lower bounds for bounded-depth Frege proofs of $PHP_n^{n+1}$ [5] used *restrictions*, partial assignments of values to input variables, and iteratively applied "switching lemmas" with respect to random choices of these restrictions. The first switching lemmas showed that after one applies a randomly chosen restriction that assigns values to many, but far from all, of the input variables with high probability one can convert an arbitrary DNF formula with small terms into a CNF formula with small clauses (hence the name). More generally, such switching lemmas allow one to convert arbitrary DNF formulas with small terms into small height decision trees (which implies the conversion to CNF formulas with small clauses). The basic idea is that for each level of the formulas/circuits, one proves that a randomly chosen restriction will succeed with positive probability for all sub-formulas/gates at that level. One then fixes such a restriction for that level and continues to the next level. To obtain a lower bound one chooses a family of restrictions suited to the target of the analysis. In the case of $PHP^m_n$, the natural restrictions to consider correspond to partial matchings between pigeons and holes.

The form of the argument by which switching lemmas are proven generally depends on the property that the ratio of the probability that an input variable remains unassigned to the probability that it is set to 0 (respectively, to 1) is sufficiently less than 1. In the case of a random partial matching that contains $(1-p)n$ edges applied to the variables of $PHP^m_n$, there are $pn$ unmatched holes and at least $pm$ unmatched pigeons. Hence, the probability that any edge-variable remains unassigned (i.e. neither used nor ruled out by the partial matching) is at least $p^2$. However, the partial matching restrictions set less than a $1/m$ fraction of variables to 1. Thus the proofs required that $p^2n < p^2m < 1$ and thus $p < n^{-1/2}$. This compares with choices of $p = n^{-O(1/h)}$ for depth $h$ circuit lower bounds in the best arguments for parity proven in [9]. Hence, the best known lower bounds on the size of depth-$h$ circuits computing parity is of the form $2^{n^{1/h}}$, while the best known lower bound on the size of depth-$h$ proofs of $PHP^{n+1}_n$ is of the form $2^{n^{2^{-h}}}$.

A problem with extending the lower bounds to $PHP^m_n$ for larger $m$ is that, after a partial matching restriction is applied, the absolute difference between the number pigeons and holes does not change but the number of holes is dramatically reduced. This can qualitatively change the ratio between pigeons and holes. If this is too large then the probability that variables remain unassigned grows dramatically and, in the next level, the above argument does not work at all. For example, with the above argument, if the difference between the number of pigeons and holes is as large as $n^{3/4}$ then after only one round the above argument will fail. The extension in [3] to lower bound proofs for $PHP^{n+n^{\varepsilon_h}}_n$ for formulas of depth $h$ relies on the fact that even after $h$ rounds of restrictions the gap is small enough that there is no such qualitative change; but this is the limit using the probabilities as above.

We are able to resolve the above difficulties for $m$ as large as $n + n/\text{polylog}\, n$. In particular, we increase the probability that variables are set to 1 to $1/\text{polylog}\, n$ from $1/m$ by restricting the matchings to be contained in bipartite graphs $G$ of polylog $n$ degree. Thus we can keep as many as $n/\text{polylog}\, n$ of the holes unmatched in each round. This means that, by choosing the exponents in the polylog $n$ carefully as a function of the depth of the formulas, we can tolerate gaps between the number of pigeons and the number of holes that are also $n/\text{polylog}\, n$.

A difficulty with this outline is that one must be careful throughout the argument that the restrictions one chooses do not remove all the neighbors of a node without matching it, which would simplify the pigeonhole principle to a triviality. It is not at all clear how one could explicitly construct low degree graphs such that some simple additional condition on the restrictions that we choose at each stage could enforce the desired property. It is unclear even how one might do this non-constructively because it is not clear what property of the random graph would suffice.

Instead, unlike previous arguments, we do not fix the graph in advance; we keep the input graph random throughout the argument, and consider for each such graph $G$ its associated proof of the pigeonhole principle restricted to $G$. Since we do not know what $G$ is at each stage we cannot simply fix the restriction as we deal with each level; we must keep that random as well. Having done this, we can use simple Chernoff bounds to show that, for almost all combinations of graphs and restrictions, the degree at each level will not be much smaller than the expected degree, so the pigeonhole principle will remain far from trivial. We adjust parameters to reduce the probability that a restriction fails to simplify a given level so that it is much smaller than the number of levels. Then we apply the probabilistic method to the whole experiment involving the graph $G$ as well as the sequence of restrictions.

There is one other technical point that is important in the argument. It does not merely suffice that the degrees of vertices are small in the original graph $G$. In order for the probabilities in the switching lemma argument to work out it is critical that the degrees of vertices in the graph after each level of restriction is applied are decreased significantly at each step. Using another simple Chernoff bound we show that the degrees of vertices given almost all combinations of graphs and restrictions will not be much larger than their expected value and this suffices to yield the decrease in degree.

Overall, our argument is expressed in much the same terms as those in [4, 19], although we find it simpler to omit formally defining $k$-evaluations as separate entities. One way of looking at our technique is that we apply two very different kinds of random restrictions to a proof of $PHP^m_n$: first, one that sets many variables to 0, corresponding to the restriction of the problem to the graph $G$, and then, one that sets partial matchings for use with the switching lemma.

# 3 The proof system

Let $G = (V_1 \cup V_2, E)$ be a bipartite graph where $|V_2| = n$ and $|V_1| = m > n$.

DEFINITION 3.1. $L(G)$ is the language built from the set of propositional variables $\{X_e : e \in E\}$, the connectives $\{\vee, \neg\}$ and the constants 0 and 1.

DEFINITION 3.2. A *formula* is a tree whose internal nodes are labelled by either $\vee$ (fanin 2) or $\neg$ (fanin 1) and whose leaves are labelled by variables. Given a node in this tree, the full tree rooted at that node is called a (not necessarily proper) *subformula* of the original formula.

DEFINITION 3.3. If the formula contains no connectives, then it has *depth* 0. Otherwise, the *depth* of a (sub)formula $A$ is the maximum number of alternations of connectives along any path from the root to leaf, plus one.

A Frege proof system is specified by a finite set of sound and complete rules. Our results hold for any Frege system, but one common variant is the following, due to Schoenfield ($p,q,r$ stand for formulas and we will use the notation $p,q \vdash r$ to denote that $p$ and $q$ yield $r$ in one step of the Frege system):
Excluded Middle: $\vdash \neg p \vee p$, Expansion Rule: $p \vdash q \vee p$, Contraction Rule: $p \vee p \vdash p$, Associative Rule: $p \vee (q \vee r) \vdash (p \vee q) \vee r$, Cut Rule: $p \vee q,\ \neg p \vee r \vdash q \vee r$.

Let $\mathcal{F}$ be one such Frege system.

DEFINITION 3.4. The *size* of a Frege rule is the number of distinct subformulas mentioned in the rule.

For example, the size of the cut rule above is 7; the subformulas mentioned are: $p, q, r, \neg p, p \vee q, \neg p \vee r, q \vee r$. We'll assume that the rules of $\mathcal{F}$ all have size at most some constant $f$.

