Forcing in finite structures

(Revised version of ML–96–02)

Domenico Zambella[1]

Abstract. We present a simple and completely model-theoretical proof of a strengthening of a theorem of Ajtai's: the independence of the pigeonhole principle from $I\Delta_0(R)$. Qua strength, the theorem proved here corresponds to the complexity/proof-theoretical results of [10] and [14] but a different combinatorics is used. Techniques inspired by Razborov [11] replace those derived from Håstad [8]. This leads to a much shorter and very direct construction.

Mathematics Subject Classification. 03C62, 03F35 and 03H15.

Keywords. Models of atithmetic, bounded arithmetic, finite model theory.

1 Introduction

The (Δ_0-)pigeonhole principle is the assertion that there is no Δ_0-definable injective map from $[n{+}1)$ to $[n)$. Its provability in $I\Delta_0$ is one of the most famous open problems of bounded arithmetic. There is a second-order version of the same question. Can the pigeonhole principle be falsified in a second-order model of Σ^p_0-*comp*? (A theory that is, for $I\Delta_0$, what ACA_0 is for PA.) That is, is there a model of Σ^p_0-*comp* that contains set that is the graph of an injection from $[n{+}1)$ to $[n)$? This corresponds to asking whether the $\Delta(R)_0$ pigeonhole principle is provable in $I\Delta_0(R)$: the theory that has an induction axiom for all $\Delta(R)_0$ formulas, for R a new binary relation (in this form the question has been posed by Paris and Wilkie in [12]). The second-order formulation of the question is interesting not only because of its similarity with the first-order problem but also because Σ^p_0-*comp* is the bottom level of a hierarchy whose union is BA (i.e., $I\Delta_0{+}\Omega_1$).

The second question appears to be more tractable than the original problem. A solution has been given by Ajtai [2]. Subsequently, Krajíček, Pudlák and Woods

[1]Research supported by the Netherlands Foundation for Scientific Research (NWO) grant PGS 22-262

[10] and, independently, Pitassi, Beame and Impagliazzo [14] improved Ajtai's result. Ajtai's exhibited a model of $I\Delta_0(R)$ where R is a bijection between n and $n{+}1$. Ajtai's model contains only standard powers of n, while, using the results in [10] and [14], one can obtain that $\exp n^{\epsilon}$ exist for every infinitesimal rational ϵ (notation: "$\exp n$" and "2^n" are interchangeable). This is an optimal result since the pigeonhole principle between n and $n{+}1$ follows from the existence of $\exp n^{1/k}$ for k is any standard number [13]. This result implies also that assuming the totality of any Δ_0-definable function $g(x)$ such that all finite iterations of $g(x)$ are eventually majorized by 2^x does not help in proving the principle. In particular this holds for the functions $\omega_0(x) := x^2$ and $\omega_{n+1}(x) := \exp \omega_n(|x|)$ where $|x|$ is $\lceil \log(x{+}1) \rceil$.

Our proof differs from those in [2], [10] and [14] in many respects. We prove the result directly, dealing exclusively with formulas and with models of arithmetic. The proof in [2] makes a detour via Boolean circuits while in [10] and [14] the problem is translated in a purely complexity/proof-theoretical formalism. (The connection between bounded arithmetic and propositional proof systems has been first pointed out by [6] . See also [12] or, for a complete account: Krajíček's monograph [9].) Another difference involve the combinatorial methods used: [10] and [14] apply a variation of Håstad switching lemma that is proved with the probabilistic techniques developed in [8]. In place of these, we apply lighter methods inspired by some recent work of Razborov [11].

Recently, we obtained an unpublished note of Beame [5] where techniques inspired by [11] are applied to a variety of problems including those presented here. Our proof has been obtained independently of [5] before we could get access to it. Indeed, our exposition is different from [5] even in the very statement of the results: here we focus in a model-theoretical framework and avoid complexity/proof-theory. The referee of this paper has drawn to our attention that [9] uses techniques of [11] combined with forcing constructions.

Acknowledgments. I am grateful to Michiel van Lambalgen for his interest. His questions and remarks have contributed to a better presentation of the article. This paper elaborates on the material of a talk held in Warsaw during the meeting on Foundations of Arithmetic in April 1996. I would like to thank the organizers Zosia Adamowitcz and Roman Murawski for the invitation.

2 Preliminaries

We can choose to work either with relativized version of "strong" first-order theories such us $I\Delta_0(R)$ or over weak second-order theories such us Σ^p_0-*comp*. The difference is immaterial. We opt for this second form. The language consists of the symbols 0, 1, $+$, $\cdot$ $<$ and $\in$. These are interpreted as usual in second-order arithmetic. There are two sorts variables: first-order, $x, y, z,\ldots$ and second-order, $X, Y, Z,\ldots$ that are meant to range over numbers and, respectively, *finite* sets of numbers. A formula is[2] Σ^p_0 if all of its quantifiers are first-order and bounded, that is, they appear in the context: $(Qx{\in}X)\varphi$ or $(Qx{<}t)\varphi$ where Q is either $\forall$ or $\exists$ and t is a term in which x does not occur. In this paper we concentrate on the class Σ^p_0. This class is the ground level of a hierarchy of formulas, Σ^p_i, Π^p_i that is obtained by counting the alternations of second-order bounded quantifiers: $(QX{<}t)\varphi$ (where $X{<}y$ stands for: all elements of X are less than y). Note that terms are just polynomial on first-order variables. For $i{>}0$ this hierarchy coincides with the polynomial time hierarchy. In complexity theory Σ^p_0 corresponds to (a uniform version of) AC_0. The theory Σ^p_0-*comp* is axiomatized by

- the axioms of Robinson arithmetic,
- extensionality: $A{=}B \leftrightarrow \forall x\ (x{\in}A \leftrightarrow x{\in}B)$
- the least number principle: $A{\neq}\emptyset \rightarrow (\exists x{\in}A)(\forall y{<}x)(y{\notin}A)$,
- the axiom of finiteness: $\exists x\ (A{<}x)$,
- comprehension for all Σ^p_0-formulas: $(\exists X{<}a)(\forall x{<}a)\ [x{\in}X \ \leftrightarrow\ \varphi(x)]$

A side-note. This theory is the ground level of a hierarchy of theories that is obtained extending the schema of comprehension to Σ^p_i formulas for $i{>}0$ (see [17]). We will not consider higher levels of this hierarchy. Any techniques that can allow to handle the case of Σ^p_1 comprehension would be extremely interesting.

