Higher complexity search problems for bounded arithmetic and a formalized no-gap theorem

Neil Thapen*

April 27, 2011

Abstract

We give a new characterization of the strict $\forall\Sigma^b_j$ sentences provable using Σ^b_k induction, for $1 \leq j \leq k$. As a small application we show that, in a certain sense, Buss's witnessing theorem for strict Σ^b_k formulas already holds over the relatively weak theory PV.

We exhibit a combinatorial principle with the property that a lower bound for it in constant-depth Frege would imply that the narrow CNFs with short depth j Frege refutations form a strict hierarchy with j, and hence that the relativized bounded arithmetic hierarchy can be separated by a family of $\forall\Sigma^b_1$ sentences.

Keywords: bounded arithmetic, proof complexity, search problems
Mathematics subject classification: 03F30, 68Q15, 03F20

1 Introduction

Let L_{PV} be a language for arithmetic containing a function symbol for every polynomial time machine. We work over a universal base theory PV which fixes the basic properties of these symbols [11, 18]. Define a $\hat{\Sigma}^b_k$ (or *strict* Σ^b_k) formula to be a formula consisting of k or fewer alternating blocks of bounded quantifiers, with the first one existential, followed by a quantifier-free formula, where a bounded quantifier has the form $\forall x < t$ or $\exists x < t$ for t a term not containing x. We are interested in Buss's [3] hierarchy $(T^k_2)_{k \in N}$ of bounded arithmetic theories, which we may take to be defined as

$$T^k_2 := \mathrm{PV} + \hat{\Sigma}^b_k\text{-IND}$$

*Institute of Mathematics, Academy of Sciences of the Czech Republic, Žitná 25, CZ-115 67 Praha 1, `thapen@math.cas.cz`. Partially supported by institutional research plan AV0Z10190503 and grant IAA100190902 of GA AV ČR, grants LC505 (Eduard Čech Center) and 1M0545 (ITI) of MŠMT and by a grant from the John Templeton Foundation.

where Γ-IND stands for the usual induction axiom restricted to formulas from the class Γ.

Whether or not this hierarchy collapses to a finite level is a long-standing open question, closely connected to a similar question in complexity theory [18, 6, 28]. It is expected that it does not collapse, and that in fact the theories prove different $\forall\hat{\Sigma}^b_1$ (or even $\forall\hat{\Pi}^b_1$) sentences, by analogy with the behaviour of classical fragments of Peano arithmetic. If we expand the language so that induction hypotheses may contain new, undefined relation or (bounded) function symbols, it is known, using oracle separation results from complexity theory, that the hierarchy does not collapse. However it is still open whether this separation of "relativized" theories can be done using sentences of low, fixed complexity. The best general relativized separation that is known is that T^{i+1}_2 is not $\forall\hat{\Sigma}^b_{i+1}$ conservative over T^i_2 [8].

The $\forall\hat{\Sigma}^b_j$ sentences provable in T^k_2, for $j \leq k$, were first characterized in [17] in terms of reflection principles for systems of quantified propositional logic. Other characterizations of at least the provable $\forall\hat{\Sigma}^b_1$ sentences have appeared in [21, 12, 9, 13, 24, 20, 25, 1, 2]. In [26], building on [20], Alan Skelley and this author presented a simple, combinatorial characterization of the $\forall\hat{\Sigma}^b_1$ sentences provable in T^k_2 in terms of a *game induction principle* GI_k. In this paper we extend the work of [26] by three small results which make use of versions of the principle GI_k with higher quantifier complexity.

In Section 2, slightly generalizing a construction from [26], we define the *j-initial game induction principle* j-GI_k and show that it captures, in a strong way, the $\forall\hat{\Sigma}^b_j$ sentences provable in T^k_2 for all $1 \leq j \leq k$. Another recent characterization of these sentences appears in [1] and [2].

In Section 3 we use this characterization to give a strengthening of one direction of Buss's witnessing theorem for S^k_2 [3]. We show that any $\forall\hat{\Sigma}^b_k$ sentence provable in S^k_2 can be witnessed by a $\square^p_k$ function, provably in PV (although we have to be careful about how this is expressed, because we do not expect PV to be able to prove that this $\square^p_k$ function is total). Previously the witnessing was only known to be provable in T^{k-1}_2 [5].

