Quasipolynomial Size Proofs of the Propositional Pigeonhole Principle

Sam Buss*
Department of Mathematics
University of California, San Diego
La Jolla, CA 92093-0112, USA
sbuss@ucsd.edu

January 24, 2015

Abstract

Cook and Reckhow proved in 1979 that the propositional pigeonhole principle has polynomial size extended Frege proofs. Buss proved in 1987 that it also has polynomial size Frege proofs; these Frege proofs used a completely different proof method based on counting. This paper shows that the original Cook and Reckhow extended Frege proofs can be formulated as quasipolynomial size Frege proofs. The key point is that st-connectivity can be used to define the Cook-Reckhow construction.

1 Introduction

One of the central questions in proof complexity is whether Frege proofs can polynomially simulate extended Frege (e$\mathcal{F}$) proofs. Frege proofs are the usual "textbook" propositional proof system with modus ponens as the only rule of inference. Extended Frege systems are Frege systems augmented with an extension rule that allows introducing new variables which abbreviate more complex formulas. Frege proofs can be viewed as proofs that reason about polynomial size Boolean formulas, and extended Frege proofs as proofs that reason about polynomial size Boolean circuits. For this reason, questions about the complexity of Frege and extended Frege proofs are

*Supported in part by NSF grants DMS-1101228 and CCF-1213151 and by the Simons Foundation award 306202.

generally thought to be closely related to questions about Boolean computational complexity.

The propositional pigeonhole principle (PHP) has played a central role as an example in proof complexity. Cook and Reckhow [8] showed that the "$n+1$ into n" versions of the pigeonhole principle, PHP_n^{n+1}, have polynomial size extended Frege proofs. These extended Frege proofs work by formalizing a proof by induction on n, reducing PHP_n^{n+1} to an instance of PHP_{n-1}^{n}. The corresponding Frege proofs, obtained by unwinding the uses of the extension rule, are exponential size: thus this left open the question of whether PHP_n^{n+1} requires exponential size Frege proofs.

Buss [5] subsequently gave polynomial size Frege proofs of PHP_n^{n+1}. These Frege proofs used a very different proof method based on the formalization of counting with polynomial size formulas. This left still open the question of whether the Cook-Reckhow extended Frege proofs could be naturally translated into sub-exponential size Frege proofs. Indeed, the apparent difficulty of finding such a translation has been taken as evidence that there may be an exponential speedup of extended Frege proofs over Frege proofs.

Theorem 1 of the present paper, however, shows there are quasipolynomial size Frege proofs of PHP_n^{n+1} which are essentially direct translations of Cook and Reckhow's extended Frege proofs. We give two different constructions of quasipolynomial size Frege proofs. Both are based on st-connectivity. The first one uses the weak n^2 to n pigeonhole principle, $\mathrm{PHP}_n^{n^2}$, which is well-known to have constant depth, polynomial size Frege proofs [15]. In essence, this proof is giving a quasipolynomial reduction from the pigeonhole principle to PPADS (the "sink" version of the parity principle for directed graphs, see Papadimitriou [14]), and then applies the $\mathrm{PHP}_n^{n^2}$ principle to the PPADS problem. It may seem a bit of cheat to reduce the PHP_n^{n+1} principle to the $\mathrm{PHP}_n^{n^2}$: this is partially true, but the existence of constant depth, polynomial size Frege proofs for $\mathrm{PHP}_n^{n^2}$ was known before the proof [5] that PHP_n^{n+1} has polynomial size Frege proofs. At any rate, our second construction of quasipolynomial size Frege proofs avoids this "cheat" since it uses neither $\mathrm{PHP}_n^{n^2}$ nor counting.

Bonet, Buss and Pitassi [4] investigated candidates for tautologies that might provide exponential separations for Frege and extended Frege systems. They did not find very many other than tautologies which are complete for extended Frege proofs such as partial consistency statements. A number of these were based on linear algebra, including the Oddtown Theorem, the Graham-Pollack Theorem, the Fisher Inequality and the Ray-

Chaudhuri-Wilson Theorem; another linear-algebra-based tautology stating that $AB = I \Rightarrow BA = I$ for Boolean matrices was subsequently proposed by Cook. These all have polynomial size extended Frege proofs using simple facts from linear algebra. Recently, Hrubeš and Tzameret [9] showed that many identities about the determinant, including the $AB = I \Rightarrow BA = I$ tautologies, have quasipolynomial size Frege proofs (as was already conjectured by [4]).