Observe that Σ^p_0-*comp* is a conservative extension of $I\Delta_0$. In fact every model of $I\Delta_0$ can be expanded to a model of Σ^p_0-*comp*.

The set of natural numbers together with the set of its finite subsets is the standard

[2]Add to justify notation. Semantically, Σ^p_0 coincides with the class of #-free $\Sigma^{1,b}_0$ formulas of [9] provided that in $\Sigma^{1,b}_0$ the second-order variables are restricted to range over the finite sets. Syntactically, $\Sigma^{1,b}_0$ formulas contains three more primitives. In our context these are unnatural and redundant.

model. Our construction is based on a countable elementary extension $\mathcal{M}$ of the standard model. This model $\mathcal{M}$ will be our ground model and it will stay fixed throughout this notes, the truth value of a formula is always evaluated in $\mathcal{M}$ (unless we specify differently). We shall expand an initial segment of $\mathcal{M}$ to a model $\mathcal{N}$ that contains the graph of a bijection G between $M, N \in \mathcal{M}$ that have cardinality $n+1$ and respectively n.

The language is expanded to include a name for every element of $\mathcal{M}$ (both fisrt- and second-order). We call these new constants parameters. The class of Σ_0^p-formulas is naturally extended to include formulas with parameters. (Parameters play an essential role below as the truth of the theorems do essentially depends on the size of these. Though, we will not diplay them in formulas unless really necessary.) A set of numbers is called *internal* when it belongs to $\mathcal{M}$. Otherwise, we will call it *external*. The graph of the bijection G that we will add to $\mathcal{M}$ is, clearly, an external object.

3 Antipasto

This section is added for expository reasons. Here we aim to present the combinatorial structure of the method in its simplest form. Still, the corollary below is of major importance because it implies a famous result on the undefinability of parity. Large parts of the proof of the lemma below will be taken over to the proof of the main lemma of next section.

1. We write $< n^{o(1)}$ and $< 2^{n^{o(1)}}$ when $< n^{\epsilon}$ and, respectively, $< 2^{n^{\epsilon}}$ holds for some infinitesimal rational ϵ. Fix a set N of non-standard cardinality n. In this section X and Y range over the subsets of N. A *condition* is the graph of a partial map P from the set N into $\{0,1\}$. The letters P, Q and R always range over conditions. We call $\|P\|$ (the cardinality of P), the *length* of the condition. We denote by $\mathcal{P}_p$ the set of conditions of length p and by $\mathcal{P}_{<p}$ the set of conditions of length $< p$. If the union of P and Q is still a condition, this is denoted by $Q^\frown P$. ∎

2. Let $\varphi(X)$ have no free variable but X. We write $P \Vdash \varphi(X)$, when $\varphi(X)$ holds for all X such that $C_X \supseteq P$ where C_X is the characteristic function of X. In words we say that P *forces* $\varphi(X)$. We say that the condition Q *decides* $\varphi(X)$ if either $Q \Vdash \varphi(X)$ or $Q \Vdash \neg\varphi(X)$. In the first case we say that Q *decides* $\varphi(X)$ *positively* in the second case Q *decides* $\varphi(X)$ *negatively*. We say that a set $\mathcal{D}$ of conditions *decides* $\varphi(X)$ if all conditions in $\mathcal{D}$ decide $\varphi(X)$. We write $\mathcal{D}^+$ for the conditions in $\mathcal{D}$ that decide $\varphi(X)$ positively and $\mathcal{D}^-$ for those that decide $\varphi(X)$ negatively. ∎

A side-note. We will show that for every Σ_0^p formula there is a condition P of length $< n$ that decides $\varphi(X)$. This is in itself an interesting result. It implies that no Σ_0^p formula counts modulo 2 the cardinality of X. In fact, suppose that $\varphi(X)$ counts modulo 2 the cardinality of X. Let P be a condition of length $< n$. Take any Y', Y'' compatible with P that differ from each other only for some $a \notin \mathrm{dom}(P)$. Then $\varphi(Y')$ iff $\neg\varphi(Y'')$. So, P does not decide $\varphi(X)$.

A side-note. The proof will proceed by induction on the complexity of φ. In general, the fact that for every x some (not too large) condition P_x decides $\varphi(x, X)$ is not sufficient to guarantee that some (not too large) condition P decides $(\forall x{<}a)\varphi(x, X)$. One needs some uniformity in the induction hypothesis. On the opposite side, to assume that one single condition decides all $\varphi(x, X)$ is easily seen to be false already for the atomic formula $x{\in}X$. One can weaken the notion of decidability as in the definition of s-decidability given below.

3. A set of conditions $\mathcal{D}$ is *complete* if

$$\{X \ : \ (\exists Q \in \mathcal{D})\ Q \subseteq C_X\}$$

is the set of all X. The set of all maps with domain A for some $A \subseteq N$ is a (natural but almost trivial) example of a complete set.