In Section 4 we show how $\overline{\mathrm{GI}}_{|||a|||}(a)$, the (negated) game induction principle for $\log^{(3)}$-turn games, can be written as a narrow CNF and show that a superquasipolynomial lower bound for constant-depth refutations of $\overline{\mathrm{GI}}_{|||a|||}(a)$ would imply a separation between the narrow CNFs with short refutations in depth k and in depth $k+1$ Frege systems, for all $k \in \mathbb{N}$. Via a standard correspondence between first-order and propositional proofs ([22] or see e.g. [15]), this would imply a $\forall\hat{\Sigma}^b_1$ separation of the relativized bounded arithmetic hierarchy, as discussed above.

One reason this last result is interesting is that we *do* have a subexponential lower bound on constant-depth refutations of $\overline{\mathrm{GI}}_{|a|}(a)$. This follows from the lower bounds known for the pigeonhole principle PHP^a_{a-1} [19, 23], since the pigeonhole principle is reducible to $\mathrm{GI}_{|a|}(a)$ – the reduction is essentially by the construction in Section 2.3 of [26], which is based on the way counting can be done in Frege proof systems and in the theory U^1_1 [4, 15].

At the end of Section 4 we briefly discuss a possible approach to a low-level separation using $\mathrm{GI}_{|||a|||}(a)$, based on an idea from [16].

We will assume familiarity with [26] and will make heavy use of the notation, definitions and results from there.

In this paper we will say that a formula ϕ is a *Herbrandization* of a formula ψ if ϕ is obtained from ψ by replacing some or all of the existential quantifiers in ψ with explicit PV functions (we use this name because formulas of this kind arise from Herbrand's theorem for PV, and reserve *Skolemization* for the replacement of existential quantifiers with new, undefined function symbols). To make it easier to talk about long alternating sequences of quantifiers, we will often use three dots $\dots$ in a formula to stand for a sequence of finite length.

The author would like to thank Jan Krajíček and the anonymous referee for helpful comments on earlier versions of this paper.

2 The $\forall\hat{\Sigma}^b_j$ consequences of T^k_2

We first observe that a simple way to characterize the $\forall\hat{\Sigma}^b_{j+1}$ consequences of T^{k+j}_2 for $k \geq 1$, $j \geq 0$ would be to take the principle GI_k, which captures the $\forall\hat{\Sigma}^b_1$ consequences of T^k_2, and relativize everything to a complete $\hat{\Pi}^b_j$ oracle. Our construction in this section has a few advantages over this. One is that it is tidier, and in particular is built up out of games, effective strategies and game reductions that are polynomial time rather than $\square^p_{j+1}$. Others are that we get a stronger notion of reducibility, and that it works provably over PV.

Definition 1 *An instance of the* j-initial k-game induction principle j-GI_k *is given by size parameters a and b, a uniform sequence $G_0, \dots, G_{a-1}$ of polynomial time relations, a polynomial time function V and a uniform sequence $W_0, \dots, W_{a-2}$ of polynomial time functions.*

The instance $\mathrm{GI}_k(G, V, W, a, b)$ states that, interpreting $G_0, \dots, G_{a-1}$ as k-turn games in which all moves are bounded by b, the following cannot all be true:

1. *Deciding the winner of game G_0 depends only on the first j moves;*

2. *Player B can always win G_0 (expressed as a $\hat{\Pi}^b_j$ property);*

3. For $i = 0, \ldots, a-2$, W_i gives a game-reduction of G_{i+1} to G_i;

4. V is an explicit winning strategy for Player A in G_{a-1}.

The statement that this holds for all a and b can be written as a $\forall\hat{\Sigma}^b_j$ formula (see below). It is provable in T^k_2 by $\hat{\Pi}^b_k$-IND on i with the inductive hypothesis "Player B can always win G_i".

To explain the sense in which this captures the provable $\forall\hat{\Sigma}^b_j$ sentences of T^k_2, we will need a technical definition. We repeat a definition from [26], since it is used here in a slightly different way.