DEFINITION 3.5. A *proof* of a formula $A$ in $\mathcal{F}$ is a sequence of formulas $A_1, \dots, A_r = A$ such that $\vdash A_1$ and for all $i > 1$ there is some (possibly empty) subset $\mathcal{A} \subset \{A_1, \dots, A_{i-1}\}$ such that $\mathcal{A} \vdash A_i$.

DEFINITION 3.6. For a Frege proof $\Pi$, let $cl(\Pi)$ denote the closure of the set of formulas in $\Pi$ under subformulas.

For convenience, we introduce the following normal form of a formula:

DEFINITION 3.7. Consider a formula of depth $h$, $A$. The *merged form* of $A$ is the tree such that all $\vee$'s of depth $h$ (with the same parent) are identified into a single node of unbounded fanin, also labelled $\vee$.

DEFINITION 3.8. The *size* of a proof is the total number of distinct subformulas that appears in the proof. The *depth* of a proof is the maximum of the depths of the formulas in the proof.

The following is a formulation of the onto and functional weak pigeonhole principle on the graph $G$. Note that if $G$ is not the complete graph $K_{m,n}$, then this principle is weaker than the standard onto and functional weak pigeonhole principle.

DEFINITION 3.9. $WPHP(G)$ is the OR of the following four (merged forms of) formulas in $L(G)$. In general, $i,j,k$ represent vertices in $G$ and $\Gamma(i)$ represents the set of neighbors of $i$ in $G$.

1. $\bigvee_{(e,e') \in I} \neg(\neg X_e \vee \neg X_{e'})$ for $I = \{(e,e') : e,e' \in E; e = \{i,k\}, e' = \{j,k\}; i,j \in V_1; i \neq j; k \in V_2\}$: two different pigeons go to the same hole.
2. $\bigvee_{(e,e') \in I} \neg(\neg X_e \vee \neg X_{e'})$ for $I = \{(e,e') : e,e' \in E; e = \{k,i\}, e' = \{k,j\}; i,j \in V_2; i \neq j; k \in V_1\}$: one pigeon goes to two different holes.
3. $\bigvee_{i \in V_1} \neg \bigvee_{j \in \Gamma(i)} X_{\{i,j\}}$: some pigeon doesn't have a hole.
4. $\bigvee_{j \in V_2} \neg \bigvee_{i \in \Gamma(j)} X_{\{i,j\}}$: some hole remains empty.

In fact, we consider an arbitrary orientation of the above formula whereby each $\vee$ is binary.

# 4 Representing matchings by trees

In this section we make minor modifications to standard definitions from [4, 19] to apply to the edge variables given by bipartite graphs and not just complete bipartite graphs.

Let $G = (V_1 \cup V_2, E)$ be a bipartite graph where $|V_2| = n$ and $|V_1| = m > n$. Assume there is an ordering on the nodes of $G$. Assume further that we have a Boolean variable $X_e$ for each $e \in E$. Call this set of variables $D$.

DEFINITION 4.1. Two edges of $G$ are said to be *inconsistent* if they share exactly one endpoint. Two partial matchings $\rho_1, \rho_2$ on the graph $G$ are said to be *consistent* if no edge in $\rho_1$ is inconsistent with an edge in $\rho_2$.

DEFINITION 4.2. For $\rho$ a partial matching on the graph $G$ that matches nodes $V_1' \subset V_1$ to nodes $V_2' \subset V_2$, we define $G|_\rho$ as the bipartite graph $((V_1 \setminus V_1') \cup (V_2 \setminus V_2'), E - (V_1' \times V_2 \cup V_1 \times V_2'))$.

DEFINITION 4.3. A *matching decision tree* $T$ for $G$ is a tree where each internal node $u$ is labelled by a node of $G$, $v$, and each edge from a node $u$ is labelled by an edge of $G$ that touches $v$. Furthermore, given any path in the tree from the root to a node $u$, the labels of the edges along the path constitute a partial matching on $G$, called $path(u)$. Let $path(T) = \{path(u)\ :\ u \text{ is a leaf of } T\}$. If $v$ is a node of $G$ that appears as a label of some node in $T$, then $T$ is said to *mention* $v$.

Furthermore, each leaf of $T$ is labelled by 0 or 1 (if a tree satisfies the above conditions but its leaves remain unlabelled, we will call it a *leaf-unlabelled* matching decision tree). Let $T^c$ be the same as $T$ except with the value of each leaf-label flipped. If $U$ is the set of leaves of $T$ labelled 1, let $disj(T)$ be the DNF formula $\bigvee_{u \in U} \bigwedge_{e \in path(u)} X_e$.

DEFINITION 4.4. A *complete* (leaf-unlabelled) matching decision tree for $G$ is one in which, for each internal node $u$ labelled $v$, the set $\{path(u')\ :\ u' \text{ a child of } u\}$ constitutes all matchings in $G$ of the form $path(u) \cup \{\{v, v'\}\}$ for all $v'$ such that $\{v, v'\} \in E$.

DEFINITION 4.5. Let $K$ be a subset of the nodes in $G$. The *full matching tree* for $K$ over $G$ is a leaf-unlabelled matching decision tree for $G$ defined inductively: if $K = \{k\}$, then the root of the tree is labelled by $k$ and, for each edge $e$ in $G$ that touches $k$, there is an edge from the root of the tree labelled $e$. If $K$ contains more than one node, let $k$ be its largest node under the ordering of nodes and assume we have a full matching tree for $K \setminus \{k\}$. For each (unlabelled) leaf $u$ of this tree, let $p$ be the path from the root to $u$. The labels of the edges along $p$ constitute a partial matching on $G$. If this partial matching touches $k$, leave $u$ unlabelled. Otherwise, label $u$ by $k$ and and attach an edge to $u$ for each edge in $G$ that touches $k$ and that extends the partial matching.

Note that the full matching tree for any subset $K$ is complete. If the degree of each node in $K$ is at least $|K|$, then the full matching tree for $K$ is guaranteed to mention all nodes in $K$. Otherwise, it might not.

**Lemma 1.** *Let $T$ be a complete matching tree for $G$ and let $\rho$ be any partial matching on $G$. Let $d$ be the minimal degree of any node in $G$ mentioned by $T$. If $d > \max\{|\rho|, height(T)\}$, then there is a matching in $path(T)$ that is consistent with $\rho$.*

*Proof.* Assume we have found an internal node $u$ in $T$ labelled by $v$ in $G$ such that $path(u)$ is consistent with $\rho$. We will find a child $u'$ of $u$ such that $path(u')$ is still consistent with $\rho$. Since the degree of $v$ is greater than the size of $\rho$, there is an edge $\{v, v'\}$ in $G$ such that $\{v, v'\}$ is either included in $\rho$ (if $\rho$ touches $v$) or extends $\rho$ (if $\rho$ does not touch $v$). Since $T$ is complete and the degree of $v$ is greater than $height(T)$, the edge $\{v, v'\}$ appears as a label of an edge from $u$ in $T$. □

DEFINITION 4.6. We call $F$ a *matching disjunction* if it is one of the constants 0 or 1, or it is a DNF formula with no negations over the variables $D$ such that the edges of $G$ corresponding to the variables in any one term constitute a partial matching. In the latter case, order the terms lexicographically based on the nodes they touch and the order of the nodes in $G$.

