Improved Bounds on the Weak Pigeonhole Principle and Infinitely Many Primes from Weaker Axioms

Albert Atserias*

Departament de Llenguatges i Sistemes Informàtics
Universitat Politècnica de Catalunya
Barcelona, Spain
atserias@lsi.upc.es

Abstract. We show that the known bounded-depth proofs of the Weak Pigeonhole Principle PHP_n^{2n} in size $n^{O(\log(n))}$ are not optimal in terms of size. More precisely, we give a size-depth trade-off upper bound: there are proofs of size $n^{O(d(\log(n))^{2/d})}$ and depth $O(d)$. This solves an open problem of Maciel, Pitassi and Woods (2000). Our technique requires formalizing the ideas underlying Nepomnjaščij's Theorem which might be of independent interest. Moreover, our result implies a proof of the unboundedness of primes in $I\Delta_0$ with a provably weaker 'large number assumption' than previously needed.

Keywords. Proof Complexity, Weak Pigeonhole Principle, Bounded Arithmetic.

1 Introduction

The Pigeonhole Principle PHP_n^{n+1} is a fundamental statement about cardinalities of finite sets. It says that $\{0,\ldots,n\}$ cannot be mapped injectively into $\{0,\ldots,n-1\}$. The Pigeonhole Principle is at the heart of many mathematical arguments, although implicitely quite often. As a matter of fact, it implies the Induction Principle, and so its apparent self-evidence makes it even more interesting.

The general form of the Pigeonhole Principle PHP_n^m states that if $m > n$, then $\{0,\ldots,m-1\}$ cannot be mapped injectively into $\{0,\ldots,n-1\}$. When m is substantially larger than n, say $m = 2n$, the principle is called the Weak Pigeonhole Principle. Still, it is quite often the case that this weaker form is enough to carry over many arguments, notably in finite combinatorics and number theory.

The complexity of proving a propositional encoding of the Pigeonhole Principle has been investigated in depth since the problem was proposed by Cook and Reckhow in connection with the $\mathbf{NP} \stackrel{?}{=} \mathbf{coNP}$ question [5]. Haken proved that PHP_n^{n+1} requires exponential-size proofs in Resolution [10]. On the other hand,

* Supported by the CUR, Generalitat de Catalunya, through grant 1999FI 00532. Partially supported by ALCOM-FT, IST-99-14186.

Buss proved that PHP_n^{n+1} has polynomial-size proofs in Frege systems [3]. Ajtai proved a superpolynomial lower bound for bounded-depth Frege systems [1], and that was improved to an exponential lower bound by Pitassi, Beame and Impagliazzo [19], and Krajícek, Pudlák and Woods [13] independently. For the Weak Pigeonhole Principle PHP_n^{2n}, the situation is quite different. While it is known that PHP_n^{2n} requires exponential-size proofs in Resolution [4], Paris, Wilkie and Woods [18, 12] proved that it has quasipolynomial-size $(n^{O(\log(n))})$ proofs in bounded-depth Frege systems. More recently, Maciel, Pitassi and Woods [15] gave a new quasipolynomial-size proof of optimal depth. Both papers left open, however, whether depth could be traded for size; that is, whether allowing more depth in the proof would allow us reduce the size below $n^{\log(n)}$. We note that such a trade-off is known for the even weaker Pigeonhole Principle $\mathrm{PHP}_n^{n^2}$ [18].

The proofs of Paris, Wilkie and Woods [18] and Maciel, Pitassi and Woods [15] consist in reducing PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$. In both cases, they build an injective map from $\{0,\ldots,n^2-1\}$ to $\{0,\ldots,n-1\}$ by repeatedly composing a suposedly injective map from $\{0,\ldots,2n-1\}$ to $\{0,\ldots,n-1\}$. The difference in their proofs is, essentially, in the proof of $\mathrm{PHP}_n^{n^2}$. Our new contribution is showing that the repeated composition technique can be made more efficient in terms of size. That is, we reduce PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$ in size $n^{o(\log(n))}$ (notice the small oh). The price we need to pay for that is an increase in depth. More precisely, we show that PHP_n^{2n} reduces to $\mathrm{PHP}_n^{n^2}$ in size $n^{O(d(\log(n))^{2/d})}$ and depth d. This gives us the desired size-depth trade-off upper bound for PHP_n^{2n} since $\mathrm{PHP}_n^{n^2}$ is provable in size $n^{O(\log^{(2)}(n))}$ and depth $O(1)$ [18, 12].

The most interesting particular case of our size-depth trade-off is when $d = O(1)$ since it proves that the previously known upper bound in bounded-depth Frege is not optimal. Indeed, $n^{O(d(\log(n))^{2/d})}$ grows slower than $n^{c\log(n)}$ for any constants $d > 2$ and $c > 0$. Thus, any lower bound proof will have to focus on a bound weaker than $n^{\log(n)^{\epsilon}}$ for any $\epsilon > 0$. We believe this is valuable information. The other interesting particular case is when $d = O(\log\log(n))$. In that case we obtain a proof of size $n^{O(\log\log(n))}$ and depth $O(\log\log(n))$. The bound $n^{O(\log\log(n))}$ is new in the context of PHP_n^{2n}.