Several combinatorial principles have been suggested as possibilities for separating Frege and extended Frege proofs; these include Frankl's theorem, local improvement principles, and the Kneser-Lovász coloring principle. Bonet, Buss, and Pitassi [4] suggested Frankl's theorem on the trace of sets, and showed these tautologies have polynomial size extended Frege proofs. But, subsequently to the results of the present paper, Aisenberg, Bonet, and Buss [1] gave quasipolynomial size Frege proofs of the tautologies expressing Frankl's theorem. They also showed that Frankl's Theorem with constant parameter t has polynomial size Frege proofs; this was earlier shown for $t = 1, 2$ by [4, 13]. Kołodziejczyk, Nguyen, and Thapen [11] suggested the propositional translations of various local improvement principles LI, $\text{LI}_{\log}$ and LLI as candidates, motivated by results on their provability in the bounded arithmetic theory V^1_2. They proved the LI principle is equivalent to partial consistency statements for extended Frege systems, but the other two remained as candidates. However, Beckmann and Buss [3] showed that the LLI principles are provable in the bounded arithmetic theory U^1_2; thus they also have quasipolynomial size Frege proofs. They also showed $\text{LI}_{\log}$ to be equivalent to LI. Finally, Aisenberg, Bonet, Buss, Crãciun, and Istrate [in preparation] have given quasipolynomial size Frege proofs for the Kneser-Lovász coloring principle; Istrate and Crãciun [10] earlier gave polynomial size extended Frege proofs for a special case of these tautologies.

As mentioned, our new Frege proofs constructed for Theorem 1 use the same underlying construction as Cook and Reckhow's extended Frege proofs. Likewise, all of the above-mentioned quasipolynomial size Frege proofs use the same underlying constructions as the prior extended Frege proofs; although the details become substantially more difficult in most cases (including in the case of the Frege proofs constructed in the present paper). This raises the possibility that Frege systems actually do quasipolynomially simulate extended Frege systems. This seems quite unlikely, however, as the techniques do not seem to apply to general extended Frege proofs.

We presume the reader has basic familiarity with Frege and extended Frege proofs, and with the kinds of arguments that can be formalized with Frege proofs. The reader unfamiliar with proof complexity should consult

the papers cited above. In addition, [6, 7, 2, 16] give surveys of propositional proof complexity.

A Frege proof system is an implicationally sound and complete propositional proof system, axiomatized by a finite set of schematic axioms and inference rules [8]. An example of a schematic axiom is the set of formulas $A \to (B \to A)$. An example of a schematic inference rule is *modus ponens*: from A and $A \to B$, infer B. We assume w.l.o.g. that Frege systems use the connectives $\neg$, $\wedge$, $\vee$, $\to$, $\leftrightarrow$, $\top$, and $\bot$. The length of a Frege proof P, denoted $|P|$, is the number of symbols occurring in P.

Let $n \geq 1$. We write $[n]$ for $\{0, 1, \ldots, n-1\}$. We use propositional variables $p_{i,j}$ to denote the condition that $f(i) = j$. For $i \in [n+1]$, define

$$\mathrm{Tot}_i^n \;:=\; \bigvee_{j \in [n]} p_{i,j}$$

stating that $f(i)$ is defined. For $j \in [n]$, define

$$\mathrm{Inj}_j^n \;:=\; \bigwedge_{0 \leq i_1 < i_2 \leq n} \neg(p_{i_1,j} \wedge p_{i_2,j})$$

stating that f is injective at j. Then PHP_n^{n+1} is the formula

$$\neg\Big(\bigwedge_{i \in [n+1]} \mathrm{Tot}_i^n \wedge \bigwedge_{j \in [n]} \mathrm{Inj}_j^n \Big)$$

stating that there is no total injective $f : [n+1] \to [n]$.

Big conjunctions ($\bigwedge$) and disjunctions ($\bigvee$) always denote (nearly) balanced trees of two input $\wedge$- and $\vee$-gates, respectively. Thus, these formulas have logarithmic depth.

A function $s : \mathbb{N} \to \mathbb{N}$ is *quasipolynomial* if $s(n) = 2^{\log^{O(1)}(n)}$.

2 Quasipolynomial Frege proofs of PHP_n^{n+1}

Theorem 1. *There are quasipolynomial size proofs of the tautologies* PHP_n^{n+1}.

Of course, Theorem 1 is weaker than what is already known, namely PHP_n^{n+1} has polynomial size proofs [5]. As discussed already, the point is that the quasipolynomial size Frege proofs are based on the same construction as Cook and Reckhow's extended Frege (e$\mathcal{F}$) proofs.

The rest of the paper is dedicated to sketching the new proof of Theorem 1.

2.1 Polynomial size extended Frege proofs.

We recall Cook and Reckhow's e$\mathcal{F}$ proofs of PHP_n^{n+1}. These e$\mathcal{F}$ proofs start with the assumptions that Tot_i^n holds for all $i \in [n+1]$ and that Inj_j^n holds for all $j \in [n]$, and obtain a contradiction. The e$\mathcal{F}$ proof of PHP_n^{n+1} introduces new variables $q_{i,j}^k$ by the extension rule, for $k = 1, \ldots, n$ and $i \in [k+1]$ and $j \in [k]$. It then proves, for each value of k, that the variables $q_{i,j}^k$ satisfy the Tot_i^k formulas and the Inj_j^k formulas. Once this is established for $k = 1$, a contradiction is readily obtained, since Tot_0^1 is $p_{0,0}^1$, Tot_1^1 is $p_{1,0}^1$, and Inj_0^1 is $\neg(p_{0,0}^1 \vee p_{1,0}^1)$.