DEFINITION 4.7. For $F$ a matching disjunction, the *restriction* $F|_\rho$ for $\rho$ a partial matching is another matching disjunction generated from $F$ as follows: set any variable in $F$ corresponding to an edge of $\rho$ to 1 and set any variable corresponding to an edge not in $\rho$ but incident to one of $\rho$'s nodes to 0. If a variable in term $t$ is set to 0, remove $t$ from $F$. Otherwise, if a variable in term $t$ is set to 1, remove that variable from $t$.

The DNF $disj(T)$ for a matching decision tree $T$ is always a matching disjunction.

DEFINITION 4.8. A matching decision tree $T$ is said to *represent* a matching disjunction $F$ if, for $l$ a leaf of $T$, $F|_{path(l)} \equiv 1$ when $l$ is labelled 1 and $F|_{path(l)} \equiv 0$ when $l$ is labelled 0.

A matching decision tree $T$ always represents $disj(T)$. Furthermore, if $\rho$ extends some matching $path(l)$ for $l$ a leaf of $T$, then $disj(T)|_\rho \equiv 0$ (1, respectively) if $l$ is labelled 0 (1).

DEFINITION 4.9. Let $F$ be a matching disjunction. We define a tree $Tree_G(F)$ called the *canonical decision tree* for $F$ over $G$: if $F$ is constant, then $Tree_G(F)$ is one node labelled by that constant. Otherwise, let $C$ be the first term of $F$. Let $K$ be the nodes of $G$ touched by variables in $C$. The top of $Tree_G(F)$ is the full matching tree on $K$ over $G$. We replace each leaf $u$ of that tree, with the tree $Tree_{G|_{path(u)}}(F|_{path(u)})$.

The tree $Tree_G(F)$ will have all of its leaves labelled. It is designed to represent $F$ and to be complete.

DEFINITION 4.10. For $T$ a matching decision tree and $\rho$ a matching, $T$ *restricted by* $\rho$, written $T|_\rho$, is a matching decision tree obtained from $T$ by first removing all edges of $T$ that are inconsistent with $\rho$, and retaining only those nodes of $T$ that remain connected to the root of $T$. Each remaining edge that corresponds to an element of $\rho$ is then contracted (its endpoints are identified and labelled by the label of the lower endpoint).

**Lemma 2.** *([19], Lemma 4.8) For $T$ a matching decision tree and $\rho$ a matching:*

*(a)* $disj(T)|_\rho \equiv disj(T|_\rho)$.
*(b)* $(T|_\rho)^c = T^c|_\rho$.
*(c) If* $T$ *represents a matching disjunction* $F$*, then* $T|_\rho$ *represents* $F|_\rho$.

# 5 The lower bound

Let $m = n + n/\log^c n$ for some integer $c > 0$ and let $h > 0$ be an integer. We assume for simplicity that $n$ is large compared to $c$ and that all subsequent expressions are integers. We will show that for any $a$ such that $8^h(a+3) < c$, any proof of $PHP^m_n = WPHP(K_{m,n})$ of depth $h$ is of size greater than $2^{\log^a n}$. To do this we do not work directly with proofs of $WPHP(K_{m,n})$ but rather we work with proofs of $WPHP(G)$ for randomly chosen subgraphs $G$ of $K_{m,n}$.

More precisely, let $b = 8^h(a+3)$, define $d = \log^b n$ and observe that $a < b < c$.

Let $\mathcal{G}(m,n,d/n)$ be the uniform distribution on all bipartite graphs from $m$ nodes to $n$ nodes where each edge is present independently with probability $d/n$.

DEFINITION 5.1. Let $M^\ell(H)$ be the set of all partial matchings of size $\ell$ on a fixed graph $H$.

We now define several sequences of parameters for a probabilistic experiment. The meanings of these parameters will be explained after the definition of the experiment. For initial values, let $m_0 = m$, $n_0 = n$, $b_0 = b$, and $k_0 = 7b_0/8$, $\ell_0 = n_0 - n_0/\log^{k_0} n$. Then, for $1 \le i \le h$, we define recursively:
$m_i = m_{i-1} - \ell_{i-1}$, $n_i = n_{i-1} - \ell_{i-1}$, $b_i = b_{i-1} - k_{i-1}$, $k_i = 7b_i/8$, and $\ell_i = n_i - n_i/\log^{k_i} n$.
In closed form,
$n_i = n/(\log n)^{\sum_{j=0}^{i-1} k_j} = n/(\log n)^{b-b/8^i}$, $m_i = n_i + (m-n)$, $b_i = b - \sum_{j=0}^{i-1} k_j = b/8^i$, $k_i = 7b/8^{i+1}$, and $\ell_i = (1 - 1/\log^{k_i} n)(n/(\log n)^{b-b/8^i})$.

Now we are ready to define the experiment: let $G_0 = G$ be a graph chosen randomly from the distribution $\mathcal{G}(m,n,d/n)$ and choose $\rho_0$ randomly and uniformly from $M^{\ell_0}(G_0)$. Again, for $1 \le i \le h$, let $G_i = G_{i-1}|_{\rho_{i-1}}$ and choose $\rho_i$ randomly and uniformly from $M^{\ell_i}(G_i)$.

First note that the distribution induced on each $G_i$ by this experiment is $\mathcal{G}(m_i,n_i,d/n)$. Furthermore, the expected degree of any pigeon in $G_i$ is $n_i d/n = \log^{b_i} n$. The expected degree of any hole in $G_i$ is $m_i d/n$, which is between $\log^{b_i} n$ and $2\log^{b_i} n$ since $n_i < m_i < 2n_i$ (because $c > b$).

We make several observations about "bad" events in this experiment.

**Lemma 3.** *The probability, for fixed i such that* $0 \le i \le h$*, that any node in* $G_i$ *has degree greater than* $\Delta_i \stackrel{def}{=} 6\log^{b_i} n$ *is at most* $(m_i+n_i)2^{-\log^{b_i} n} < 2^{-\log^{b_i-1} n}$.

*Proof.* The expected degree of any hole at stage $i$ will be $m_i d/n$. The expected degree of any pigeon will be $n_i d/n$. By the conditions on $m$ and $n$, the former quantity is at most twice the latter, which is equal to $\log^{b_i} n$. Fix a node in the graph and let $X$ be the random variable that represents its degree. This is a sum of Bernoulli trials since the edges occur independently with the same probability. Chernoff's bound tells us that $\mathbf{Pr}(X > 3\mu(X)) < (e^2/27)^{\mu(X)} < (1/2)^{\mu(X)}$. We know that $\log^{b_i} n \le \mu(X) \le 2\log^{b_i} n$ and $b_i \ge 2$, so we have the bound. $\square$

**Lemma 4.** *Suppose that n is sufficiently large and let* $0 \le i \le h$*. The probability that any node in* $G_i$ *has degree less than* $\frac{1}{2}\log^{b_i} n$ *is at most* $(m_i+n_i)2^{-\frac{1}{16}\log^{b_i} n} < 2^{-\log^{b_i-1} n}$.