All notions concerning forcing can be relativized to the cone above some condition P. We say that $\mathcal{D}$ *decides* $\varphi(X)$ *above* P if $P^\frown Q$ decides $\varphi(X)$ for every $Q \in \mathcal{D}$ that is compatible with P. A set of conditions $\mathcal{D}$ is *complete above* P if

$$\{X \ : \ (\exists Q \in \mathcal{D})\ Q \subseteq C_X\} \quad \supseteq \quad \{X \ : \ P \subseteq C_X\}.$$

(The specification "above P" will be omitted whenever P is clear from the context.) We say that $\varphi(X)$ is *s-decidable* iff there is a complete set of conditions of length $< s$ that decides $\varphi(X)$. Roughly, to be s-decidable means that the truth value of $\varphi(X)$ only depends on a small subset of X of cardinality $< s$. ∎

A side-note. In the following lemma we isolate the induction step needed for the proof of the corollary below. We want a condition P that s-decides $(\forall x{<}a)\varphi(x, y, X)$

for all $y<b$ (where P and s are sufficiently small). It is crucial that P has to be the same for all $y<b$. We use a standard trick. We prove that for all y there are many P's that s-decides $(\forall x<a)\varphi(x,y,X)$. If the fraction of P's is sufficiently large compared with b we may immediately conclude that there is a P that works for all simultaneously $y<b$.

Warning. When we write that $\varphi(x,X)$ is s-decidable for all $x<a$, we intend "for every $x<a$ there is a complete set..".

Lemma 1 Fix t, s, l such that $t<s<l<n$ and let $a<2^{n^{o(1)}}$. Suppose that $\varphi(x,X)$ is t-decidable for all $x<a$. There is a set $\mathcal{S}\subseteq\mathcal{P}_{n-l}$ of conditions that s-decide $(\forall x<a)\varphi(x,X)$ such that[3]

$$\frac{\|\mathcal{S}\|}{\|\mathcal{P}_{n-l}\|} \quad\geq\quad 1-\left(\frac{2lt^2}{n-l}\right)^s.$$

Proof. For every x fix a complete set $\mathcal{D}_x\subseteq\mathcal{P}_{<t}$ that decides $\varphi(x,X)$. Let $P\in\mathcal{P}_{n-l}$ be given. We show how to construct a set $\mathcal{D}_P\subseteq\mathcal{P}_{<s}$ that decides $(\forall x<a)\varphi(x,X)$ above P. Afterwards we define $\mathcal{S}\subseteq\mathcal{P}_{n-l}$ so that for every condition P in $\mathcal{S}$ the set $\mathcal{D}_P$ is complete above P. In the second part we prove the claimed bound on the cardinality of $\mathcal{S}$.

4. The set $\mathcal{D}_P$ is defined setting $\mathcal{D}_P^-$ and $\mathcal{D}_P^+$ separately. The definition of $\mathcal{D}_P^-$ is straightforward:

$$\mathcal{D}_P^- \quad := \quad \bigcup_{x<a}\mathcal{D}_x^-.$$

The definition of $\mathcal{D}_P^+$ requires more work. To every condition P we associate a game between Prover and Verifier. Verifier has a deterministic behavior while Prover is free to guess among all possible strategies. The game is named P after the condition. Prover guess a condition R to be put in $\mathcal{D}_P^+$. At the start R is P and at each round Prover concatenates to R a new 1-condition. She tries to diagonalize against all conditions that force $\neg(\forall x<a)\varphi(x,X)$. Verifier provides evidence that her task is not accomplished yet: he presents some new condition in $\mathcal{D}_P^-$ that is compatible with what Prover has constructed so far. The game stops when one of the two players has achieved his/her goal. That is, either when Prover has been trapped in constructing

[3]This is not an optimal bound but it minimalizes the details of the proof.

a condition R in $\mathcal{D}_P^-$ or when Verifier can not find in $\mathcal{D}_P^-$ any condition compatible with R. In this second case, R forces $\neg(\forall x{<}a)\varphi(x,X)$, so, it can be enumerated in $\mathcal{D}_P^+$. If neither one of the two players can win after s moves, the play is draw. If no draws were possible, this could be the end of the story. In the second part of the proof we shall see that for most of the P this is actually the case. We give the formal definition of the game.

After i moves Prover has constructed the condition $P^\frown P_0^\frown\cdots^\frown P_{i-1}$ in replay to the conditions $V_0,\cdot\cdot,V_{i-1}$ produced by Verifier. If $i{=}s$ the play stops with a draw. If $i<s$, the play continues. Verifier plays the the least condition V_i in $\mathcal{D}_P^-$ that is compatible with $P^\frown P_0^\frown\cdots^\frown P_{i-1}$. If Verifier can not move the play stops: Prover wins. Otherwise, the play proceeds. Prover replays with *any* 1-condition P_i with $\mathrm{dom}(P_i)\subseteq\mathrm{dom}(V_i)$ that is compatible with $P^\frown P_0^\frown\cdots^\frown P_{i-1}$ but is incompatible with V_i. If Prover can not find such condition P_i, (this may happen iff V_i is a subcondition of $P^\frown P_0^\frown\cdots^\frown P_{i-1}$), the play stops: Prover has lost. This completes the description of the game. We say that a condition $P^\frown P_0^\frown\cdots^\frown P_{i-1}$ is *winning* if it is the condition that Prover has constructed in a winning play. Finally, we define

$$\mathcal{D}_P^+ \;:=\; \{R \;:\; \text{R is winning in the game } P\}. \qquad \blacksquare$$

5. Let $\mathcal{S}$ be the set of conditions P of length $n{-}l$ such that no draw is possible in the game P. It is clear that, for every P in $\mathcal{P}_{n-l}$, the set $\mathcal{D}_P^-$ forces $\neg(\forall x{<}a)\varphi(x,X)$ while $\mathcal{D}_P^+$ forces $(\forall x{<}a)\varphi(x,Y)$. In the next paragraph we check that, for every P in $\mathcal{S}$, the set $\mathcal{D}_P^-\cup\mathcal{D}_P^+$ is complete above P. $\blacksquare$