The method that we use to reduce the size of the composition technique is inspired from the theory of automata. We observe that checking whether b is the image of a under repeated composition of a function f is a reachability problem in a graph. Therefore, one can use (an analogue of) Savitch's Theorem to efficiently solve the reachability problem. We are more ambitious and we use ideas from an old theorem of Nepomnjaščij that achieves a size-depth trade-off for the same problem [17]. We formalize the ideas in Nepomnjaščij's Theorem into a theorem of Bounded Arithmetic with an automatic translation into propositional Gentzen Calculus. This formalization may be of independent interest. We note that Nepomnjaščij's Theorem has received a renewed deal of attention recently in the context of time-space trade-off lower bounds for the satisfiability problem [6, 14, 7].

The new bounds on the Weak Pigeonhole Principle that we obtain have some consequences for Feasible Number Theory whose aim is to develop as much number theory as possible without exponentiation. The main open problem of the field is whether the bounded arithmetic theory $I\Delta_0$ can prove that there are unboundedly many primes [18, 16, 2]. Of course, Euclides' proof cannot be carried over in $I\Delta_0$ since exponentially large numbers are required in the proof. In a major breakthrough, Woods [21] showed that exponentiation can be replaced by a combinatorial argument using the Pigeonhole Principle PHP_n^{n+1}, and Paris, Wilkie and Woods [18] realized that the Weak Pigeonhole Principle PHP_n^{2n} was enough for that proof. As a corollary to their results, they show that $I\Delta_0$ augmented with the statement that $x^{\log(x)}$ exists proves that $(\exists y)(y > x \wedge prime(y))$. Our results improve this to show that, for every standard natural number k, the theory $I\Delta_0$ augmented with the statement that $x^{\log(x)^{1/k}}$ exists proves that $(\exists y)(y > x \wedge prime(y))$. Therefore, the large number assumption "$x^{\log(x)}$ exists" is not optimal. Indeed, for $k > 1$, one can build a model of $I\Delta_0$ with a non-standard element a such that $a^{\log(a)^{1/k}}$ exists in the model but $a^{\log(a)}$ does not.

2 The Proof of the WPHP in Propositional Logic

The propositional form of the Weak Pigeonhole Principle PHP_n^m that we use is formalized by the following sequent:

$$\bigwedge_{i=1}^{m} \bigvee_{j=1}^{n} p_{i,j} \vdash \bigvee_{k=1}^{n} \bigvee_{\substack{i,j=1 \\ i \neq j}}^{m} p_{i,k} \wedge p_{j,k}.$$

We will work with the propositional fragment of the sequent calculus LK. We refer the reader to any standard textbook for a definition [20, 12].

Our goal is to prove a size-depth trade-off upper bound for proofs of PHP_n^{2n}. As mentioned in the introduction, our technique consists in reducing PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$. The difference with previous reductions is that our composition of mappings is done efficiently mimicking the proofs of Savitch's and Nepomnjaščij's Theorems in Complexity Theory. To complete the proof, we will use the fact that $\mathrm{PHP}_n^{n^2}$ has LK proofs of size $n^{O(\log^{(d)}(n))}$ and depth $O(d)$, where $\log^{(d)}(n)$ is the d-wise composition of log with itself [18, 12].

For the sake of clarity of exposition, it is more convenient to prove the following extreme case of the trade-off first.

Theorem 1. PHP_n^{2n} *has* LK *proofs of size* $n^{O(\log\log(n))}$ *and depth* $3\log\log(n)$.

Proof: Let $d = \log\log(n)$. For every $\alpha \in \{0,1\}^{\leq d}$, we define numbers L_α and R_α inductively as follows: Let $L_\lambda = \log(n)$, $R_\lambda = 0$, and $L_{\alpha 0} = L_\alpha$, $R_{\alpha 1} = R_\alpha$, $L_{\alpha 1} = R_{\alpha 0} = \frac{1}{2}(L_\alpha + R_\alpha)$. Now we define sets A_α, B_α and C_α as follows: Let $A_\alpha = \{0, \ldots, 2^{L_\alpha} n - 1\}$, $B_\alpha = \{0, \ldots, 2^{R_\alpha} n - 1\}$ and $C_\alpha = \{0, \ldots, 2^{\frac{1}{2}(L_\alpha + R_\alpha)} n - 1\}$. Observe that $A_\lambda = \{0, \ldots, n^2 - 1\}$ and $B_\lambda = \{0, \ldots, n - 1\}$.