Definition 2 *A formula is $\tilde{\Sigma}^b_k$ if it consists of k bounded quantifiers, beginning with an existential quantifier and then strictly alternating in type. The bounds on the quantifiers may only contain free variables, not bound variables. $\tilde{\Pi}^b_k$ is defined dually.*

Any $\hat{\Sigma}^b_k$ formula Φ can be made into a $\tilde{\Sigma}^b_k$ formula Ψ by using pairing to combine adjacent quantifiers, finding a common bounding term, and possibly adding dummy quantifiers. Clearly Ψ is equivalent to Φ in a strong sense. In particular, witnessing results about Ψ can be transferred to Φ provably in PV.

We may write a sentence from j-GI$_k$ as a $\forall\tilde{\Sigma}^b_j$ sentence as follows:

$$\forall(a,b)\,\exists(w,x_1)\!<\!(ab^{2k},b)\,\forall x_2\!<\!b\,\exists x_3\!<\!b\,\ldots\,\mathrm{Q}x_j\!<\!b \qquad [\neg G_0(x_1,\ldots,x_j,0,\ldots,0)\vee\psi(w)]$$

where Q stands for $\exists$ if j is odd and $\forall$ if j is even, and where we are (rather informally) using a pairing function (u,v) to avoid repeated quantifiers. Here $\psi(w)$ stands for a PV formula expressing that w witnesses that condition 1, 3 or 4 from Definition 1 fails. Notice that, under the assumption that condition 1 holds, the formula

$$\exists x_1\!<\!b\,\forall x_2\!<\!b\,\exists x_3\!<\!b\,\ldots\,\mathrm{Q}x_j\!<\!b\,\neg G_0(x_1,\ldots,x_j,0,\ldots,0)$$

is equivalent to condition 2 failing.

Definition 3 *Let Φ and Ψ be $\forall\tilde{\Sigma}^b_k$ sentences respectively of the form*

$$\forall x\exists y_1\!<\!s_1\,\forall y_2\!<\!s_2\,\exists y_3\!<\!s_3\,\ldots\,\phi(x,\bar{y})$$

and

$$\forall u\exists v_1\!<\!t_1\,\forall v_2\!<\!t_2\,\exists v_3\!<\!t_3\,\ldots\,\psi(x,\bar{y}).$$

Then we say that Ψ is reducible to Φ *if there are* PV *functions* $f_0(u)$, $f_1(u, y_1)$, $f_2(u, y_1, v_2)$, *... such that* $f_i(u, y_1, v_2, \ldots, y_{i-1}, v_i) < s_i$ *for all even* $0 < i \leq k$, $f_i(u, y_1, v_2, \ldots, v_{i-1}, y_i) < t_i$ *for all odd* $i \leq k$, *and*

$$\phi(f_0(u), y_1, f_2(u, y_1, v_2), y_3, \ldots) \to \psi(u, f_1(u, y_1), v_2, f_3(u, y_1, v_2, y_3), \ldots).$$

For classes Γ and Δ of $\forall\tilde{\Sigma}^b_k$ sentences, we write $\Gamma \leq \Delta$ if every sentence in Γ is reducible to some sentence in Δ, and $\Gamma \equiv \Delta$ if this holds in both directions.

This is a natural extension to higher complexity classes of the definition of reducibility between NP search problems. Notice that it is a Herbrandization of a prenex form of the implication $\Phi \to \Psi$, and is essentially the same thing as what [26] calls a game-reduction between games represented by Ψ and Φ.

We can now state our characterization result.

Theorem 4 *For $k \geq 1$ and $1 \leq j \leq k$, $\forall\tilde{\Sigma}^b_j(T^k_2) \equiv j$-$\mathrm{GI}_k$, provably in* PV.

Proof For one direction, each sentence in j-GI_k is provable in T^k_2, so j-$\mathrm{GI}_k \subseteq \forall\tilde{\Sigma}^b_j(T^k_2)$. The other direction will follow from Lemmas 5 and 6 below. □

Lemma 5 *For all $k \geq 1$, $\forall\tilde{\Sigma}^b_1(T^k_2) \leq 1$-$\mathrm{GI}_k$, provably in* PV.