6. Take an arbitrary Y with characteristic function C_Y that extends P. Firstly, suppose $\neg(\forall x{<}a)\varphi(x,Y)$. Then $\neg\varphi(a,Y)$ for some a, so, some $Q\subseteq C_Y$ belongs to $\mathcal{D}_a^-\subseteq\mathcal{D}_P^-$. This condition Q forces $\neg(\forall x{<}a)\varphi(x,X)$ (the hypothesis $P\in\mathcal{S}$ is superfluous in this case). Secondly, suppose that $(\forall x{<}a)\varphi(x,Y)$. Consider the strategy where at each move Prover ensures that $P_i\subseteq C_Y$. Prover can loose only if she constructs a condition forcing $\neg(\forall x{<}a)\varphi(x,X)$ and this can not be the case with this strategy. Since there are no draws for $P\in\mathcal{S}$, Prover wins. The winning condition is compatible with Y and forces $(\forall x{<}a)\varphi(x,X)$. Then $\mathcal{D}_P^-\cup\mathcal{D}_P^+$ is complete. $\blacksquare$

7. We are left to prove the claim on the cardinality of $\mathcal{S}$. It suffices to define an injective map $\mathcal{F}$ from $\mathcal{P}_{n-l}\setminus\mathcal{S}$ into a set of cardinality

$$\leq \quad \|\mathcal{P}_{n-l}\|\cdot \left(\frac{2lt^2}{n-l}\right)^s$$

Fix a P that admits a draw. Let $V_0, \cdot\cdot, V_{s-1}$ and $P_0, \cdot\cdot, P_{s-1}$ be the moves of such a play. The the value of $\mathcal{F}(P)$ is defined in the next paragraphs. It is a triple. The first component of the triple is the condition $P^\frown V$. Ideally, $P^\frown V$ is the concatenation of P and $V_0, \cdot\cdot, V_{s-1}$ but, since these may be incompatible, we first erase the overlappings. Also, we cut off a part of the condition to obtain that $P^\frown V$ has length exactly s (this just to simplify counting). The second component of the triple $\mathcal{F}(P)$ is a string S that is meant to code $V_i \cap V$ modulo the knowledge of V_i. Finally, T (the third component of the triple) codes P_i modulo the knowledge of V_i. ∎

8. For $i < s$ let V_i^* be the condition

$$V_i^* \quad := \quad V_i \ \setminus \ V_{i\lceil \mathrm{dom}(P^\frown V_0^{*\frown}\cdots^\frown V_{i-1}^*)}$$

(in particular, V_0^* is $V_0 \setminus P$.) These conditions are mutually compatible. Let V be the subcondition $V_0^{*\frown}\cdots^\frown V_{s-1}^*$ obtained with the following procedure. Order the elements in $V_0^{*\frown}\cdots^\frown V_{s-1}^*$ so that the elements of V_i^* come before the elements of V_{i+1}^* and within a single V_i^* order the elements in an arbitrary way (e.g., by magnitude). Let V be the first s elements of $V_0^{*\frown}\cdots^\frown V_{s-1}^*$. We claim that $V_0^{*\frown}\cdots^\frown V_{s-1}^*$ has length $\geq s$ (so, the definition of V is legitimate). To prove the claim consider that V_i is compatible with $P^\frown P_0^\frown\cdots^\frown P_{i-1}$ but is incompatible with P_i, so, $P_0^\frown\cdots^\frown P_{s-1}$ has length s. For all $i < s$, $\mathrm{dom}(P_i)$ is contained in $\mathrm{dom}(V_i)$. Consequently $\bigcup_{i<s} \mathrm{dom}(V_i)$ has cardinality at least s, so, the claim is proved. The first component of $\mathcal{F}(P)$ is $P^\frown V$. ∎

9. Let S be a string of length $t{\cdot}s$ in the alphabet $\{0,1\}$ obtained concatenating $S_0, \cdot\cdot, S_{s-1}$. For $h < t$, the h-th entry of S_i is 1 if the h-th element of V_i exists and belongs to V, otherwise it is 0. The first component of $\mathcal{F}(P)$ is S. ∎

10. Finally, let T be a string of length s of numbers $< t$. The h-th entry of T is h if the h-th element of V_i is incompatible with P_i. (Since P_i has length 1 this element characterize P_i.) The third component of $\mathcal{F}(P)$ is T. ∎

11. Let $\mathcal{F}(P)$ be the triple consisting of the condition $P^\frown V$ the string S and the string T as defined above. We claim that $\mathcal{F}$ is one-to-one. To show this we show the invertibility of $\mathcal{F}$. Given a $P^\frown V$ and S, suppose we could reconstruct the first

j moves of the game: $V_0, \cdot\cdot, V_{j-1}$ and $P_0, \cdot\cdot, P_{j-1}$. Examining the first j blocks of t entries of S we can compute the condition $P^\frown[V \setminus V_0 \setminus, \cdots, \setminus V_{j-1}]$ (set difference associates to the left). Look for the minimal condition in $\mathcal{D}^-(P)$ that is compatible with

$$P^\frown P_0{}^\frown \cdots {}^\frown P_{j-1}{}^\frown[V \setminus V_0 \setminus, \cdots, \setminus V_{j-1}].$$

This condition is V_j. Now, looking at the j-th entry of T we reconstruct P_j. ∎

To find a bound on the cardinality of $\mathcal{P}_{n-l} \setminus \mathcal{S}$ it is sufficient to count the range of $\mathcal{F}$. There are

$$\binom{n}{l-s} 2^{n-l+s}$$

conditions of length $n-l+s$. The strings like S above have exactly s entries 1. In fact, to every 1 there correspond a different elements of V. By construction, these are at most s. So, at most

$$\binom{st}{s}$$

strings that are second component of a possible output of $\mathcal{F}$. Finally, there are t^s string that may appear as third component of an output of $\mathcal{F}$. Now, from straightforward approximations we get

$$\frac{\|\mathcal{P}_{n-l} \setminus \mathcal{S}\|}{\|\mathcal{P}_{n-l}\|} \quad \leq \quad \frac{\binom{n}{l-s} 2^{n-l+s} \binom{st}{s} t^s}{\binom{n}{l} 2^{n-l}} \quad \leq \quad \left(\frac{2lt^2}{n-l}\right)^s.$$

This yields the lemma. Q.E.D.