For $\alpha \in \{0,1\}^d$, $x \in A_\alpha$, $y \in B_\alpha$, let R^α_{xy} be defined as follows: If $\lfloor x/2n \rfloor \neq \lfloor y/n \rfloor$, define $R^\alpha_{xy} = 0$; otherwise, define $R^\alpha_{xy} = P_{x'y'}$ where $x' = x \bmod 2n$ and $y' = y \bmod n$. For $\alpha \in \{0,1\}^{<d}$, $x \in A_\alpha$, $y \in B_\alpha$, define $R^\alpha_{xy} = \bigvee_{z \in C_\alpha} R^{\alpha 0}_{xz} R^{\alpha 1}_{zy}$. It is easy to see that the size of R^α_{xy} is bounded by $n^{2(d-|\alpha|)}$ and the depth is bounded by $2(d-|\alpha|)$. In particular, R^λ_{xy} has size bounded by $n^{2\log\log(n)}$ and depth bounded by $2\log\log(n)$.

We want to prove the following sequents

$$\bigwedge_{\alpha \in \{0,1\}^k} \bigwedge_x \bigvee_y R^\alpha_{xy} \to \bigvee_{\alpha \in \{0,1\}^k} \bigvee_y \bigvee_{x_1 \neq x_2} R^\alpha_{x_1 y} R^\alpha_{x_2 y} \tag{1}$$

in size bounded by n^{ck} and depth bounded by $3k$, where c is a sufficiently large constant independent of n ($c = 20$ should work). When $k = d = \log\log(n)$, it is easy to see that sequent (1) is equivalent to PHP^{2n}_n after contraction of repeated formulas, and the theorem will follow. We fix a sufficiently large n, and proceed by induction on k.

Observe that the base case $k = 0$ is precisely the sequent $\mathrm{PHP}^{n^2}_n(R^\lambda)$. Since the sequent $\mathrm{PHP}^{n^2}_n$ has LK proofs of size $n^{O(\log^{(3)}(n))}$ and constant depth, it follows that $\mathrm{PHP}^{n^2}_n(R^\lambda)$ has LK-proofs of size bounded by $n^{c\log\log(n)}$ and depth bounded by $3\log\log(n)$. Here is where we need n to be sufficiently large.

Suppose next that we have proved sequent (1) for a $k > 0$ in size n^{ck} and depth $3k$. We prove it for $k+1$. We first prove the following sequents for every $\alpha \in \{0,1\}^k$ and w:

$$\bigwedge_x \bigvee_y R^{\alpha 0}_{xy} \wedge \bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \to \bigvee_y R^\alpha_{wy}. \tag{2}$$

Recall that R^α_{wy} stands for $\bigvee_z R^{\alpha 0}_{wz} R^{\alpha 1}_{zy}$. Start with the sequents $R^{\alpha 0}_{wz} \to R^{\alpha 0}_{wz}$ and $\bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \to \bigvee_y R^{\alpha 1}_{zy}$, and apply right $\wedge$ introduction to obtain

$$R^{\alpha 0}_{wz}, \bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \to R^{\alpha 0}_{wz} \wedge \bigvee_y R^{\alpha 1}_{zy}. \tag{3}$$

By distributivity we easily get

$$R^{\alpha 0}_{wz}, \bigwedge_x \bigvee_y R^{\alpha 1}_{xy} \to \bigvee_y R^{\alpha 0}_{wz} R^{\alpha 1}_{zy}. \tag{4}$$

By left $\vee$-introduction, left weakening, left $\wedge$-introduction, right $\vee$-introduction and commutativity of $\vee$, in this order, we get the desired sequent (2).

Next we prove the following sequents for every $\alpha \in \{0,1\}^k$, $w_1 \neq w_2$ and y:

$$R^\alpha_{w_1 y} R^\alpha_{w_2 y} \to \bigvee_z \bigvee_{x_1 \neq x_2} R^{\alpha 0}_{x_1 z} R^{\alpha 0}_{x_2 z} \vee \bigvee_z \bigvee_{x_1 \neq x_2} R^{\alpha 1}_{x_1 z} R^{\alpha 1}_{x_2 z}. \tag{5}$$

Recall that $R^\alpha_{w_i y}$ stands for $\bigvee_z R^{\alpha 0}_{w_i z} R^{\alpha 1}_{zy}$. Using distributivity, derive the sequent

$$R^\alpha_{w_1 y} R^\alpha_{w_2 y} \to \bigvee_{z_1, z_2} R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y}. \tag{6}$$

For $z_1 = z_2$, derive the sequent $R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y} \to R^{\alpha 0}_{w_1 z_1} R^{\alpha 0}_{w_2 z_2}$. For $z_1 \neq z_2$, derive the sequent $R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y} \to R^{\alpha 1}_{z_1 y} R^{\alpha 1}_{z_2 y}$. Left $\vee$-introduction and right $\vee$-introduction gives the sequent

$$\bigvee_{z_1, z_2} R^{\alpha 0}_{w_1 z_1} R^{\alpha 1}_{z_1 y} R^{\alpha 0}_{w_2 z_2} R^{\alpha 1}_{z_2 y} \to \bigvee_{z_1 = z_2} R^{\alpha 0}_{w_1 z_1} R^{\alpha 0}_{w_2 z_2} \vee \bigvee_{z_1 \neq z_2} R^{\alpha 1}_{z_1 y} R^{\alpha 1}_{z_2 y}. \tag{7}$$

A cut with sequent (6), right weakening, right $\vee$-introduction and commutativity of $\vee$, in this order, give the desired sequent (5).