The variables $q_{i,j}^k$ are defined as follows. First,

$$q_{i,j}^n \;\leftrightarrow\; p_{i,j}.$$

(Alternately, $q_{i,j}^n$ is just another name for $p_{i,j}$.) Then, successively for $k = n-1, \ldots, 2, 1$, define $q_{i,j}^k$ by

$$q_{i,j}^k \;\leftrightarrow\; q_{i,j}^{k+1} \vee (q_{i,k}^{k+1} \wedge q_{k+1,j}^{k+1}) \tag{1}$$

for $i \in [k+1]$ and $j \in [k]$. The intuitive idea for the definitions (1) of $q_{i,j}^k$ is shown in Figure 1.

It is clear by inspection that, under the assumption that the variables $\vec{q}^{k+1}$ violate the PHP_{k+1}^{k+2} pigeonhole principle, the new variables $\vec{q}^k$ violate the PHP_k^{k+1} pigeonhole principle. More formally, write $\mathrm{Tot}_{i,j}^k(\vec{q}^k)$ for the result of substituting the variables $q_{i,j}^k$ for the variables $p_{i,j}$ in $\mathrm{Tot}_{i,j}^k$. Then given the hypotheses (1), it is straightforward to prove

$$\mathrm{Tot}_i^{k+1}(\vec{q}^{k+1}) \wedge \mathrm{Tot}_{k+1}^{k+1}(\vec{q}^{k+1}) \to \mathrm{Tot}_i^k(\vec{q}^k)$$

for each $i \in [k]$, and

$$\mathrm{Inj}_j^{k+1}(\vec{q}^{k+1}) \wedge \mathrm{Inj}_k^{k+1}(\vec{q}^{k+1}) \to \mathrm{Inj}_j^k(\vec{q}^k)$$

for each $j \in [k-1]$. These proofs can be carried out with polynomial size Frege proofs given the equivalences (1) as hypotheses. Therefore, there are polynomial size Frege proofs of $\mathrm{PHP}_{k+1}^{k+2}(\vec{q}^{k+1}) \to \mathrm{PHP}_k^{k+1}(\vec{q}^k)$ given the hypotheses (1). Putting these together gives Cook and Reckhow's proof of PHP_n^{n+1}.

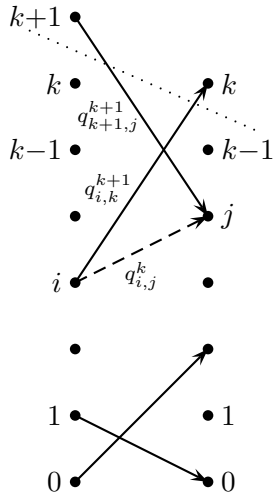

Figure 1: Discarding pigeon $k{+}1$ and hole k, using extension to define $q^k_{i,j}$ to equal $q^{k+1}_{i,j} \vee (q^{k+1}_{i,k} \wedge q^{k+1}_{k+1,j})$.

2.2 Preliminaries for quasipolynomial size Frege proofs.

The problem with directly translating the e$\mathcal{F}$ proof into a Frege proof is that unwinding the definitions (1) into formulas defining $q^k_{i,j}$ yields exponential size formulas in the original variables $p_{i,j}$. We shall give an alternate definition of the values $q^k_{i,j}$ using formulas $\varphi^k_{i,j}$ of *quasipolynomial* size. The trick is to define the variables $\vec{\varphi}^k$ independently for each k, rather than inductively for successive values of k. This will be done by defining suitable directed acyclic graphs (dags) $G_{n,k}$ with out-degree at most 1, and using an st-connectivity property.

Fix a value for n. The $n{+}1$ *pigeon nodes* are the pairs $\langle 0, i\rangle$ for $i \in [n{+}1]$, and the n *hole nodes* are the pairs $\langle 1, j\rangle$, for $j \in [n]$. Now fix $k \leq n$. The nodes of $G_{n,k}$ consist of all $n+1$ pigeon nodes and n hole nodes. The edges of $G_{n,k}$ are the edges as given by the presumed violation of the pigeonhole principle, plus the "back edges" from each hole node $\langle 1, \ell\rangle$ to pigeon node $\langle 0, \ell{+}1\rangle$ for $k \leq \ell < n$. This is pictured in Figure 2. More formally, we use variables α and β to denote nodes of $G_{n,k}$, and define formulas $\gamma^k_{\alpha,\beta}$ indicating which edges are present in $G_{n,k}$. For $\alpha = \langle 0, i\rangle$ and $\beta = \langle 1, j\rangle$ with $i \in [n{+}1]$ and $j \in [n]$, the formula $\gamma^k_{\alpha,\beta}$ is