Corollary 2 Let N and n be as above. Let $\varphi(\bar{y}, X)$ be a Σ^p_0 formula of with parameters $< \exp n^{o(1)}$. Fix an $a < \exp n^{o(1)}$. For every standard rational $0<\nu<1$ there is an $s< n^{o(1)}$ and a condition P of length $< n - n^\nu$ that s-decides $\varphi(\bar{y}, X)$ for all $\bar{y} < a$.

Proof. The claim is obvious when φ is open. The inductive step for negation is trivial, so, it suffices to prove inductive step for the universal quantifier. Taking, if necessary, some larger $a < \exp n^{o(1)}$ we may assume that all quantifiers are of the form $\forall x{<}a$. (Indeed, recall the hypothesis $N < \exp n^{o(1)}$ that we have assumed in paragraph 1, this is essential to replace quantifiers of the form $\forall x{\in}X$ with $(\forall x{<}a)[x{\in}X \to \cdot\,]$.)

So, let φ have the form $(\forall x{<}a)\psi(x,\bar{y},X)$. As induction hypothesis we assume the existence of set $\mathcal{D}_{x,\bar{y}} \subseteq \mathcal{P}{<}t$ that decides $\psi(x,\bar{y},X)$ and is a complete above Q. The length of Q is $< n{-}n^{\nu}$ for some ν standard and $t < n^{o(1)}$. There is no loss of generality if we assume that Q is the empty condition. Otherwise, we can redefine

$$n' := n^{\nu} \qquad\qquad N' := N \setminus \mathrm{dom}(Q)$$

$$\psi'(x,\bar{y},X) := \psi(x,\bar{y},Q\cup X) \qquad \mathcal{D}'_{x,\bar{y}} := \{R \subseteq N'{\times}\{0,1\} \;\;:\;\; Q^\frown R \in \mathcal{D}_{x,\bar{y}}\}.$$

Observe that N' has cardinality n' and $N' < 2^{n'^{o(1)}}$. If we find a P' (with domain $\subseteq N'$) of length $n'{-}n'^{\nu}$ that s-decides $(\forall x{<}a)\psi'(x,\bar{y},X)$, then the condition $Q\cup P'$ s-decides $\varphi(\bar{y},X)$. The condition $Q\cup P'$ has length $n{-}n^{\nu\nu}$. The claim follows by the arbitrariness of ν.

So, assume that the condition Q above is the empty condition. From the lemma it follows that for every fixed $\bar{y} < a$ there are at least $1-2^{-s}$ conditions of length $n{-}n^{\nu}$ that s-decide $\varphi(\bar{y},X)$ (for any s such that $t{<}s{<}l$). Choose s large enough (but still $< n^{o(1)}$). By counting, there is a condition that s-decides all formulas $\varphi(\bar{y},X)$ (simultaneously for all $\bar{y} < a$). Q.E.D.

A bibliographical note. This Corollary implies that there is no Σ^p_0-formula that defines the parity of X. This has been proved by Ajtai [1] and independently by Furst, Saxe and Sipser [7] for formulas with parameters $< n$. Subsequently, the result has been extended by Yao [16] to formulas as in the theorem above. The setting in these articles in not quite the same as ours: in [1], the formalism used is that of finite model theory, while [7] and [16] prove a theorem on Boolean circuits complexity. In Boolean circuits complexity jargon Yao's theorem is called "exponential lower bound for parity". The proofs that appear in the most recent expositions make use of a so-called "switching lemma" and use probabilistic techniques developed in [8]. The switching lemma corresponds (in spirit) to the induction lemma proved above.

4 Pigeonhole principle

The construction of a model falsifying the pigeonhole principle is divided in two parts. The second part is a standard model theoretical argument of compactness. We shall consider an initial segment of $\mathcal{M}$, (those elements that are $< \exp n^{o(1)}$), add to it a new second-order object (the graph of a bijection falsifying the pigeonhole principle) and, finally, take the closure under Σ^p_0 definability. A forcing-like argument is used to ensure that the least number principle holds in the new model.

The first part is of combinatorial nature. It is similar to the proof presented in the previous section. Here X will be a bijection. There would not be much new if we considered X to be a bijection between sets of equal cardinality. But this is not the case we are most interested in. Since there is no internal bijection between sets of different cardinality, a definition of forcing similar to that of the previous section would quantify on the empty set. So, we have to allow X to be an external object. Fortunately, to have the proof working it is sufficient to require that X is locally internal. (A suitable notion of genericity would imply “local internal”. It is a matter of taste which notion is preferable.)

4.1 The forcing lemma

Some notation. Fix a set M of non-standard cardinality $n{+}1$ and a set N of cardinality n such that $M,\ N < \exp n^{o(1)}$. In this section X and Y range over the bijections with domain M and codomain N. We allow X to be an external object but we require that X is *locally internal*. That is, for every $s < n^{o(1)}$ and every formula $\varphi(x, X)$ with parameters $< \exp n^{o(1)}$ there is an internal sequence $a_0, \cdot\cdot, a_{s-1}$ such that, for all $i < s$,

$$(\exists x{<}n)\varphi(x, X_{\lceil\{a_0,\cdot\cdot,a_{i-1}\}}) \ \rightarrow\ \varphi(a_i, X_{\lceil\{a_0,\cdot\cdot,a_{i-1}\}})$$

It is routine to show that locally internal sets exists. No object considered in the sequel other than X and Y is external. Since X is an external object, the truth value of $\varphi(X)$, has to be computed in some expansion of the ground model $\mathcal{M}$, i.e., the variable X is considered as a fresh symbol for an unary predicate (on first-order objects) and is interpreted in some external subset of first-order elements of $\mathcal{M}$. In our notation we confuse X and its interpretation.