Proof This is by a straightforward reduction of GI_k to 1-GI_k. Suppose we have an instance of GI_k given by games $G_0, \ldots, G_{a-1}$, strategies U and V and game-reductions $W_0, \ldots, W_{a-2}$. This is reducible to an instance of 1-GI_k formed by adding an extra game G_{-1} at the start in which B wins every play, and using the strategy V to define a game-reduction W_{-1} of G_0 to G_{-1}. □

Lemma 6 *For $k \geq 0$ and $2 \leq j \leq k+2$, $\forall\tilde{\Sigma}^b_j(T^{k+2}_2) \leq j$-$\mathrm{GI}_{k+2}$, provably in* PV.

Proof The argument is essentially that of Theorems 4 and 5 of [26]. Suppose that $\forall u \Psi(u)$ is provable in T^{k+2}_2, for Ψ a $\tilde{\Sigma}^b_j$ formula. The first step is to replace Ψ with an equivalent (under reducibility) $\tilde{\Sigma}^b_j$ formula Φ of the form $\exists v_1 < t(u)\, \forall v_2 < t(u) \ldots \phi(u, \bar{v})$, where t is a term with u as its only free variable.

We have that $\forall u \Phi(u)$ is provable in T^{k+2}_2. Let $\Phi^\Delta(u)$ be the dual $\forall v_1 < t(u)\, \exists v_2 < t(u) \ldots \neg\phi(u, \bar{v})$ of $\Phi(u)$. By free-cut elimination (see e.g. [7]) there is a first order derivation, in the sequent calculus for T^{k+2}_2, of the sequent $\Phi^\Delta(u) \longrightarrow \emptyset$ in which every formula is of complexity $\tilde{\Pi}^b_{k+2}$ or lower. Hence by Theorem 21 of [26] there is a family of PK^0_k refutations, of size quasipolynomial in a, of the table of cedents $(\Phi^\Delta(a))^\circ + A$, where $(\Phi^\Delta(a))^\circ$ is the propositional translation of $\Phi^\Delta(a)$ and A is a sequence of true, polylogarithmic width "auxiliary" clauses. Furthermore these refutations are polynomial-time definable using the parameter a.

From now on we will write t for $t(a)$. By simple changes to the refutation, we may build a new, at most quasipolynomially larger, refutation which begins not with the initial cedents $(\Phi^\Delta(a))^\circ$ but with a slightly different translation of $\Phi^\Delta(a)$ into a table of propositional cedents, namely

$$(\{\langle \forall v_3 < t\, \exists v_4 < t \ldots \neg\phi(a, s_1, s_2, v_3, \ldots, v_j)\rangle : s_2 < t\})_{s_1 < t}.$$

If $j = k+1$ or $k+2$, this is the same thing as $(\Phi^\Delta(a))^\circ$. If $j \leq k$, then $(\Phi^\Delta(a))^\circ$ is the cedent consisting just of the symbol $\langle\Phi(a)\rangle$, and our version differs in that the formula is broken down so that we can access its structure.

So our refutation now starts with exactly the t initial cedents above, which we will call $B_0, \ldots, B_{t-1}$. These are followed by the auxiliary clauses and then the body of the refutation; we will call these two sets of cedents together $C_1, \ldots, C_e$, so that C_1 is the first auxiliary clause and C_e is the final, empty cedent of the refutation.

We can now define our instance of j-initial k-game induction. For convenience we will use a slightly different notation from the definition and call our first game G_{-1} rather than G_0, and our last game G_e. So the instance will consist of games $G_{-1}, \ldots, G_e$, a strategy V and reductions $W_{-1}, \ldots, W_{e-1}$.

The games $G_0, G_1, \ldots, G_e$ are defined from our refutation $B_0, \ldots, B_{t-1}$, $C_1, \ldots, C_e$ as in the proof of Theorem 4 of [26], except that this time we do not have games corresponding to the first $t-1$ cedents $B_0, \ldots, B_{t-2}$ but instead begin with B_{t-1}. So G_0 is a game which starts with player A choosing a cedent from $B_0, \ldots, B_{t-1}$ and claiming that all formulas in it are false; player B then picks a formula from the cedent and claims that it is true, and then the game continues as in [26]. For $1 \leq i \leq e$, G_i is a game which starts with player A choosing a cedent from $B_0, \ldots, B_{t-1}, C_1, \ldots, C_i$ and then proceeds in the same way.