Now combine all sequents (2) by right $\wedge$-introduction, left $\wedge$-introduction and commutativity of $\wedge$, in this order, to obtain

$$\bigwedge_{\alpha \in \{0,1\}^{k+1}} \bigwedge_x \bigvee_y R^{\alpha}_{xy} \to \bigwedge_{\alpha \in \{0,1\}^k} \bigwedge_x \bigvee_y R^{\alpha}_{xy}. \tag{8}$$

Similarly, combine all sequents (5) by left $\vee$-introduction, right $\vee$-introduction and commutativity of $\vee$, in this order, to obtain

$$\bigvee_{\alpha \in \{0,1\}^k} \bigvee_y \bigvee_{x_1 \neq x_2} R^{\alpha}_{x_1 y} R^{\alpha}_{x_2 y} \to \bigvee_{\alpha \in \{0,1\}^{k+1}} \bigvee_y \bigvee_{w_1 \neq w_2} R^{\alpha}_{w_1 y} R^{\alpha}_{w_2 y}. \tag{9}$$

Finally, two cuts using sequents (1), (8) and (9) give the desired result for $k+1$. If c is sufficiently large, it is easy to check that the size of this proof is bounded by $n^{c(k+1)}$ and its depth is bounded by $3(k+1)$. This completes the induction step. □

An analogous argument mimicking the proof of Nepomnjaščij's Theorem, instead of Savitch's Theorem as above, would give a general size-depth trade-off upper bound. Since the notation in the proof would get fairly tedious, we prefer to state it without proof and get it as a corollary to Theorem 6 below (see the end of Section 4).

Theorem 2. PHP^{2n}_n *has* LK *proofs of size* $n^{O(d(\log(n))^{2/d})}$ *and depth* d.

3 Formalization of Nepomnjaščij's Theorem

Let us briefly recall the proof of a general form of Nepomnjaščij's Theorem. This will be of help later.

Theorem 3. (General Form of Nepomnjaščij's Theorem) *Let* $K = K(n)$, $T = T(n)$ *and* $S = S(n)$ *be time-constructible functions such that* $K(n) \geq 2$. *For every non-deterministic Turing machine running in simultaneous time* T *and space* S, *there exists an equivalent alternating Turing machine running in time* $O(SK \log(T)/\log(K))$ *and* $2\log(T)/\log(K)$ *alternations.*

Proof: Let M be a non-deterministic Turing machine running in simultaneous time T and space S. The idea is to divide the reachability problem between

configurations of M into many equivalent subproblems of smaller size. Hence, configuration C_K is reachable from configuration C_0 in N steps if and only if there exist $K-1$ intermediate configurations $C_1, \ldots, C_{K-1}$ such that for every $i \in \{0, \ldots, K-1\}$, configuration C_{i+1} is reachable from configuration C_i in N/K steps. An alternating Turing machine can existentially quantify those intermediate configurations, and universally branch to check that every two consecutive configurations are reachable from each other in the appropriate number of steps. Applying this recursively yields an alternating machine that checks whether an accepting configuration is reachable from the initial configuration of M.

The details of the calculations follow. Since M runs in space S, each configuration can be coded by a binary word of length $O(S)$. The outermost reachability problem has length T since M runs in time T. The second level of reachability subproblems have length T/K. In general, the subproblems at level i have length T/K^i. After $\log(T)/\log(K)$ levels, we reached a trivial reachability subproblem. Therefore, the whole computation of the simulating machine takes time $O(SK\log(T)/\log(K))$ and $2\log(T)/\log(K)$ alternations. □

Our next goal is to formalize the ideas of Theorem 3 into a theorem of the bounded arithmetic theory $I\Delta_0$. We will use the beautiful arithmetization in $I\Delta_0$ of Chapter V, Section 3, in the book of Hájek and Pudlák [9]. In summary, the arithmetization allows us to manipulate sequences provably in $I\Delta_0$. Thus, there are formulas $Seq(s)$ meaning that s is the code of a sequence, $(s)_i = x$ meaning that the i-th element of the sequence s is x, and $s \frown t = p$ meaning that sequence p is the result of appending sequence t to the end of sequence s. The coding is so that if $s = (x)$, the sequence with x as its only element, then $I\Delta_0$ proves $s \leq 9 \cdot x^2$. Morevoer, $I\Delta_0$ proves that for every sequence s, the bound $s \frown (x) \leq 9 \cdot x^2 \cdot s$ holds (see Lemma 3.7 in page 297 of [9]). It follows that $I\Delta_0$ proves that if l^r exists below n and s is a sequence of length r all of whose elements are smaller than l, then $s \leq 9^r \cdot l^{2r}$ (by Δ_0-induction on r). Therefore, the coding is fairly close to its information theoretic bound. Of course, $I\Delta_0$ can prove several other obvious facts about $Seq(s)$, $(s)_i$ and $s \frown t$ (see [9] for details).