A side note. The requirement of being local internal becomes immaterial if we restrict it to standard s (every finite sequence is internal). A weaker version of the theorem below (where the claim that $\exp n^{\epsilon}\!\downarrow$ in $\mathcal{N}$ is dropped) could be proved woithout restrictions on X.

In the case considered here, it is necessary to restrict ourself to locally internal to have the proof of the lemma below work. This is not a loss of generality. We shall construct a model $\mathcal{N}$ that contains all elements of $\mathcal{M}$ that are $<\exp n^{o(1)}$ and containing a bijection G between sets of different cardinality. On its first-order part $\mathcal{N}$ coincides with $\mathcal{M}$. A modest amount of Σ^p_0 comprehension guarantees that any sequence like $a_0,\cdot\cdot,a_{s-1}$ above (with G for X) exists and is codable by numbers $<\exp n^{o(1)}$. So, this sequence need be internal. Hence, every set in $\mathcal{N}$ is necessarily locally internal.

In this section *conditions* are graphs of partial injections of M into N (but often we speak of them as functions). The letters P, Q and R denote always conditions. We say that P and Q are compatible if there is an injection extending both P and Q, i.e., if their set-theoretical union is still a condition. If P and Q are compatible we denote their union by $P^\frown Q$. We write $P\Vdash\varphi(X)$, when $\varphi(X)$ holds for all X such that $X\supseteq P$. The definitions of decidability and completeness given in Paragraphs 2 of the previous section are taken over to here. (Obviously, the interpretation of X, P and Q is the one in force now!).

Let us agree that in the statement and the proof of the following lemma *definable* means Σ^p_0 definable over parameters $<\exp n^{o(1)}$. Moreover, we say that a complete set $\mathcal{D}$ of conditions is definable if both $\mathcal{D}^+$ and $\mathcal{D}^-$ are definable.

Lemma 3 Fix t, s, and l such that $t<s<n^{o(1)}<l<n$. Let $a<2^{n^{o(1)}}$. Suppose we are given a definable sequence of complete sets $\mathcal{D}_x\subseteq\mathcal{P}_{<t}$ that decides $\varphi(x,X)$ for all $x<a$. Then there is a definable sets $\mathcal{S}\subseteq\mathcal{P}_{n-l}$ and $\mathcal{D}_P\subseteq\mathcal{P}_{<s}$ such that for every P in $\mathcal{S}$ the set $\mathcal{D}_P$ is complete above P and decides $(\forall x<a)\varphi(x,X)$. Moreover

$$\frac{\|\mathcal{S}\|}{\|\mathcal{P}_{n-l}\|}\quad\geq\quad 1\ -\ \left(\frac{(2tl)^2}{n-l}\right)^s.$$

If $\varphi(x,X)$ depends on parameters, then $\mathcal{D}_P$ and $\mathcal{S}$ are uniform on these, provided that also the sequence $\mathcal{D}_x$ is give uniformly.

Proof. For most of the P of length $n{-}l$ we can define a set $\mathcal{D}_P \subseteq \mathcal{P}_{<s}$ that is complete above P and that decides $(\forall x{<}a)\varphi(x,X)$ this will give the sets $\mathcal{S}$. Uniformity is straightforward, the reader can check it at the second reading.

The definition of $\mathcal{D}_P$ and $\mathcal{S}$ by means of a game are literally the same as in Paragraph 5 above. To prove the completeness of $\mathcal{D}_a^- \subseteq \mathcal{D}_P^-$, take an arbitrary $Y \supseteq P$. Firstly, suppose $\neg(\forall x{<}a)\varphi(x,Y)$. We reason exactly as in Paragraph 6: $\neg\varphi(a,Y)$ for some a, so, some $Q \subseteq Y$ belongs to $\mathcal{D}_a^- \subseteq \mathcal{D}_P^-$. This condition Q forces $\neg(\forall x{<}a)\varphi(x,X)$. Secondly, suppose that $(\forall x{<}a)\varphi(x,Y)$. The idea is the same as in Paragraph 6, however a few words need to be added to justify the claim of the definability of $\mathcal{D}_P$. At each move Prover ensures that $P_i \subseteq Y$. Apparently to define these moves one needs to quantify over the elements of $\mathcal{D}_P^-$ that is, over second-order objects. But $\mathcal{D}_P^- \subseteq \mathcal{P}_{<t}$ and $t < n^{o(1)}$. So, quantification over elements of $\mathcal{D}_P^-$ can be transformed into first-order quantification with bound 2^t. This guarantee the definability of Prover's strategy. The proof proceeds as above: Prover can loose only if she constructs a condition forcing $\neg(\forall x{<}a)\varphi(x,X)$ and this can not be the case with this strategy. Since there are no draws for $P \in \mathcal{S}$, Prover wins. The winning condition is compatible with Y and forces $(\forall x{<}a)\varphi(x,X)$.

The proof that for all P in $\mathcal{S}$ the set $\mathcal{D}_P$ is complete can be taken over here. For the second part of the argument in 5, one needs also to check that the strategy defined using Y is internal. Here we use that Y is locally internal and that $s < n^{o(1)}$.