*Proof.* The expected degree of any particular node in $G_i$ is at least $\log^{b_i} n$. Applying a Chernoff bound in the form $\mathbf{Pr}(X < \frac{1}{2}\mu(X)) < \exp(-\frac{1}{8}\mu(X))$, we have the result. $\square$

We now develop the switching lemma argument. The overall structure uses the simplified counting techniques of [16] and [4], however the statement and proof are both complicated by the need to use probabilistic properties of the formulas themselves as well as the relationship of those properties to the restrictions under consideration. In order to discuss these properties we need some definitions:

DEFINITION 5.2. For a graph $H$ and integers $\ell$ and $\Delta$, let $N^{\ell,\Delta}(H)$ be the set of all $\rho$ in $M^\ell(H)$ such that all nodes of $H|_\rho$ have degree at most $\Delta$. Let $N^{\ell,\Delta}_{+j}(H)$ be the set of partial matchings of size $\ell + j$ on $H$ that are extensions of matchings in $N^{\ell,\Delta}(H)$.

For a particular $i$, the set $N^{\ell_i,\Delta_{i+1}}(G_i)$ represents in some sense the usable or "good" portion of all the matchings in $M^{\ell_i}(G_i)$. The next lemma shows that with high probability most of the matchings at any particular stage $i$ are good:

**Lemma 5.** *Let* $0 \le i < h$*. Then*
$\mathbf{Pr}\left[\frac{|N^{\ell_i,\Delta_{i+1}}(G_i)|}{|M^{\ell_i}(G_i)|} < 1 - 2^{-\log^{b_{i+1}-2} n}\right] < 1/n.$

*Proof.* Observe that $\mathbf{E}\left[\frac{|N^{\ell_i,\Delta_{i+1}}(G_i)|}{|M^{\ell_i}(G_i)|}\right]$ equals $\mathbf{Pr}[\rho_i \in N^{\ell_i,\Delta_{i+1}}(G_i)] > 1 - 2^{-\log^{b_{i+1}-1} n}$ by Lemma 3. Now, since $\frac{|N^{\ell_i,\Delta_{i+1}}(G_i)|}{|M^{\ell_i}(G_i)|}$ is bounded above by 1, we can apply Markov's inequality to yield that $\mathbf{Pr}[\frac{|N^{\ell_i,\Delta_{i+1}}(G_i)|}{|M^{\ell_i}(G_i)|} \le 1 - n \cdot 2^{-\log^{b_{i+1}-1} n}] < 1/n$. The result follows by observing that $n \cdot 2^{-\log^{b_{i+1}-1} n} < 2^{-\log^{b_{i+1}-2} n}$. $\square$

We are now read to state the switching lemma.

**Lemma 6 (Switching Lemma).** *Consider $G$ and $\rho_0, \ldots, \rho_{h-1}$ chosen according to the experiment defined above. Let $i, s, r$ be any integers such that $0 \le i < h$, $0 < s \le \Delta_{i+1}/\log^3 n$ and $r > 0$. Finally, let $F$ be any matching disjunction with conjunctions of size $\le r$ over the edge-variables of $G_i$. The probability that $Tree_{G_{i+1}}(F|_{\rho_i})$ has height $\ge s$ conditioned on the events*

- *$G_i$ has maximum degree at most $\Delta_i$*
- $\frac{|N^{\ell_i, \Delta_{i+1}}(G_i)|}{|M^{\ell_i}(G_i)|} \ge 1 - 2^{-\log^{b_{i+1}-2} n}$
- $\rho_i \in N^{\ell_i, \Delta_{i+1}}(G_i)$

*is at most* $2\left(720r/\log^{b_i/2} n\right)^{s/2}$.

DEFINITION 5.3. Let $stars(r, j)$ be the set of all sequences $\beta = (\beta_1, \ldots, \beta_k)$ such that for each $i$, $\beta_i \in \{*, -\}^r \setminus \{-\}^r$ and the total number of $*$'s in $\beta$ is $j$.

**Lemma 7 ([4]).** $|stars(r, j)| < (r/\ln 2)^j$.

**Lemma 8.** *For $H$ a fixed bipartite graph with an ordering on its nodes, let $F$ be a matching disjunction with conjunctions of size $\le r$ over the edge-variables of $H$ and let $S \subset N^{\ell,\Delta}(H)$ be the set of matchings $\rho$ such that $Tree_{H|_\rho}(F|_\rho)$ has height $\ge s$. There is an injection from the set $S$ to the set* $\bigcup_{s/2 \le j \le s} N^{\ell,\Delta}_{+j}(H) \times stars(r, j) \times [\Delta]^s$.

*Proof.* Let $F = C_1 \vee C_2 \vee \ldots$. If $\rho \in S$, then let $\pi$ be the partial matching labelling the first path in $Tree_{H|_\rho}(F|_\rho)$ of length $\ge s$ (actually, we consider only the first $s$ edges in $\pi$, starting from the root, and hence we assume $|\pi| = s$). Let $C_{\nu_1}$ be the first term in $F$ not set to 0 by $\rho$ and let $K_1$ be the variables of $C_{\nu_1}$ not set by $\rho$. Let $\sigma_1$ be the unique partial matching over $K_1$ that satisfies $C_{\nu_1}|_\rho$ and let $\pi_1$ be the portion of $\pi$ that touches $K_1$.

Now define $\beta_1 \in \{*, -\}^{|K_1|} \setminus \{-\}^{|K_1|}$, so that the $p$-th component of $\beta_1$ is a $*$ if and only if the $p$-th variable in $C_{\nu_1}$ is set by $\sigma_1$.

Continue this process to define $\pi_i$, $\sigma_i$, $K_i$, etc. (replacing $\rho$ with $\rho\pi_1 \ldots \pi_{i-1}$ and $\pi$ with $\pi \setminus \pi_1 \ldots \pi_{i-1}$ until some stage $k$ when we've exhausted all of $\pi$. Let $\sigma$ be the matching $\sigma_1 \ldots \sigma_k$, and $\beta$ be the vector $(\beta_1, \ldots, \beta_k)$. Let $j = |\sigma|$ be the number of edges in $\sigma$. Note that $s/2 \le j \le s$. Observe that $\beta \in stars(r, s)$ and $\rho\sigma \in N^{\ell,\Delta}_{+j}(H)$.

We now encode the differences between all the corresponding $\pi_i$ and $\sigma_i$ pairs in a single vector $\delta$ consisting of $|\pi| = s$ components, each in $\{1, \ldots, \Delta\}$. Let $u_1$ be the smallest numbered node in $K_1$ and suppose that $\pi$ (in particular $\pi_1$) matches $u_1$ with some node $v_1$. Then the first component of $\delta$ is the natural number $x$ such that $v_1$ is the $x$-th neighbor (under the ordering of nodes) of $u_1$ in the graph $H|_{\rho\sigma_2\sigma_3\ldots\sigma_k}$. More generally, until the mates of all nodes in $K_1$ under $\pi_1$ have been determined, we determine the $p$-th component of $\delta$ by finding the smallest numbered node $u_p$ of $K_1 \setminus \{u_1, \ldots, u_{p-1}, v_1, \ldots, v_{p-1}\}$ and then we find its mate $v_p$ under $\pi_1$ and encode the position $x$ of $v_p$ in the order of the neighbors of $u_p$ in $H|_{\rho\sigma_2\sigma_3\ldots\sigma_k}$. Once $K_1$ (and thus $\pi_1$) has been exhausted the next component is based on the mates of the smallest numbered nodes in $K_2$ under $\pi_2$, until that is exhausted, etc. where the ordering about each vertex when dealing with $K_i$ is with respect to the graph $H|_{\rho\sigma_{i+1}\sigma_{i+2}\ldots\sigma_k}$.