Let $\theta(x,y)$ be a Δ_0-formula in a language L extending the usual language of arithmetic $\{+, \times, \leq\}$. Obviously, $\theta(x,y)$ defines a binary relation on any model for the language L that may be interpreted as an infinite directed graph. We define Δ_0-formulas $\Theta_i(x,y)$, with certain parameters, meaning that y is reachable from x under certain conditions that depend on the parameters. More precisely, let $\Theta_0(x,y,t,r,l,n) = \theta(x,y)$ (note that the parameters t, r and l are not used for the moment). Inductively, we define $\Theta_{k+1}(x,y,t,r,l,n)$ as follows:

$$(\exists z \leq n)(Seq(z) \wedge (z)_0 = x \wedge (z)_r = y \wedge (\forall i < r+1)((z)_i \leq l) \wedge$$
$$\wedge(\forall i < r)(\exists c, c' \leq z)((z)_i = c \wedge (z)_{i+1} = c' \wedge \Theta_k(c,c',t,t,l,n))).$$

Informally, the formula $\Theta_k(x,y,t,t,l,n)$ says that y is reachable from x in t^k steps according to the directed graph defined by $\theta(x,y)$ as long as each number in the path is bounded by l. The following theorem states this in the form of recursive equations:

Theorem 4. *Let $M \models I\Delta_0(L)$, and let $x, y, t, r, l, n \in M$ be such that $x, y \leq l$, l^t exists in M, $r < t$, and $9^{t+1} \cdot l^{2(t+1)} \leq n$. Then,*

(i) $\Theta_{k+1}(x, y, t, 0, l, n) \leftrightarrow x = y$,
(ii) $\Theta_{k+1}(x, y, t, r+1, l, n) \leftrightarrow (\exists c \leq l)(\Theta_k(x, c, t, t, l, n) \wedge \Theta_{k+1}(c, y, t, r, l, n))$

hold in M.

Proof: (i) Assume $\Theta_{k+1}(x, y, t, 0, l, n)$ holds. Then, for some $z \leq n$, we have $(z)_0 = x$ and $(z)_0 = y$. Hence, $x = y$. Conversely, if $x = y$, then $z = (x) \leq 9 \cdot x^2 \leq 9 \cdot l^2 \leq n$ is a witness for the existential quantifier in $\Theta_{k+1}(x, y, t, 0, l, n)$. (ii) Assume $\Theta_{k+1}(x, y, t, r+1, l, n)$ holds. Let $z \leq n$ be the witness for its existential quantifier. Let $c = (z)_1$ and note that $(z)_0 = x$. Obviously, $c \leq l$ and $\Theta_k(x, c, t, t, l, n)$ holds. Now let z' be the sequence that results from z when $(z)_0$ is dropped (definable in $I\Delta_0$). It is easily seen that $z' \leq z \leq n$, and $\Theta_{k+1}(c, y, t, r, l, n)$ holds with z' witnessing its existential quantifier. Conversely, if $\Theta_k(x, c, t, t, l, n)$ and $\Theta_{k+1}(c, y, t, r, l, n)$ hold, let $z' \leq n$ be a witness for the existential quantifier in the latter. We can assume that z' codes a sequence of $r+1$ numbers bounded by l each; if not, just trim z' to the first $r+1$ numbers (definable in $I\Delta_0$) and the result is still a witness of the existential quantifier. Moreover, $z' \leq 9^{r+1} \cdot l^{2(r+1)}$. Since $x \leq l$ and $r < t$, the sequence $(x) \frown z'$ is coded by some $z \leq 9^{r+2} \cdot l^{2(r+2)} \leq n$. such a z is a witness for the existential quantifier in $\Theta_{k+1}(x, y, t, r+1, l, n)$ and we are done. □

The reader will notice that the recursive equations in Theorem 4 correspond to the inductive definition of the transitive closure of the graph defined by Θ_k. The innermost level of stratification, namely $k = 0$, is the inductive definition of the transitive closure of the graph defined by θ. It is in this sense that we interpret Theorem 4 as a formalization of Nepomnjaščij's Theorem.

4 The Proof of the WPHP in Bounded Arithmetic

The graph of the exponentiation function $x = y^z$ is definable by a Δ_0-formula on the natural numbers. Moreover, Pudlák gave a definition with the basic properties being provable in $I\Delta_0$. Similarly, one can define $y = \lceil \log(x) \rceil$, and $y = \lceil \log^{(k)}(x) \rceil$ in $I\Delta_0$. We make the convention that when expressions such as $\log(a)$ or $(\log(a))^\epsilon$ do not come up integer numbers, the nearest larger integer is assumed unless specified otherwise. Thus, $(\log(a))^\epsilon$ really stands for $\lceil \log(a)^\epsilon \rceil$.