Now it is left to prove the bound on the cardinality of $\mathcal{P}_{n-l} \setminus \mathcal{S}$ as claimed by the lemma. We define $\mathcal{F}$ on $\mathcal{P}_{n-l} \setminus \mathcal{S}$, show that here it is one-to-one and bound the cardinality of the co-domain. Exactly as in Paragraph 7 of the proof of Lemma 3, for P in $\mathcal{P}_{n-l} \setminus \mathcal{S}$, $\mathcal{F}(P)$ is a triple. The first component of the triple is defined as in Paragraphs 8 if we modify the definition of V_i^* (to guarantee compatibility we have to ensure that both the ranges and the domains do not overlap so, we take the two-sided restrictions)

$$V_i^* \;:=\; V_i \setminus \Big(\mathrm{dom}(P^\frown V_0^{*\frown}\cdots^\frown V_{i-1}^*)\times N \;\cup\; M\times \mathrm{range}(P^\frown V_0^{*\frown}\cdots^\frown V_{i-1}^*)\Big)$$

The second component can be defined literally as in Paragraph 9. The definition of the third component of $\mathcal{F}(P)$ differs of that reported in Paragraph 10. (The reason is clear: knowing that V_i differs of P_i on $a{\in}M$ is not sufficient to know the value that P_i assumes on a.) The condition P_i has length 1 and the domain of P_i is contained in V_i. So the domain of P_i may be coded by a number $<t$. It is clear that the range of P_i does not intersect that of P. We claim that the range of P_i does not intersect that of V_j^* for $j>i$. To prove the claim consider that, since $\mathrm{dom}(P_i) \subseteq \mathrm{dom}(V_i)$ then $\mathrm{dom}(P_i) \subseteq \mathrm{dom}(V_0^{*\frown}\cdots{}^\frown V_i^*)$. By construction, V_j is compatible with P_i, so, if the range of P_i intersect that of V_j then $P_i \subseteq V_j$. The claim follows. Consequently, the range of P_i can be coded (modulo the knowledge of $P^\frown V$ and of $V_0^{*\frown}\cdots{}^\frown V_i^*$) by a number $<l$. That is, a number $<s$ if it belongs to $\mathrm{range}(V_0^{*\frown}\cdots{}^\frown V_i^*)$, a number between s and l if it belongs to $P^\frown V$. So, let the third component of $\mathcal{F}(P)$ be a string T of length s whose i-th entry is a number $<tl$ that code domain and range of P_i up to the knowledge of $\mathrm{dom}(V_i)$ and $\mathrm{range}(P^\frown V)$.

A bound on the cardinality of $\mathcal{P}_{n-l} \setminus \mathcal{S}$ follows from a bound on the cardinality of the image of $\mathcal{P}_{n-l} \setminus \mathcal{S}$ under $\mathcal{F}$. There are

$$\binom{n}{l-s}\binom{n+d}{l+d-s}(n-l+s)!$$

conditions of length $n-l+s$. The second components of $\mathcal{F}(P)$ do not differs from Those of Lemma 3. Finally, there are $(tl)^s$ string that may appear as third component of an output of $\mathcal{F}$. Now, from straightforward approximations we get

$$\frac{\|\mathcal{P}_{n-l} \setminus \mathcal{S}\|}{\|\mathcal{P}_{n-l}\|} \quad \leq \quad \frac{\binom{n}{l-s}\binom{n+1}{l+1-s}(n-l+s)!\binom{st}{s}(tl)^s}{\binom{n}{l}\binom{n+1}{l+1}(n-l)!} \quad \leq \quad \left(\frac{(2tl)^2}{n-l}\right)^s.$$

This yields the lemma. Q.E.D.

Corollary 4 Let M and N be as above. Let $\varphi(\bar{y}, X)$ be a Σ_0^p formula with parameters $<\exp n^{o(1)}$. Fix an $a<\exp n^{o(1)}$. There is a condition P and a definable sequence of sets $\mathcal{D}_{\bar{y}} \subseteq \mathcal{P}_{<s}$ that decide $\varphi(\bar{y}, X)$ for all $\bar{y}<a$ and is complete above P. The condition P may be chosen so that $\|P\|<n-n^\nu$ for some standard rational $\nu>0$ and s can be found $<n^{o(1)}$.

Proof. The claim is obvious when φ is open. The inductive step for negation is trivial, so, it suffices to prove inductive step for the universal quantifier. It is easy to see that all universal quantifiers may be assumed to be of the form $(\forall x{<}a)$ where

a is some constant $< n^{o(1)}$. (Indeed, the hypothesis $N, M < \exp n^{o(1)}$ has been added to eliminate quantifiers like $\forall x{\in}X$.)

So, let $\varphi(\bar{y}, X)$ have the form $(\forall x{<}a)\psi(x, \bar{y}, X)$. As induction hypothesis we assume the existence of a set $\mathcal{D}_{x,\bar{y}} \subseteq \mathcal{P}_{<t}$ that is complete above Q and that decides $\psi(x, \bar{y}, X)$, where $t < n^{o(1)}$ and Q has length $< n{-}n^{\nu}$ for some ν standard. There is no loss of generality if we assume that the condition Q is empty. Otherwise, we can redefine

$$N' := N \setminus \mathrm{dom}(Q) \qquad M' := M \setminus \mathrm{range}(Q) \qquad n' := n^{\nu}$$

$$\psi'(x, \bar{y}, X) := \psi(x, \bar{y}, Q \cup X) \qquad \mathcal{D}'_{x,\bar{y}} := \{R \subseteq N' {\times} M' \ : \ Q^\frown R \in \mathcal{D}_{x,\bar{y}}\}.$$

M' has cardinality $n'^{\nu}{+}1$ and N' has cardinality n'^{ν}. If we find a P' (with domain $\subseteq N'$) of length $n'{-}n'^{\nu}$ that s-decides $(\forall x{<}a)\psi'(x, \bar{y}, X)$, then the condition $Q \cup P'$ has length $n{-}n^{\nu\nu}$ and s-decides $\varphi(\bar{y}, X)$.

So, assume that the condition Q above is empty and apply the lemma. This implies immediately that for every fixed $\bar{y} < a$ there are at least $1 - 2^{-s}$ conditions of length $n{-}n^{\nu}$ that s-decide $\varphi(\bar{y}, X)$ (for any s such that $t{<}s{<}l$). Choose s large enough (but still $< n^{o(1)}$). By counting, there is a condition that s-decides all formulas $\varphi(\bar{y}, X)$ (simultaneously for all $\bar{y} < a$). Q.E.D.