$$p_{i,j} \ \wedge \ \bigwedge_{j'<j} \neg p_{i,j'}.$$

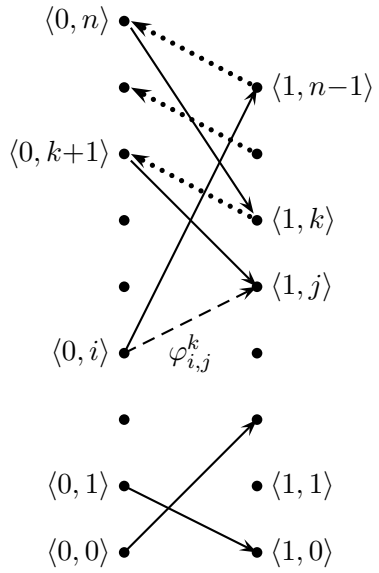

Figure 2: The graph $G_{n,k}$ has as nodes the pigeon nodes $\langle 0, i\rangle$, $i \in [n{+}1]$, and the hole nodes $\langle 1, j\rangle$, $j \in [n]$. The edges of $G_{n,k}$ are the solid lines from left-to-right as indicated by the variables $p_{i,j}$, and the dotted lines from $\langle 1, j\rangle$ to $\langle 0, j{+}1\rangle$ for $k \leq j < n$. The dashed line indicates an edge from $\langle 0, i\rangle$ to $\langle 1, j\rangle$, as defined by $\varphi^k_{i,j}$ for $i \in [k{+}1]$ and $j \in [k]$: this is in the transitive closure of $G_{n,k}$, but not necessarily in $G_{n,k}$.

With the totality hypothesis Tot^n_i, this ensures that for each i, $\gamma^k_{i,j}$ is true for exactly one value j. For $\alpha = \langle 1, j\rangle$ and $\beta = \langle 0, j+1\rangle$ with $k \leq j < n$, the formula $\gamma^k_{\alpha,\beta}$ is the constant true ($\top$). For all other α and β, $\gamma^k_{\alpha,\beta}$ is the constant false ($\bot$).

From $G_{n,k}$ we further define formulas $\varphi^k_{i,j}$ which define an instance of the PHP^{k+1}_k pigeonhole principle. The idea is simple: We trace out maximal length paths in $G_{n,k}$ starting at a pigeon node $\langle 0, i\rangle$ and terminating at a hole node $\langle 1, j\rangle$. The formula $\varphi^k_{i,j}$ is true if and only if this path exists. Then, we prove by "brute-force" induction on $k = n, n{-}1, \ldots, 1$ that the formulas $\vec{\varphi}^k$ falsify PHP^{k+1}_k. At $k = 1$, this yields a contradiction.

The formulas $\varphi^k_{i,j}$ are defined in terms of formulas $\mathrm{Path}^k[\ell, \alpha, \beta]$ which express the property that there is a path of length ℓ from α to β in $G_{n,k}$. Note that there is a different formula $\mathrm{Path}^k[\ell, \alpha, \beta]$ for each choice of values for $n, k, \ell, \alpha, \beta$. (The dependency on n is suppressed in the notation.) For $\ell = 0$, $\mathrm{Path}^k[0, \alpha, \alpha]$ is the constant $\top$, and for $\alpha \neq \beta$, $\mathrm{Path}^k[0, \alpha, \beta]$ is

the constant $\bot$. For $\ell = 1$, $\mathrm{Path}^k[1, \alpha, \beta]$ is the formula $\gamma^k_{\alpha,\beta}$. For $\ell > 1$, $\mathrm{Path}^k[\ell, \alpha, \beta]$ is defined to equal

$$\bigvee_{\alpha' \in G_{n,k}} \Big(\mathrm{Path}^k[\lfloor \tfrac{\ell}{2} \rfloor, \alpha, \alpha'] \wedge \mathrm{Path}^k[\lceil \tfrac{\ell}{2} \rceil, \alpha', \beta] \Big) .$$

Let $N = 2n + 1$. Define $\varphi^k_{i,j}$, for $i \in [k{+}1]$ and $j \in [k]$, to be

$$\bigvee_{\ell \leq N} \mathrm{Path}^k[\ell, \langle 0, i\rangle, \langle 1, j\rangle]. \tag{2}$$

By inspection, $\varphi^k_{i,j}$ has size $2^{O(\log^2 n)}$ and depth $O(\log^2 n)$, i.e., quasipolynomial size and polylogarithmic depth.

The definition of $\varphi^k_{i,j}$ is illustrated in Figure 2, and it not hard to convince oneself that, if the $p_{i,j}$'s define a violation of PHP^{n+1}_n, then the $\varphi^k_{i,j}$'s define a violation of PHP^{k+1}_k. Our second construction of quasipolynomial size Frege proofs will give formal proofs of this fact.

Lemma 2. *The following formulas have quasipolynomial size Frege proofs from the hypothesis $\neg\mathrm{PHP}^{n+1}_n$, for all appropriate k, α, β and β', and all $\ell' \leq \ell$. Quasipolynomial size means quasipolynomial in n and ℓ.*

(a) $\mathrm{Path}^k[\ell, \alpha, \beta] \;\rightarrow\; \bigvee_{\alpha'} \mathrm{Path}^k[\ell', \alpha, \alpha'] \wedge \mathrm{Path}^k[\ell{-}\ell', \alpha', \beta]$.