We let L be the usual language of arithmetic $\{+, \times, \leq\}$ extended by a unary function symbol α. We denote $I\Delta_0(L)$ (see the previous section) by $I\Delta_0(\alpha)$. The Weak Pigeonhole Principle PHP_n^m is formalized by the following statement:

$$(\forall x < m)(\alpha(x) < n)) \rightarrow (\exists x, y < m)(x \neq y \wedge \alpha(x) = \alpha(y)).$$

We will abbreviate this statement by $\neg\alpha : m \xrightarrow{1-1} n$. We will make use of the following result:

Theorem 5. [18] *For every $K > 0$,*

$$I\Delta_0(\alpha) \vdash (\exists y)(y = x^{\log^{(K)}(x)}) \rightarrow \neg\alpha : x^2 \xrightarrow{1-1} x.$$

Here, $\log^{(0)}(x) = x$ *and* $\log^{(k+1)}(x) = \log(\log^{(k)}(x))$.

Next, we formalize the reduction of PHP_n^{2n} to $\mathrm{PHP}_n^{n^2}$ using the result of Section 3.

Theorem 6. *For every $K > 0$,*

$$I\Delta_0(\alpha) \vdash (\exists y)(y = x^{(\log(x))^{1/K}}) \rightarrow \neg\alpha : 2x \xrightarrow{1-1} x.$$

Proof: Let $\epsilon = 1/K$. Let $\mathbf{M} = (M, F)$ be a model of $I\Delta_0(\alpha)$, let $a \in M$ be such that $a^{(\log(a))^\epsilon}$ exists in M, and assume for contradiction that $F : 2a \xrightarrow{1-1} a$. Let $T = (\log(a))^\epsilon$, $L = a^2$, and $N = 9^{T+1} \cdot L^{2(T+1)}$. Observe that N exists in M since $a^{(\log(a))^\epsilon}$ exists, and M is closed under multiplication. Define $\theta(x, y)$ as follows (see the text that follows the formula for the intuition):

$$(\exists r < 2a)(\exists r' < a)(\exists q, q' < a)(x = 2aq + r \wedge y = aq' + r' \wedge \alpha(r) = r' \wedge q = q').$$

Note that this Δ_0-formula could be informally abbreviated by

$$y \bmod a = \alpha(x \bmod 2a) \wedge \lfloor y/a \rfloor = \lfloor x/2a \rfloor$$

when $x, y \in \{0, \ldots, a^2 - 1\}$.

Lemma 1. $\theta(x, y)^{\mathbf{M}} : 2^{i+1}a \xrightarrow{1-1} 2^i a$ *for every* $i < \log(a)$.

Proof: Given $u \in 2^{i+1}a$, let $v = \lfloor u/2a \rfloor \cdot a + \alpha(u \bmod 2a)$. It is not hard to see that $v \in 2^i a$ and $\theta(u, v)$ holds. Moreover, if $w \in 2^i a$ is such that $\theta(u, w)$ holds, then $\lfloor v/a \rfloor = \lfloor u/2a \rfloor = \lfloor w/a \rfloor$ and $v \bmod a = \alpha(u \bmod 2a) = w \bmod a$. Hence, $v = w$. This shows that $\theta(x, y)^{\mathbf{M}}$ is the graph of a function from $2^{i+1}a$ to $2^i a$.

We show next that the function is one-to-one. Let $u, v \in 2^{i+1}a$ and $w \in 2^i a$ be such that $\theta(u, w)$ and $\theta(v, w)$. Then, $\lfloor u/2a \rfloor = \lfloor v/2a \rfloor = \lfloor w/a \rfloor$ and $F(u \bmod 2a) = F(v \bmod 2a) = w \bmod a$. Since F is one-to-one, it must be then that $u \bmod 2a = v \bmod 2a$. Hence, $u = 2a \cdot \lfloor u/2a \rfloor + (u \bmod 2a) = 2a \cdot \lfloor v/2a \rfloor + (v \bmod 2a) = v$. □

Lemma 2. $\Theta_K(x, y, T, T, L, N)^{\mathbf{M}} : a^2 \xrightarrow{1-1} a$.

Proof: We prove that for every $k \leq K$, the formula $\Theta_k(x, y, T, T, L, N)$ defines a one-to-one mapping $\Theta_k : 2^{(i+1)T^k}a \rightarrow 2^{iT^k}a$ for every $i < \log(a)/T^k$. The lemma will be proved since $T^K = ((\log(a))^{1/K})^K = \log(a)$ (in fact, $T^K \geq \log(a)$ by our convention on rounding). The proof is by induction on k (this induction is outside M).