Finally, we define the image of $\rho \in S$ under the injection to be $(\rho\sigma, \beta, \delta)$. To prove that this is indeed an injection, we show how to invert it: Given $\rho\sigma_1 \ldots \sigma_k$, we can identify $\nu_1$ as the index of the first term of $F$ that is not set to 0 by it. Then, using $\beta_1$ we can reconstruct $\sigma_1$ and $K_1$. Next, reading the components of $\delta$ and the graph $H|_{\rho\sigma_2\ldots\sigma_k}$, until all of $K_1$ is matched, we can reconstruct $\pi_1$. Then we can derive $\rho\pi_1\sigma_2 \ldots \sigma_k$.

At a general stage $i$ of the inversion, we will know $\pi_1, \ldots, \pi_{i-1}$ and $\sigma_1, \ldots, \sigma_{i-1}$ and $K_1, \ldots, K_{i-1}$. We use $\rho\pi_1 \ldots \pi_{i-1}\sigma_i \ldots \sigma_k$ to identify $\nu_i$ and, hence, $\sigma_i$ and $K_i$ (using $\beta$). Then we get $\pi_i$ from $\delta$, $K_i$, and $\rho\sigma_{i+1} \ldots \sigma_k$. After $k$ stages, we will know all of $\sigma$ and can recover $\rho$. $\square$

*Proof of Lemma 6.* We prove this by obtaining an upper bound on $\frac{|N^{\ell_i, \Delta_{i+1}}_{+j}(G_i)|}{|N^{\ell_i, \Delta_{i+1}}(G_i)|}$ and then applying Lemma 8.

By definition, $N^{\ell_i, \Delta_{i+1}}_{+j}(G_i) \subseteq M^{\ell_i + j}(G_i)$ so we develop an upper bound on $|M^{\ell_i+j}(G_i)|$. We can count the number of matchings in $M^{\ell_i+j}(G_i)$ by choosing matchings in $M^{\ell_i}(G_i)$ and extending them by $j$ edges. A matching $\rho \in M^{\ell_i}(G_i)$ can be extended by $j$ edges by choosing $j$ unmatched holes and then choosing the sequence of their $j$ neighbors in $G_i|_\rho$. For each such $\rho$, there are $\binom{n_i - \ell_i}{j}$ ways to choose the unmatched holes. Then, for $\rho \in N^{\ell_i, \Delta_{i+1}}(G_i)$, there are at most $\Delta_{i+1}$ ways to choose a neighbor for each of these holes since that is the degree of $G_i|_\rho$. If $\rho \in M^{\ell_i}(G_i) \setminus N^{\ell_i, \Delta_{i+1}}(G_i)$, then the degree of $G_i|_\rho$ is at most the degree of $G_i$, which is $\Delta_i$ by assumption, so there are at most $\Delta_i$ ways to choose a neighbor for each of the unmatched holes. This procedure overcounts each matching of size $\ell_i + j$ by a factor of $\binom{\ell_i + j}{j}$. Hence we have that $|M^{\ell_i+j}(G_i)|$ is at most

$$\left(|N^{\ell_i, \Delta_{i+1}}(G_i)| \cdot \Delta_{i+1}^j + |M^{\ell_i}(G_i) \setminus N^{\ell_i, \Delta_{i+1}}(G_i)| \cdot \Delta_i^j\right) \frac{\binom{n_i - \ell_i}{j}}{\binom{\ell_i + j}{j}}.$$

Now by assumption

$$|M^{\ell_i}(G_i) \setminus N^{\ell_i, \Delta_{i+1}}(G_i)| \le \frac{2^{-\log^{b_{i+1}-2} n}|N^{\ell_i, \Delta_{i+1}}(G_i)|}{1 - 2^{-\log^{b_{i+1}-2} n}}$$

$$\leq \quad 2^{1-\log^{b_{i+1}-2} n}|N^{\ell_i,\Delta_{i+1}}(G_i)|.$$

Thus, using the fact that $|N_{+j}^{\ell_i,\Delta_{i+1}}(G_i)| \leq |M^{\ell_i+j}(G_i)|$, we have

$$\begin{aligned} &|N_{+j}^{\ell_i,\Delta_{i+1}}(G_i)| \\ &\leq \quad |N^{\ell_i,\Delta_{i+1}}(G_i)| \left(\Delta_{i+1}^j + 2^{1-\log^{b_{i+1}-2} n}\Delta_i^j\right) \frac{\binom{n_i-\ell_i}{j}}{\binom{\ell_i+j}{j}}. \end{aligned}$$

Then, $\dfrac{|N_{+j}^{\ell_i,\Delta_{i+1}}(G_i)|}{|N^{\ell_i,\Delta_{i+1}}(G_i)|}$ is at most

$$\begin{aligned} &\left[1 + 2^{1-\log^{b_{i+1}-2} n}\left(\frac{\Delta_i}{\Delta_{i+1}}\right)^j\right]\left(\frac{\Delta_{i+1}(n_i-\ell_i)}{\ell_i}\right)^j \\ &< \left[1+2^{1-\Delta_{i+1}/(6\log^2 n)}(\log n)^{k_i s}\right]\left(\frac{\Delta_{i+1} n_i}{\ell_i \log^{k_i} n}\right)^j \\ &< \left[1+2^{1-\Delta_{i+1}/(6\log^2 n)}(\log n)^{k_i \Delta_{i+1}/\log^3 n}\right]\left(\frac{\Delta_{i+1} n_i}{\ell_i \log^{k_i} n}\right)^j \\ &< \left(\frac{2\Delta_{i+1}}{\log^{k_i} n}\right)^j = \left(\frac{12\log^{b_{i+1}} n}{\log^{k_i} n}\right)^j. \end{aligned}$$

The first two inequalities and follow from $j \leq s \leq \Delta_{i+1}/\log^3 n$ and the definitions of $\Delta_i$ and $\Delta_{i+1}$. The third inequality follows since $12k_i \log\log n < \log n$ for $n$ sufficiently large and the fact that $n_i/\ell_i = 1/(1-1/\log^{k_i} n)$ which is close to 1.