(b) $\mathrm{Path}^k[\ell', \alpha, \alpha'] \wedge \mathrm{Path}^k[\ell{-}\ell', \alpha', \beta] \;\rightarrow\; \mathrm{Path}^k[\ell, \alpha, \beta]$.

(c) $\mathrm{Path}^k[\ell, \alpha, \beta] \;\rightarrow\; (\mathrm{Path}^k[\ell', \alpha, \alpha'] \leftrightarrow \mathrm{Path}^k[\ell{-}\ell', \alpha', \beta])$.

(d) *For $\beta' \neq \beta$,*

$$\neg\big(\mathrm{Path}^k[\ell, \alpha, \beta] \wedge \mathrm{Path}^k[\ell, \alpha, \beta']\big),$$

and

$$\neg\big(\mathrm{Path}^k[\ell, \beta, \alpha] \wedge \mathrm{Path}^k[\ell, \beta', \alpha]\big).$$

Proof. The proof is a standard "brute-force" induction. Fix a value for k. The statements (a), (b) and (c) for all α and β, and all $\ell' \leq \ell$, are proved simultaneously, first for $\ell = 0$, then for $\ell = 1$, then for $\ell = 2$, etc. The base cases for (c) depend on the fact that every node has in- and out-degrees ≤ 1.

Part (d) follows from the $\ell' = 0$ and $\ell' = \ell$ cases of (c). □

Define $\mathrm{Src}_{n,k}$ to be the set of *source* nodes in $G_{n,k}$, namely the nodes $\langle 0, i\rangle$ for $i \in [k{+}1]$. Likewise, the set $\mathrm{Snk}_{n,k}$ of *sink* nodes of $G_{n,k}$ contains the nodes $\langle 1, j\rangle$ for $j \in [k]$.

Lemma 3. *The following have quasipolynomial size (in n and ℓ) Frege proofs from the hypothesis that $\neg\mathrm{PHP}_n^{n+1}$, for all $k \leq n$, all nodes α and β, and all $\ell' < \ell$.*

(a) *For $\alpha \in \mathrm{Src}_{n,k}$,*

$$\neg\big(\mathrm{Path}^k[\ell', \alpha, \beta] \wedge \mathrm{Path}^k[\ell, \alpha, \beta]\big).$$

(b) *For $\beta \in \mathrm{Snk}_{n,k}$,*

$$\neg\big(\mathrm{Path}^k[\ell', \alpha, \beta] \wedge \mathrm{Path}^k[\ell, \alpha, \beta]\big).$$

Note that the lemma can fail for general α and β since $G_{n,k}$ can have a cycle in the nodes above $\langle 0, k\rangle$ and $\langle 1, k{-}1\rangle$.

Proof. (Sketch) We assume for sake of a contradiction that $\mathrm{Path}^k[\ell', \alpha, \beta]$ and $\mathrm{Path}^k[\ell, \alpha, \beta]$ both hold, and argue as can be formalized with quasipolynomial size Frege proofs.

First suppose $\ell' = 0 < \ell$. Since $\mathrm{Path}^k[0, \alpha, \beta]$ is a hypothesis, if $\alpha \neq \beta$, we immediately obtain a contradiction. For $\alpha = \beta$, we have $\mathrm{Path}^k[\ell, \alpha, \alpha]$ as a hypothesis. By Lemma 2(a), there is an α' such that $\mathrm{Path}^k[1, \alpha', \alpha]$ holds. But this is impossible since $\alpha \in \mathrm{Src}_{n,k}$, so we again obtain a contradiction.

The $\ell' > 0$ cases are argued as follows. By Lemma 2(c), $\mathrm{Path}^k[\ell{-}\ell', \alpha, \alpha]$. Since also $\mathrm{Path}^k[0, \alpha, \alpha]$, this contradicts the second case handled in the previous paragraph.

Part (b) is proved similarly to (a). □

It is interesting to note that the Frege proofs of Lemmas 2 and 3 have only polynomially many formulas; in fact, the sizes of the Frege proofs are polynomially bounded by the size of the tautologies being proved.

2.3 The first quasipolynomial size Frege proofs.

We describe a family of quasipolynomial size Frege proofs for PHP_n^{n+1}, giving a first proof of Theorem 1. We will informally prove PHP_n^{n+1}, using arguments that can be formalized as quasipolynomial size Frege proofs. As before, we start with the assumptions that Tot_i^n holds for all $i \in [n{+}1]$ and that Inj_j^n holds for all $j \in [n]$.