Lemma 1 takes care of the base case $k = 0$. We turn to the inductive case $0 < k \leq K$. Fix $i < \log(a)/T^k$. We prove that for every $r \leq T$, the formula

$\Theta_k(x, y, T, r, L, N)$ defines a one-to-one mapping $\Theta_k^r : 2^{(iT+r)T^{k-1}}a \to 2^{iT^k}a$. That is, we prove that for every $r \leq T$, $x < 2^{(iT+r)T^{k-1}}a$ and $y < 2^{iT^k}a$,

$$\Theta_k(x, z, T, r, L, N) \wedge \Theta_k(y, z, T, r, L, N) \to x = y$$

holds in $\mathbf{M}$. We use the schema of Δ_0-induction on r in the Δ_0-formula above. The base case $r = 0$ is immediate since $\Theta_k(x, y, T, 0, L, N)$ defines the identity by Theorem 4. Suppose that $0 < r \leq T$, and that $\Theta_k(x, y, T, r-1, L, N)$ defines a one-to-one mapping $\Theta_k^{r-1} : 2^{(iT+r-1)T^{k-1}}a \to 2^{iT^k}a$. Since $\Theta_{k-1}(x, y, T, T, L, N)$ defines a one-to-one mapping $\Theta_{k-1} : 2^{(iT+r)T^{k-1}}a \to 2^{(iT+r-1)T^{k-1}}a$ by induction hypothesis on k, and since $\Theta_k(x, y, T, r, L, N)$ defines the composition of Θ_{k-1} and Θ_k^{r-1} by Theorem 4 (observe that $x, y \leq L$, $r - 1 < T$ and $9^{T+1} \cdot L^{2(T+1)} \leq N$), it follows that $\Theta_k(x, y, T, r, L, N)$ defines a one-to-one mapping $\Theta_k^r : 2^{(iT+r)T^{k-1}}a \to 2^{iT^k}a$ as required. □

Since Θ_K is a $\Delta_0(\alpha)$ formula, we have that $(M, \Theta_K(x, y, T, T, L, N)^{\mathbf{M}}) \models I\Delta_0(\alpha)$. Moreover, $a^{\log^{(2)}(a)} < a^{(\log(a))^\epsilon}$ exists in M. It follows from Theorem 5 that $\Theta_K(x, y, T, T, L, N)^{\mathbf{M}}$ is not a one-to-one mapping from $a^2 \to a$; a contradiction to Lemma 2. □

It is well-known that proofs in $I\Delta_0(\alpha)$ translate into bounded-depth LK proofs of polynomial-size. When statements of the form "$f(x)$ exists" are required as in Theorem 4, the translations come up of size $f(n)^{O(1)}$ (see [12], for example). This gives us Theorem 2 as a corollary.

5 Infinitude of Primes

The existence of infinitely many primes is not guaranteed in weak fragments of arithmetic. For example, it is known that I_{open}, Peano Arithmetic with induction restricted to open formulas, has models with a largest prime [16]. It is an open problem whether $I\Delta_0$ proves the infinitude of primes. It is known, however, that $I\Delta_0$ augmented with the axiom $(\forall x)(\exists y)(y = x^{\log(x)})$ proves it. In addition, a *single* application of this axiom suffices. More precisely,[1]

Theorem 7. [18] $I\Delta_0 \vdash (\exists y)(y = x^{\log(x)}) \to (\exists y)(y > x \wedge \mathit{prime}(y))$.

The aim of this section is to show that a weaker axiom suffices, and so the existence of $x^{\log(x)}$ is not the optimal large number assumption. Namely,

Theorem 8. *$I\Delta_0 \vdash (\exists y)(y = x^{(\log(x))^{1/K}}) \to (\exists y)(y > x \wedge \mathit{prime}(y))$ for every $K > 0$. Moreover, there exists a model $M \models I\Delta_0$ with a non-standard element $a \in M$ such that $a^{(\log(a))^{1/K}}$ exists in M but $a^{\log(a)}$ does not.*

[1] This notion of limited use of an axiom also appears in Chapter V, Section 5, Subsection (g) of [9].

Proof: For the second part, let M be a non-standard model of true arithmetic, and let $a \in M$ be its non-standard element. Obviously, $a^{(\log(a))^{1/K}}$ exists in M since the function is total in true arithmetic. Let $N = \{n \in M : (\exists i \in \omega)(M \models n < a^{i(\log(a))^{1/K}})\}$. It is not hard to see that N is a cut of M that is closed under addition and multiplication. It follows that $N \models I\Delta_0$ (see Lemma 5.1.3 in page 64 of [12]). Finally, $a^{(\log(a))^{1/K}}$ still exists in N by absoluteness of the Δ_0-formula expressing the graph of exponentiation. However, $a^{\log(a)}$ does not exist in N because $a^{\log(a)} > a^{i(\log(a))^{1/K}}$ in M for every standard $i \in \omega$.