4.2 The construction of the model

At this point the proof proceeds with a standard argument of compactness. We state the main theorem.

Theorem 5 Fix $M, N \in \mathcal{M}$ of cardinality $n{+}1$ and n such that $M, N < 2^{n^{o(1)}}$. There is a model $\mathcal{N}$ of Σ^p_0-*comp* that contains all (first- and second-order) elements of $\mathcal{M}$ that are $< \exp n^{o(1)}$ and contains the graph of a bijection between M and N.

Proof. Let $\varphi_i(x, X)$ and a_i be some enumerations of the Σ^p_0-formula with parameters $< \exp n^{o(1)}$ and the first-order elements $< \exp n^{o(1)}$ of $\mathcal{M}$. We construct a chain $\{P_i\}_{i \in \omega}$ of conditions each of length at most $n{-}n^{\nu}$ for some standard rational $\nu{>}0$. Begin with $P_0 := \emptyset$. At stage $2i{+}1$ find a condition $P \supseteq P_{2i}$ that is of length $n{-}n^{\nu}$ and such that $\varphi_i(x, X)$ admits a definable complete set $\mathcal{D}_x \subseteq \mathcal{P}_{<s}$ above P for all

$x < a_i$. The theorem above guarantees the existence of this complete set for some $s < n^{o(1)}$. Now, find the minimal $\hat{x}$ such that $\mathcal{D}^+_{\hat{x}}$ is non empty. Let P_{2i+1} be any $P^\frown Q$ for Q in $\mathcal{D}^+_{\hat{x}}$. At even stages we prolong P_{2i+1} in order to ensure that $G := \bigcup_{i\in\omega} P_i$ is a bijection and that it is locally internal.

The model $\mathcal{N}$ contains all and only the first-order elements of $\mathcal{M}$ that are $< 2^{n^{o(1)}}$ and exactly all the sets of the form $\{x{<}a_i \ : \ \varphi_i(x, G)\}$. By construction, G is a bijection between M and N and $\mathcal{N}$ is a model of the schema of Σ^p_0-comprehension. Let us check that the least number principle is satisfied since all other axioms are evident. Every set in $\mathcal{N}$ is definable by some $\varphi_i(x, G)$, so, it suffices to have the least number principle for these formulas. At stage $2i{+}1$ we forced $\varphi_i(\hat{x}, X)$ true. For $x < \hat{x}$ with $\mathcal{D}^+_x$ is complete above P_{2i+1} so, $\varphi_i(\hat{x}, G)$ holds. Q.E.D.

References

[1] M.Ajtai, Σ^1_1 formulas on finite structures. *Annals of Pure and Applied Logic*, vol. 24, 1983, pp. 1–48

[2] M.Ajtai, The complexity of the pigeonhole principle. In the proceedings of the IEEE 29th Annual *Symposium on Foundations of Computer Science*, 1988, pp.346–355

[3] M.Ajtai, Parity and the pigeon-hole principle. In *Feasible mathematics*, eds. S.R.Buss and P.J.Scott, Birkhäuser, 1990, pp. 1–24

[4] M.Ajtai, The independence of the modulo p counting principles. In the proceedings of the 26th Annual *AMC Symposium on Theory of Computing*, 1985, pp. 402–411

[5] P.Beame, A switching lemma primer. Unpublished notes.

[6] S.A.Cook, Feasibly constructive proofs and the propositional calculus. In the proceedings of the 7th Annual *AMC Symposium on Theory of Computing*, 1975, pp. 83–97

[7] M.Furst, J.B.Saxe and M.Sipser, Parity circuits and the polynomial time hierarchy, In the proceedings of the IEEE 22nd Annual *Symposium on Foundations of Computer Science*, 1981 pp. 260–270

[8] J.Håstad, Almost optimal lower bounds for small depth circuits. In Randomness and Computation, ed. S.Micali, vol. 5 of *Advances in Computing Research*, JAI Press, 1989, pp. 143–170

[9] J.Krajíček, *Bounded Arithmetic, Propositional Logic and Complexity Theory.* Cambridge University Press. 1995.

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

[11] A.A.Razborov, Bounded arithmetic and lower bounds in Boolean complexity. *Feasible Mathematics II*, eds. P.Clote and J.Remmel, Birkhäuser, 1993 pp. 344–386

[12] J.B.Paris and A.J.Wilkie, Counting problems in bounded arithmetic. In the proceedings: Methods in Mathematical Logic, Caracas 1983, *Lecture notes in Mathematics*, vol. 1130, Springer Verlag, 1985, pp. 317–340.

[13] J.B.Paris and A.J.Wilkie, Counting Δ_0 sets. *Fundamenta Mathematicae*, vol. 127, 1987, pp. 67–76

[14] T.Pitassi, P.Beame and R.Impagliazzo, Exponential lower bound for the pigeonhole principle. In: *Computational Complexity*, vol. 3, 1993, pp. 97–140

[15] S.Riis, Count(p) does not imply Count(q). BRICS-preprint, RS-94-21, 1995, `http://www.brics.aau.dk/BRICS/`

[16] A.C.Yao, Separating the polynomial-time hierarchy by oracles. In the proceedings of the IEEE 26th Annual *Symposium on Foundations of Computer Science*, 1985, pp. 1–10

[17] D.Zambella, Notes on polynomially bounded arithmetic. *The Journal of Symbolic Logic*, vol. 61, 1996.

Domenico Zambella
Department of Mathematics and Computer Science, University of Amsterdam
Plantage Muidergracht 24, 1018 TV Amsterdam, NL
Email: domenico@fwi.uva.nl WWW: http://turing.fwi.uva.nl/~domenico/