Consider the graph $G_{n,0}$. Referring to Figure 2, every node in $G_{n,0}$ has out-degree 1. In addition, $\alpha_0 = \langle 0, 0\rangle$ has in-degree 0. (Some hole nodes $\langle 1, j\rangle$ may also have in-degree 0 since we did not assume f is onto.) Consider

paths in $G_{n,0}$ starting at α_0. We let PathExists[ℓ] be the formula stating there is a path of length ℓ starting at α_0:

$$\text{PathExists}[\ell] \ := \ \bigvee_{\beta} \text{Path}^0[\ell, \alpha_0, \beta].$$

We have PathExists[0]. And, since every node has out-degree 1, Lemma 2(b) gives

$$\text{PathExists}[\ell] \ \rightarrow \ \text{PathExists}[\ell{+}1].$$

Let $m = 2n+1$ and $M = m^2$. By induction on ℓ from 0 to M, we obtain PathExists[M]. The ℓ-th node α_ℓ on this path of length M is definable by $\text{Path}^0[\ell, \alpha_0, \alpha_\ell]$. The node α_ℓ exists by Lemma 2(a) and is unique by Lemma 3(a). Therefore the mapping $\ell \mapsto \alpha_\ell$ defines a total injective map from $M = m^2$ to the m many nodes of $G_{n,0}$. This violates the $\text{PHP}_m^{m^2}$ pigeonhole principle, and thereby gives a contradiction. Hence, PHP_n^{n+1} is proved.

It follows immediately by standard techniques that this argument can be formalized with a quasipolynomial size Frege proof. As remarked before, [15] showed there are constant depth, polynomial size Frege proofs of $\text{PHP}_m^{m^2}$. In our setting, the proofs are quasipolynomial size, since the formulas defining the nodes α_ℓ as a function of ℓ are quasipolynomial size, not polynomial size. (In fact, [15] showed there are constant depth, polynomial size Frege proofs of PHP_m^{2m}; the lowest possible depth polynomial size Frege proofs of PHP_m^{2m} have been given by Maciel, Pitassi, and Woods [12].)

2.4 The second quasipolynomial size Frege proofs.

We now construct quasipolynomial size Frege proofs of PHP_n^{n+1} that avoid any use of a weak pigeonhole principle. The difficulty is that, without any pigeonhole principle, it is hard to see how to show that a path starting at (say) $\langle 0, 0\rangle$ cannot go forever. Ideally, we would like to prove that since $G_{n,k}$ has $2n+1$ nodes, paths in $G_{n,k}$ cannot have length $> 2n$.

Let ℓ_i equal the length of the path from the i-th source node to a sink node, we would like to prove that $\sum_i \ell_i \leq 2n$, even that $\sum_i \ell_i \leq (n+1) + (n-k)$. However, it requires counting to even state these inequalities, and this would take us back to essentially the counting-based polynomial size Frege proofs of [5]. So instead, we define another kind of path (called a "δ-path") which concatenates all the paths in $G_{n,k}$ from source to sink nodes into a single path. This will allow us to replace the inequality $\sum_i \ell_i \leq (n+1)+(n-k)$ with an explicit 1-1 correspondence that is definable by quasipolynomial size formulas.

Definition 4. *Fix n, and let $k \leq n$. The formulas $\delta^k_{\alpha,\beta}$ are defined as follows. If $\alpha \notin \mathrm{Snk}_{n,k}$ or $\beta \notin \mathrm{Src}_{n,k}$, then $\delta^k_{\alpha,\beta}$ is the formula $\gamma^k_{\alpha,\beta}$. If $\alpha = \langle 1, j\rangle \in \mathrm{Snk}_{n,k}$ and $\beta = \langle 0, i{+}1\rangle \in \mathrm{Src}_{n,k}$, then $\delta^k_{\alpha,\beta}$ is the formula*

$$\bigvee_{\ell \leq 2n} \mathrm{Path}^k[\ell, \langle 0, i\rangle, \alpha] \tag{3}$$

stating that α is reachable from $\langle 0, i\rangle$ by a path of length $\leq 2n$. If β is $\langle 0, 0\rangle$, $\delta^k_{\alpha,\beta}$ is false for all α.

Define the formulas δ-$\mathrm{Path}^k[\ell, \alpha, \beta]$ exactly like $\mathrm{Path}^k[\ell, \alpha, \beta]$, except replacing the formulas $\gamma^k_{\alpha,\beta}$ with $\delta^k_{\alpha,\beta}$. Clearly, δ-$\mathrm{Path}^k[\ell, \alpha, \beta]$ is quasipolynomial size in n and ℓ. The intent is that the δ-path starting at $\langle 0, 0\rangle$ in $G_{n,k}$ is the concatenation of the $k{+}1$ many paths starting at $\langle 0, i\rangle$ for $i = 0, 1, \ldots, k$. The definition (3) ensures this by joining the end α of the path starting at $\langle 0, i\rangle$ to the node $\beta = \langle 0, i{+}1\rangle$.

Let δ-$\mathrm{Reach}^k[\ell, \beta]$ be δ-$\mathrm{Path}^k[\ell, \langle 0, 0\rangle, \beta]$, indicating that β is reachable from $\langle 0, 0\rangle$ by a δ-path of length ℓ. The assertions in the next lemma will be proved with quasipolynomial size Frege proofs from the hypothesis $\neg\mathrm{PHP}^{n+1}_n$.