The probability we wish to bound is the size of the image in Lemma 8 with $H \leftarrow G_i$, $\ell \leftarrow \ell_i$, and $\Delta \leftarrow \Delta_{i+1}$ divided by $|N^{\ell_i,\Delta_{i+1}}(G_i)|$. By the above bound, this is less than

$$\sum_{s/2 \leq j \leq s} \left(\frac{20r\log^{b_{i+1}} n}{\log^{k_i} n}\right)^j \times (6\log^{b_{i+1}} n)^s.$$

Since $b_{i+1} = b_i - k_i$ and without loss of generality $\frac{20r\log^{b_i} n}{\log^{2k_i} n} < 1/2$ (otherwise the probability bound in the lemma statement is meaningless), this quantity is at most

$$2(720r\log^{3b_i-4k_i} n)^{s/2} \leq 2(720r/\log^{b_i/2} n)^{s/2}$$

since $3b_i - 4k_i = -b_i/2$. □

DEFINITION 5.4. Let $S_G$ be a set of formulas of depth at most $h$ that is closed under subformulas and defined over the graph $G$. For $\rho = \rho_0 \ldots \rho_{h-1}$, we define, for every $0 \leq i < h$, $\mathcal{T}_{\rho_0\ldots\rho_i}$, a mapping from formulas with depth $\leq i+1$ in $S_G$ to matching decision trees. It is defined inductively as follows:

For a variable $X_e$, $\mathcal{T}_{\rho_0}(X_e)$ is $Tree_G(X_e)|_{\rho_0}$. Furthermore, $\mathcal{T}_{\rho_0}(\neg X_e)$ is $(\mathcal{T}_{\rho_0}(X_e))^c$. For $A$ a depth-1 formula with merged form $\bigvee_{e\in I} X_e$, $\mathcal{T}_{\rho_0}(A)$ is $Tree_{G_1}((\bigvee_{e\in I} X_e)|_{\rho_0})$.

For $0 < i < h$, for all formulas $A$ of depth $< i+1$, $\mathcal{T}_{\rho_0\ldots\rho_i}(A)$ is $\mathcal{T}_{\rho_0\ldots\rho_{i-1}}(A)|_{\rho_i}$. For a formula $A$ of depth $i+1$, if $A = \neg B$, then $\mathcal{T}_{\rho_0\ldots\rho_i}(A)$ is $(\mathcal{T}_{\rho_0\ldots\rho_i}(B))^c$, and otherwise, if the merged form of $A$ is $\bigvee_{j\in J} B_j$, let $F$ be the matching disjunction $\bigvee_{j\in J} disj(\mathcal{T}_{\rho_0\ldots\rho_{i-1}}(B_j))$ and let $\mathcal{T}_{\rho_0\ldots\rho_i}(A)$ be the canonical matching tree $Tree_{G_{i+1}}(F|_{\rho_i})$.

¿From the definition of $\mathcal{T}_\rho$, we have that if $\neg A$ is a formula in $S_G$, then $\mathcal{T}_\rho(\neg A) = (\mathcal{T}_\rho(A))^c$. Also, by lemma 2, if $\bigvee_{i\in I} A_i$ is the merged form of some formula $A$ in $S_G$, then $\mathcal{T}_\rho(A)$ represents $\bigvee_{i\in I} disj(\mathcal{T}_\rho(A_i))$.

**Lemma 9.** *Let $a$ and $h$ be positive integers. For each graph $G$, let $S_G$ be a set of formulas closed under subformulas defined on the variables of $G$ such that $|S_G| \leq 2^{\log^a n}$ and each formula $A \in S_G$ has depth at most $h$. For $n$ sufficiently large in $a$ and $h$, there exists a choice of $G$ and $\rho = \rho_0, \ldots, \rho_{h-1}$ as defined above such that the following conditions hold:*
*1. $\mathcal{T}_\rho(A)$ has height at most $\log^a n$ for all $A \in S_G$, and*
*2. every node in $G_h$ has degree at least $\log^{a+1} n$.*

*Proof.* We proceed using the probabilistic method. For $0 \leq i \leq h$, define the following events:
$A_i$: Every node in $G_i$ has degree at most $\Delta_i = 6\log^{b_i} n$.
$B_i$: Every node in $G_i$ has degree at least $(1/2)\log^{b_i} n$.
$C_i$: $\frac{|N^{\ell_i,\Delta_{i+1}}(G_i)|}{|M^{\ell_i}(G_i)|} \geq 1 - 2^{-\log^{b_{i+1}-2} n}$. Here $i < h$.
$D_i(A)$: $\mathcal{T}_{\rho_0\ldots\rho_{i-1}}(A)$ has height at most $\log^a n$ for some formula $A \in S_G$ of depth at most $i$. Here $i \geq 1$.
$D_i$: for all formulas $A \in S_G$ of depth at most $i$, $D_i(A)$ holds. Here $i \geq 1$.

We compute an upper bound on the probability that any of these events fails to be true and prove that this probability is strictly less than 1. Since $b_h = a+3$, if $B_h$ occurs and $D_i(A)$ occurs for each $i = 1, \ldots, h$ and each $A \in S_G$ of depth $i$ then the claims of the lemma are satisfied for that$(G, \rho)$, so this probability bound suffices.

Now by Lemma 3, $\mathbf{Pr}[\neg A_i] < 2^{-\log^{b_i-1} n}$ and by Lemma 4, $\mathbf{Pr}[\neg B_i] < 2^{-\log^{b_i-1} n}$. Furthermore, by Lemma 5, $\mathbf{Pr}[\neg C_i] \leq 1/n$. Let $A \in S_G$ be of depth $i < h$ with the merged form of $A$ equal to $\bigvee_{j\in J} Q_j$ and let $F$ be the matching disjunction $\bigvee_{j\in J} disj(\mathcal{T}_{\rho_0\ldots\rho_{i-1}}(Q_j))$. Observing that $b_h = b/8^h = (a+3)$, by Lemma 6 applied to $F$ with $r = s = \log^a n \leq \Delta_h/\log^3 n$, we have

$$\begin{aligned} &\mathbf{Pr}[\neg D_{i+1}(A) \mid A_i \wedge D_i \wedge A_{i+1} \wedge C_i] \\ &\quad \leq \quad 2(720/\log^{b_i/2-a} n)^{(\log^a n)/2} \\ &\quad \leq \quad 2(720/\log^{b_{h-1}/2-a} n)^{(\log^a n)/2} \\ &\quad \leq \quad 2(720/\log^{3a+3} n)^{(\log^a n)/2} < 2^{-\log^a n}/n \end{aligned}$$

for $n$ sufficiently large. Therefore, $\mathbf{Pr}[\neg D_{i+1} \mid A_i \wedge D_i \wedge A_{i+1} \wedge C_i] \leq 1/n$ since each $S_G$ contains at most $2^{\log^a n}$ disjunctions of depth $i+1$.

Therefore the total probability that some $A_i$, $B_i$, $C_i$, or $D_i$ fails is at most:

$$\sum_{i=0}^{h}\mathbf{Pr}[\neg A_i]+\sum_{i=0}^{h}\mathbf{Pr}[\neg B_i]+\sum_{i=0}^{h-1}\mathbf{Pr}[\neg C_i]$$
$$+\mathbf{Pr}[\neg D_1 \mid A_0\wedge A_1\wedge C_0]$$
$$+\mathbf{Pr}[\neg D_2 \mid A_1\wedge D_1\wedge A_2\wedge C_1]+\cdots$$
$$+\mathbf{Pr}[\neg D_h \mid A_{h-1}\wedge D_{h-1}\wedge A_h\wedge C_{h-1}].$$

In total there are $4h$ terms in this sum, each of which is at most $1/n$, and thus the whole probability is $< 1$. □

¿From now on, we fix a graph $G$ and a restriction $\rho = \rho_0,\ldots,\rho_{h-1}$ obeying the conditions of Lemma 9 when applied to the sets of formulas in $S_G = cl(\Pi_G)$ where each $\Pi_G$ is a proof of $WPHP(G)$ in a proof system $\mathcal{F}$ whose largest rule has size $f$.

The following three lemmas are adapted from [19].