For the first part, suppose that $a^{(\log(a))^{1/K}}$ exists in $M \models I\Delta_0$. Our goal is to show that no Δ_0-definable function $F : M \to M$ maps $9a\log(a)$ injectively into $8a\log(a)$. The result would follow from Theorem 11 of [18] since then a prime exists in M between a and a^{11}. Let $b = a^2$ and observe that $b^{(\log(b))^{1/K}} = a^{2^{1+1/K}(\log(a))^{1/K}}$ exists in M since it is closed under multiplication. By Theorem 6, no Δ_0-definable function maps $2b$ injectively into b. It follows that no Δ_0-definable functions maps $\frac{9}{8}b$ injectively into b; otherwise we could compose that function with itself a constant number of times to maps $2b$ injectively into b. We conclude that no Δ_0-definable function maps $9a\log(a)$ injectively into $8a\log(a)$; otherwise, we could juxtapose that function with itself to obtain a Δ_0-definable function mapping $\frac{9}{8}b$ injectively into b (break b and $\frac{9}{8}b$ into $a/8\log(a)$ blocks of size $8a\log(a)$ and $9a\log(a)$ respectively). □

We note that $I\Delta_0$ proves $(\forall x)(\exists y)(y = x^{\log(x)^\epsilon}) \to (\forall x)(\exists y)(y = x^{\log(x)})$. However, the second part of Theorem 8 implies that $I\Delta_0$ does not prove $(\exists y)(y = x^{\log(x)^\epsilon}) \to (\exists y)(y = x^{\log(x)})$.

6 Discussion and Open Problems

Another major open problem in Feasible Number Theory is whether Fermat's Little Theorem is provable in $I\Delta_0$. Berarducci and Intrigila [2] point out that one important difficulty is that the modular exponentiation relation $x^y \equiv z \pmod n$ is not known to be Δ_0-definable. The situation has changed, however. Very recently, Hesse [11] proved that the modular exponentiation relation on numbers of $O(\log(n))$ bits is first-order definable. A well-known translational argument shows then that $x^y \equiv z \pmod n$ is Δ_0-definable. The proof of this result, however, seems to rely on Fermat's Little Theorem, and therefore it is not clear whether the basic properties of modular exponentiation are provable in $I\Delta_0$.

Open Problem 1 *Find a Δ_0 definition of the modular exponentiation relation whose basic properties are provable in $I\Delta_0$; namely, $x^y x^z \equiv x^{y+z} \pmod n$ and $(x^y)^z \equiv x^{yz} \pmod n$.*

We believe that a positive solution to this open problem would help developping the number theory of $I\Delta_0$ in the same way that the Δ_0-definition of the (non-modular) exponentiation relation helped developping the metamathematics of $I\Delta_0$ [8, 9].

Acknowledgments. I thank J. L. Balcázar and M. L. Bonet for helpful comments.

References

1. M. Ajtai. The complexity of the pigeonhole principle. In *29th Annual IEEE Symposium on Foundations of Computer Science*, pages 346–355, 1988.
2. A. Berarducci and B. Intrigila. Combinatorial principles in elementary number theory. *Annals of Pure and Applied Logic*, 55:35–50, 1991.
3. S. R. Buss. Polynomial size proofs of the propositional pigeonhole principle. *Journal of Symbolic Logic*, 52(4):916–927, 1997.
4. S. R. Buss and G. Turán. Resolution proofs of generalized pigeonhole principles. *Theoretical Computer Science*, 62:311–317, 1988.
5. S. Cook and R. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44:36–50, 1979.
6. L. Fortnow. Time-space tradeoffs for satisfiability. In *12th IEEE Conference in Computational Complexity*, pages 52–60, 1997. To appear in *Journal of Computer and System Sciences*.
7. L. Fortnow and D. van Meldebeek. Time-space tradeoffs for non-deterministic computation. In *15th IEEE Conference in Computational Complexity*, 2000.
8. H. Gaifman and C. Dimitracopoulos. Fragments of Peano's arithmetic and the MRDP theorem. In *Logic and algorithmic*, number 30 in Monographies de l'Enseignement Mathématique, pages 187–206. Univeristé de Genève, 1982.
9. P. Hájek and P. Pudlák. *Metamathematics of First-Order Arithmetic*. Springer, 1993.
10. A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.
11. W. Hesse. Division is in uniform TC^0. To appear in ICALP'01, 2001.
12. J. Krajícek. *Bounded arithmetic, propositional logic, and complexity theory*. Cambridge University Press, 1995.
13. J. Krajícek, P. Pudlák, and A. Woods. Exponential lower bound to the size of bounded depth frege proofs of the pigeon hole principle. *Random Structures and Algorithms*, 7(1):15–39, 1995.
14. R. J. Lipton and A. Viglas. On the complexity of SAT. In *40th Annual IEEE Symposium on Foundations of Computer Science*, pages 459–464, 1999.
15. A. Maciel, T. Pitassi, and A. R. Woods. A new proof of the weak pigeonhole principle. In *32th Annual ACM Symposium on the Theory of Computing*, 2000.
16. A. J. Macintyre and D. Marker. Primes and their residue rings in models of open induction. *Annals of Pure and Applied Logic*, 43(1):57–77, 1989.
17. V. A. Nepomnjaščij. Rudimentary predicates and Turing calculations. *Soviet Math. Dokl.*, 11:1462–1465, 1970.
18. 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(4):1235–1244, 1988.
19. T. Pitassi, P. Beame, and R. Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3(2):97–140, 1993.
20. G. Takeuti. *Proof Theory*. North-Holland, second edition, 1987.
21. A. R. Woods. *Some problems in logic and number theory, and their connections*. PhD thesis, Univerity of Manchester, Department of Mathematics, 1981.