Lemma 5. *Let n be fixed and suppose $\neg\mathrm{PHP}^{n+1}_n$. Suppose $0 \leq k \leq n$, and consider the δ-path in $G_{n,k}$ starting from $\langle 0, 0\rangle$. For each $i \leq k$, there are values $\ell_1 = \ell_1(i, k)$ and $\ell_2 = \ell_2(i, k)$ such that*

- *The ℓ_1-th node on the δ-path is the source node $\langle 0, i\rangle$.*
- *The ℓ_2-th node on the path is a sink node β.*
- *There is a path in $G_{n,k}$ of length $\ell_2 - \ell_1$ from $\langle 0, i\rangle$ to β.*
- *There are no other source or sink nodes on the δ-path between positions ℓ_1 and ℓ_2.*

The length of the δ-path is $\leq 2n + 1$.

The δ-path has total length $\ell_2(k, k)$ since, by definition, the sink node reachable in $G_{n,k}$ from $\langle 0, k\rangle$ has out-degree 0 according to $\vec{\delta}^k$.

Let $\mathrm{nSnS}_{n,k}$ ("neither source nor sink") be the set of nodes in $G_{n,k}$ other than the $k{+}1$ source nodes and k sink nodes. The itemized assertions in

Lemma 5 are expressed by a quasipolynomial size propositional formula as:

$$\bigvee_{\ell_2 \leq 2n+1} \bigvee_{\ell_1 < \ell_2} \Big[\; \delta\text{-Reach}^k[\ell_1, \langle 0, i\rangle] \\ \wedge \bigvee_{\beta \in \mathrm{Snk}_{k,n}} \Big(\delta\text{-Reach}^k[\ell_2, \beta] \;\wedge\; \mathrm{Path}^k[\ell_2 - \ell_1, \langle 0, i\rangle, \beta] \Big) \\ \wedge \bigwedge_{\ell_1 < \ell < \ell_2} \bigvee_{\alpha \in \mathrm{nSnS}_{n,k}} \delta\text{-Reach}^k[\ell, \alpha] \; \Big].$$

We prove Lemma 5 informally, using arguments that can be formalized with quasipolynomial size Frege proofs. The proof is by reverse induction on k. For the base case, $k = n$, Lemma 5 is trivial, the node $\langle 0, i\rangle$ is the $(2i)$-th node on the δ-path in $G_{n,k}$. The $(2i+1)$-st node is the sink node joined by an edge to $\langle 0, i\rangle$.

Now suppose Lemma 5 holds for $k{+}1$; we must prove it for k. The induction hypothesis gives the existence of the δ-path π^{k+1} in $G_{n,k+1}$; we must define a δ-path π^k in $G_{n,k}$ and prove it enjoys the needed properties.

Suppose $i \leq k{+}1$. The induction hypothesis gives indices $\ell_1(i, k{+}1)$ and $\ell_2(i, k{+}1)$ such that the $\ell_1(i, k{+}1)$-th node of π^{k+1} is $\langle 0, i\rangle$, and such that the $\ell_2(i, k)$-th node of π^{k+1} is the sink node reachable from $\langle 1, i\rangle$ in $G_{n,k+1}$. We henceforth denote this sink node by $\beta_{i,k+1}$. Let $\ell'_1 = \ell_1(k{+}1, k{+}1)$ and $\ell'_2 = \ell_2(k{+}1, k{+}1)$. The ℓ'_1-th and ℓ'_2-th nodes on the δ-path π^{k+1} in $G_{n,k+1}$ are $\langle 0, k{+}1\rangle$ and $\beta_{n,k+1}$, respectively. The length of π^{k+1} is ℓ'_2.

The argument now splits into two cases. First, suppose there is no $\langle 0, i\rangle \in \mathrm{Src}_{n,k}$ (that is, no $i \leq k$) such that $\beta_{i,k+1}$ equals $\langle 1, k\rangle$. In this case, since the only change between $G_{n,k+1}$ and $G_{n,k}$ is the addition of the edge from $\langle 1, k\rangle$ to $\langle 0, k{+}1\rangle$, we have that the first $\ell_2(k, k{+}1)$ nodes of π^{k+1} already form the desired δ-path in $G_{n,k}$. The desired properties of π^k follow immediately from the properties of π^{k+1}. In this case, the δ-path π^k is shorter than the δ-path π^{k+1}.

Second, suppose that $\langle 1, k\rangle$ is equal to $\beta_{i,k+1}$ for some $i \leq k$. Let $\ell''_1 = \ell_1(i, k{+}1)$ and $\ell''_2 = \ell_2(i, k{+}1)$, so that the ℓ''_1-th and ℓ''_2-th nodes on π^{k+1} are $\langle 0, i\rangle$ and $\beta_{i,k+1} = \langle 1, k\rangle$, respectively. The desired δ-path π^k consists of the following: the first ℓ''_2 many nodes of π^{k+1}, followed by the nodes in positions ℓ'_1 through ℓ'_2 of π^{k+1}, and then followed by the nodes in positions $\ell''_2{+}1$ through $\ell'_1{-}1$ of π^{k+1}. It is straightforward to verify that the desired properties of π^k hold, given the properties of π^{k+1}. (When formalizing this argument as a Frege proof, the argument splits into separate cases for each of the polynomially many possible values for ℓ'_1, ℓ'_2, ℓ''_1, and ℓ''_2.)