**Lemma 10.** *Let $C$ be a line in a Frege proof $\Pi$. Let $\mathcal{A}$ be the immediate ancestors of $C$ in the proof (if there are any), so that $\mathcal{A}\vdash C$. Let $\mathcal{B}$ be the subformulas of $\mathcal{A}$ and $C$ mentioned in the application of the rule which derives $C$ from $\mathcal{A}$. Let $\Gamma=\mathcal{A}\cup\mathcal{B}\cup\{C\}$. Note that by our bound on the size of rules in $\mathcal{F}$, $|\Gamma|+1\le f$. Finally, let $\sigma$ be a matching which extends soundly some $\sigma_A\in path(\mathcal{T}_\rho(A))$ for each $A\in\mathcal{A}$, some $\sigma_B\in path(\mathcal{T}_\rho(B))$ for each $B\in\mathcal{B}$, and some $\sigma_C\in path(\mathcal{T}_\rho(C))$. If $disj(\mathcal{T}_\rho(A))|_\sigma\equiv 1$ for all $A\in\mathcal{A}$, then $disj(\mathcal{T}_\rho(C))|_\sigma\equiv 1$.*

*Proof.* First note the following facts, where $\alpha,\beta\in\Gamma$ and $D(\alpha)$ is an abbreviation for $disj(\mathcal{T}_\rho(\alpha))$:

- $D(\alpha)|_\sigma\equiv 0$ or $D(\alpha)|_\sigma\equiv 1$
- If $\neg\alpha\in\Gamma$, then $D(\neg\alpha)|_\sigma\equiv 1$ iff $D(\alpha)|_\sigma\equiv 0$.
- If $(\alpha\vee\beta)\in\Gamma$, then $D(\alpha\vee\beta)|_\sigma\equiv 1$ iff $D(\alpha)|_\sigma\equiv 1$ or $D(\beta)|_\sigma\equiv 1$

Now consider the rule $R$ used to derive $C$ formulated as in the examples from section 3. The application of $R$ substitutes subformulas $A_p, A_q, A_r,\ldots$ in $\Gamma$ for each of the atoms $p,q,r,\ldots$ in $R$ and there is a derived correspondence mapping subformulas $F$ appearing in $R$ to formulas $A_F\in\Gamma$. Define a function $\tau$ on the atoms of $R$ by $\tau(p)=D(A_p)|_\sigma$ for each such atom $p$. By the first property, $\tau$ is a truth assignment to these atoms. Furthermore, by the other two properties, the truth assignment $\tau$ extends to all subformulas $F$ in $R$ so that $\tau(F)=D(A_F)|_\sigma$. Since $R$ is sound, if $\tau$ satisfies all formulas in $\mathcal{A}$ it will satisfy $C$ and thus $D(C)|_\sigma\equiv 1$. □

**Lemma 11.** *Let $\mathcal{F}$ be a Frege system with maximum rule size $f$. Let $n$ be sufficiently large w.r.t. $f$. Let $a,h>0$. For each $G$, assume that $\Pi_G$ is a proof in $\mathcal{F}$ of $WPHP(G)$ of size at most $2^{\log^a n}$ and depth at most $h$. Let $\rho$ and $G$ be as defined in Lemma 9 applied with $S_G=cl(\Pi_G)$. If $C$ is an arbitrary line in proof $\Pi_G$ then all leaves of $\mathcal{T}_\rho(C)$ are labelled by* 1.

*Proof.* We proceed by (complete) induction on the lines in the proof. Assume every leaf of $\mathcal{T}_\rho$ for any line preceding $C$ is labelled 1. Let $\mathcal{A}$, $\mathcal{B}$, $\Gamma$ be as in Lemma 10. For any leaf $l$ of $\mathcal{T}_\rho(C)$, we use Lemma 1 to find $\sigma$ that extends $path(l)$ and extends a matching in each of the sets $path(\mathcal{T}_\rho(A))$ for all $A\in\mathcal{A}$ and $path(\mathcal{T}_\rho(B))$ for all $B\in\mathcal{B}$. This is possible since there are at most $f$ trees to consider and by Lemma 9 the sum of their heights is at most $f\log^a n\le\log^{a+1}n$ which is the degree of $G_h$.

By assumption, $disj(\mathcal{T}_\rho(A))|_\sigma\equiv 1$ for all $A$ in $\mathcal{A}$. Hence, by Lemma 10, $disj(\mathcal{T}_\rho(C))|_\sigma\equiv 1$, so $l$ must be labelled 1. □

**Lemma 12.** *All leaves of $\mathcal{T}_\rho(WPHP(G))$ are labelled by* 0.

*Proof.* It suffices to show that $\mathcal{T}_\rho$ applied to each of the following types of formulas has all leaves labelled 0:

1. $\neg(\neg X_e\vee\neg X_{e'})$ for $e,e'\in E; e=\{i,k\}, e'=\{j,k\}; i,j\in V_1; i\ne j; k\in V_2$.
2. $\neg(\neg X_e\vee\neg X_{e'})$ for $e,e'\in E; e=\{k,i\}, e'=\{k,j\}; i,j\in V_2; i\ne j; k\in V_1$.
3. $\neg\bigvee_{j\in\Gamma(i)}X_{\{i,j\}}$ for $i\in V_1$.
4. $\neg\bigvee_{i\in\Gamma(j)}X_{\{i,j\}}$ for $j\in V_2$.

In fact, we will show that $\mathcal{T}_\rho$ applied to the complement of each of these formulas has all leaves labelled 1.

For a formula of the first type, $T=\mathcal{T}_\rho(\neg X_e\vee\neg X_{e'})$ must represent $disj(\mathcal{T}_\rho(\neg X_e))\vee disj(\mathcal{T}_\rho(\neg X_{e'}))$. If $\rho$ sets the value of either $X_e$ or $X_{e'}$ then it must set one of $\neg X_e$ or $\neg X_{e'}$ to 1 and thus all leaves of $\mathcal{T}_\rho(\neg X_e\vee\neg X_{e'})$ are certainly labelled 1. Otherwise, for $l$ a leaf of $T$, $path(l)$ cannot contain both $e$ and $e'$. Without loss of generality it does not contain $e$. By Lemma 1 applied to graph $G_h$ we can find $\sigma$ that extends $path(l)$ and is an extension of some matching in $\mathcal{T}_\rho(\neg X_e)$. But then $disj(\mathcal{T}_\rho(\neg X_e))|_\sigma\equiv 1$, so $l$ must be labelled 1. The argument is the same for formulas of the second type.