That completes the proof of Lemma 5.

The proof of Theorem 1 is now easy. It is trivial to see that when $k = 0$, there can be no such δ-path π^0, since $G_{n,0}$ has no sink nodes. Alternately, the argument can stop at the case $k = 1$, and argue as in the earlier-sketched argument by Cook and Reckhow for extended Frege proofs. So this completes the proof of Theorem 1.

We claim that these quasipolynomial size Frege proofs really do intensionally simulate the extended Frege proofs of Cook and Reckhow. To prove this, it will suffice to show that there are quasipolynomial size proofs of the formulas

$$\varphi^k_{i,j} \;\leftrightarrow\; \varphi^{k+1}_{i,j} \vee (\varphi^{k+1}_{i,k} \wedge \varphi^{k+1}_{k+1,j}) \tag{4}$$

corresponding to the introduction of the variables $q^k_{i,j}$ in the e$\mathcal{F}$ proofs by the extension rule (1). The formulas (4) are shown successively for $k = n, \ldots, 2, 1, 0$: indeed they were essentially proved above as part of the proof of Lemma 5. The only subtle point is that the formulas $\varphi^k_{i,j}$ were defined by (2) in terms of the existence of a path of length at most $N = 2n+1$ whereas Lemma 5 was proved with δ-paths. It is easy to prove the equivalence of these two approaches. From Lemma 5, we know there is some path from $\langle 0, i\rangle$ to a sink node, and its length is at most $2n + 1$. And from Lemma 2, there is only one sink node reachable by a path from $\langle 0, i\rangle$.

Acknowledgements. We thank James Aisenberg and an anonymous referee for comments.

References

[1] J. Aisenberg, M. L. Bonet, and S. Buss, *Quasipolynomial-size Frege proof of Frankl's theorem on the trace of finite sets*. To appear in Journal of Symbolic Logic, 201?

[2] P. Beame and T. Pitassi, *Propositional proof complexity: Past, present and future*, in Current Trends in Theoretical Computer Science Entering the 21st Century, G. Paun, G. Rozenberg, and A. Salomaa, eds., World Scientific, 2001, pp. 42–70. Earlier version appeared in *Computational Complexity Column*, Bulletin of the EATCS, 2000.

[3] A. Beckmann and S. R. Buss, *Improved witnessing and local improvement principles for second-order bounded arithmetic*, ACM Transactions on Computational Logic, 15 (2014). Article 2, 35 pages.

[4] M. L. Bonet, S. R. Buss, and T. Pitassi, *Are there hard examples for Frege systems?*, in Feasible Mathematics II, P. Clote and J. Remmel, eds., Boston, 1995, Birkhäuser, pp. 30–56.

[5] S. R. Buss, *Polynomial size proofs of the propositional pigeonhole principle*, Journal of Symbolic Logic, 52 (1987), pp. 916–927.

[6] ——, *Propositional proof complexity: An introduction*, in Computational Logic, U. Berger and H. Schwichtenberg, eds., Springer-Verlag, Berlin, 1999, pp. 127–178.

[7] ——, *Towards NP-P via proof complexity and proof search*, Annals of Pure and Applied Logic, 163 (2012), pp. 1163–1182.

[8] S. A. Cook and R. A. Reckhow, *The relative efficiency of propositional proof systems*, Journal of Symbolic Logic, 44 (1979), pp. 36–50.

[9] P. Hrubeš and I. Tzameret, *Short proofs for determinant identities*. To appear in SIAM J. Computing, 201?

[10] G. Istrate and A. Crăciun, *Proof complexity and the Kneser-Lovász theorem*, in Theory and Applications of Satisfiability Testing (SAT), Lecture Notes in Computer Science 8561, Springer Verlag, 2014, pp. 138–153.

[11] L. A. Kołodziejczyk, P. Nguyen, and N. Thapen, *The provably total NP search problems of weak second-order bounded arithmetic*, Annals of Pure and Applied Logic, 162 (2011), pp. 419–446.

[12] A. Maciel, T. Pitassi, and A. R. Woods, *A new proof of the weak pigeonhole principle*, Journal of Computer and System Sciences, 64 (2002), pp. 843–872.

[13] A. Nozaki, T. Arai, and N. H. Arai, *Polynomial-size Frege proofs of Bollobás' theorem on the trace of sets*, Proceedings of the Japan Academy, Series A. Math. Sci., 84 (2008), pp. 159–161.

[14] C. H. Papadimitriou, *On the complexity of the parity argument and other inefficient proofs of existence*, Journal of Computer and System Sciences, 48 (1994), pp. 498–532.

[15] 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 (1988), pp. 1235–1244.

[16] N. Segerlind, *The complexity of propositional proofs*, Bulletin of Symbolic Logic, 13 (2007), pp. 417–481.