For a formula of the third type, $T=\mathcal{T}_\rho(\bigvee_{j\in\Gamma(i)}X_{\{i,j\}})$ must represent $\bigvee_{j\in\Gamma(i)}disj(\mathcal{T}_\rho(X_{\{i,j\}}))$. Hence, if $\rho$ sets $X_{\{i,j\}}$ to 1 for some $j\in\Gamma(i)$, then all leaves of $T$ are certainly labelled 1. Otherwise, for a leaf $l$ of $T$, if $path(l)$ touches node $i$, then $\bigvee_{j\in\Gamma(i)}disj(\mathcal{T}_\rho(X_{\{i,j\}}))|_{path(l)}\equiv 1$. Finally, if $path(l)$ does not touch node $i$, extend it to $\sigma=path(l)\cup\{i,j\}$ for some $j$ such that $X_{\{i,j\}}$ is not set by $\rho$. Then $disj(\mathcal{T}_\rho(X_{\{i,j\}}))|_\sigma\equiv 1$, so $l$ is labelled 1. Formulas of the fourth type follow in the same way. □

**Theorem 13.** *Given any $c$ sufficiently large, there exists a bipartite graph $G$ from $m = n + n/\log^c n$ pigeons to $n$ holes such that there is no depth-$h$, $2^{\log^a n}$-size $\mathcal{F}$-proof of $WPHP(G)$ provided that $8^h(a+3) < c$.*

*Proof.* Assume that for all such $G$, there is a proof $\Pi_G$ of the required depth and size. For the $G$ in Lemma 9 and its corresponding proof $\Pi_G$ of $WPHP(G)$, there exists a $\rho$ such that $\mathcal{T}_\rho(A)$ has all leaves labelled 1 for any $A \in cl(\Pi_G)$, but $\mathcal{T}_\rho(WPHP(G))$ has all leaves labelled 0. If $\Pi_G$ is to be a proof of $WPHP(G)$, then $WPHP(G)$ must appear in $\Pi_G$, so we have a contradiction. □

**Corollary 14.** *Given any $c$ sufficiently large, there is no depth-$h$, $2^{\log^a n}$-size $\mathcal{F}$-proof of $WPHP = WPHP(K_{m,n})$ from $m = n + n/\log^c n$ pigeons to $n$ holes, provided that $8^h(a+3) < c$.*

# 6 Open questions

The most important open problem related to our work is to resolve the complexity of the weak pigeonhole principle with $2n$ or more pigeons, and $n$ holes. As mentioned in the introduction, this has important implications for: the metamathematics of the $P$ versus $NP$ statement, the complexity of approximate counting, and the proof-theoretic strength underlying elementary number theory. In the proof presented here, we apply *direct* restrictions. That is, we limit the space of all truth assignments to a subcube where certain variables are set to 0 and other variables are set to 1. While this type of switching lemma fails with $2n$ pigeons, it may still be possible to prove a switching lemma on a more general class of restrictions. A simple extension would be to try the projections suggested in [20] which allow the identification of variables as well as the assignment of values. More complex extensions, for example a class of restrictions that limits the space of truth assignments to those satisfying a set of linear equations, may resolve the $2n$ to $n$ case.

Another possible avenue for resolving the $2n$ to $n$ case is to explore more general reductions within the bounded-depth Frege setting. In turn, lower bounds for certain versions of the weak pigeonhole principle could then be used to prove, via reductions, lower bounds for random formulas. Our result essentially shows that $WPHP(G)$ is hard with reasonably high probability when $G$ is from a certain class of bipartite graphs. A similar result for a class of graphs with more pigeons and smaller degree would yield lower bounds for random CNFs.

A potentially simpler problem that still gets to the heart of the matter is to prove quasipolynomial lower bounds for $Res(\text{polylog } n)$ proofs of the weak pigeonhole principle which would match the upper bounds in [12]. It is conceivable that this could be achieved by proving lower bounds for $Res(k)$ proofs of the weak pigeonhole principle for larger and larger $k$, extending the exponential lower bound for $Res(2)$ in [2]; but new techniques seem to be needed.

See [17] for an excellent summary of what is known and additional future directions relating to pigeonhole principle complexity. Most importantly, we seek alternative methods for proving lower bounds for Frege systems, as well as a deeper understanding of reductions between various combinatorial principles.

# 7 Acknowledgments

We are very grateful to Alan Woods for suggesting this problem a while back and for the many valuable comments and insights that he shared with us. We also want to thank Mike Molloy for some helpful discussions. Finally, the third author is grateful to Miki Ajtai for hinting that there may just not be a polynomial-size bounded-depth proof.

# References

[1] M. Ajtai. The complexity of the pigeonhole principle. *Combinatorica*, 14(4):417–433, 1994.

[2] A. Atserias, M. L. Bonet, and J. Esteban. Lower bounds for the weak pigeonhole principle beyond resolution. In J. Orejas, P. G. Spirakis, and J. van Leeuwen, editors, *Automata, Languages, and Programming: 28th International Colloquium*, volume 2076 of *Lecture Notes in Computer Science*, pages 1005–1016, Heraklion, Crete, July 2001. Springer-Verlag.

[3] Paul Beame and Søren Riis. More on the relative strength of counting principles. In Paul W. Beame and Samuel 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.

[4] Paul W. Beame. A switching lemma primer. Technical Report UW-CSE-95-07-01, Department of Computer Science and Engineering, University of Washington, November 1994.

[5] Paul W. Beame, Russell Impagliazzo, Jan Krajíček, Toniann Pitassi, Pavel Pudlák, and Alan Woods. Exponential lower bounds for the pigeonhole principle. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing*, pages 200–220, Victoria, B.C., Canada, May 1992.

[6] E. Ben-Sasson and A. Wigderson. Short proofs are narrow – resolution made simple. In *Proceedings of the*

*Thirty-First Annual ACM Symposium on Theory of Computing*, pages 517–526, Atlanta, GA, May 1999.

[7] S. Buss and G. Turán. Resolution proofs of generalized pigeonhole principles. *Theoretical Computer Science*, 62:311–317, 1988.

[8] A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–305, 1985.

[9] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, pages 6–20, Berkeley, CA, May 1986.

[10] J. Krajíček. *Bounded Arithmetic, Propositional Logic and Complexity Theory*. Cambridge University Press, 1996.

[11] *Logic and Algorithmic,* An International Symposium Held in Honor of Ernst Specker, Zürich, February 5–11, 1980. Monographie No. 30 de L'Enseignement Mathématique, Université de Genève, 1982.

[12] A. Maciel, T. Pitassi, and A. Woods. A new proof of the weak pigeonhole principle. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, pages 368–377, Portland,OR, May 2000.

[13] J. Paris and A. 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 *Lecture notes in Mathematics*, pages 317–340, Berlin, 1985. Springer-Verlag.

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

[15] R. Raz. Resolution lower bounds for the weak pigeonhole principle. Technical Report TR01-021, Electronic Colloquium in Computation Complexity, `http://www.eccc.uni-trier.de/eccc/`, 2001.

[16] A. A. Razborov. Bounded arithmetic and lower bounds in Boolean complexity. In P. Clote and J. Remmel, editors, *Feasible Mathematics II*, pages 344–386. Birkhauser, 1995.

[17] A. A. Razborov. Proof complexity of pigeonhole principles. In *Proceedings of the Fifth International Conference on Developments in Language Theory*, Vienna, Austria, July 2001.

[18] A. A. Razborov. Resolution lower bounds for perfect matching principles. In *ProceedingsSeventeenth Annual IEEE Conference on Computational Complexity*, Montreal, PQ, Canada, May 2002. To appear.

[19] Alasdair Urquhart and Xudong Fu. Simplified lower bounds for propositional proofs. *Notre Dame Journal of Formal Logic*, 37(4):523–544, 1996.

[20] Leslie G. Valiant. Reducibility by algebraic projections. *L'Enseignement Mathématique*, XXVIII:253–268, 1982. Also in [11, pages 365–380].