The strategy V is the same as in the proof of Theorem 4 of [26].

We define the reductions $W_0, \ldots, W_{e-1}$ as follows: if C_{i+1} is an auxiliary clause, then the reduction W_i of G_{i+1} to G_i is trivial. This is because C_{i+1} is true and its literals can be listed in polynomial time, so if player A chooses C_{i+1} on the first turn of G_{i+1} then B can win on the second turn by naming the first true literal in C_{i+1}. If C_{i+1} is not an auxiliary clause then the reduction of G_{i+1} to G_i is exactly as in the proof of Theorem 4 of [26].

It remains to define the game G_{-1} and the reduction W_{-1} of G_0 to G_{-1}. Notice that game G_0 has the following structure:

1. Player A first names a cedent B_{r_1}, with $r_1 \leq t-1$, and claims all formulas in it are false. By construction, B_{r_1} has the form

$$\{\langle \forall v_3 < t \ldots \neg\phi(a, r_1, s_2, v_3, \ldots, v_j)\rangle : s_2 < t\},$$

   which is a translation of $\exists v_2 < t \, \forall v_3 < t \ldots \neg\phi(a, r_1, v_2, \ldots, v_j)$.

2. Player B then names a formula $r_2 \in B_{r_1}$, claiming it is true. By the structure of B_{r_1}, r_2 must be a formula of the form

$$\langle \forall v_3 < t \ldots \neg\phi(a, r_1, r_2', v_3, \ldots, v_j)\rangle$$

   for some $r_2' < t$. So choosing r_2 is equivalent to choosing a value r_2' for the variable v_2 in the formula $\exists v_2 < t \, \forall v_3 < t \ldots \neg\phi(a, r_1, v_2, \ldots, v_j)$.

3. Player A then names a conjunct r_3 of r_2, claiming it is false. This is equivalent to choosing a value r_3' for the variable v_3 in the formula $\forall v_3 < t \ldots \neg\phi(a, r_1, r_2', v_3, \ldots, v_j)$.

4. etc.

The game ends on the jth turn, in which one of the players must name some literal $\langle \neg\phi(a, r_1, r_2', \ldots, r_j')\rangle$, with B winning if the literal is true and A if it is false.

So we define the game G_{-1} as follows: if either player plays a move $\geq t$, that player loses immediately (this captures the bounds on the quantifiers in $\Phi^{\Delta}(a)$). Otherwise, after a finished play $v_1, \ldots, v_k$, player B wins if $\neg\phi(a, v_1, \ldots, v_j)$ and player A wins if $\phi(a, v_1, \ldots, v_j)$.

The games G_0 and G_{-1} are now essentially the same, and a reduction of G_0 to G_{-1} consists simply of a sequence of functions translating moves r_m in G_0 (naming cedents or subformulas in the propositional translation) to equivalent moves r_m' in G_{-1} (naming values to assign to the variables) and vice versa. Also notice that although they are both formally k turn games, only the first j moves play a role in deciding the winner.

The sentence of j-GI_k we have built has the following form, where q is a term in a coming from the size of the PK_k^0 refutation and $\psi(w)$ expresses that w is a witness that condition 1, 3 or 4 from Definition 1 fails:

$$\forall a \, \exists (w, v_1) < (q^{2k+1}, t) \, \forall v_2 < t \, \exists v_3 < t \ldots \mathrm{Q} v_j < t \\ [\neg G_{-1}(v_1, \ldots, v_j, 0, \ldots, 0) \vee \psi(w)].$$

Observe that $\neg G_{-1}(v_1, \ldots, v_j, 0, \ldots, 0)$ is just $\phi(a, v_1, \ldots, v_j)$. Furthermore this, together with the fact that the PK_k^0 refutation we used in our construction is well-formed, provably in PV, means that PV proves that $\psi(w)$ is always false. Therefore the sentence

$$\forall u \, \exists v_1 < t(u) \, \forall v_2 < t(u) \, \ldots \phi(u, \bar{v})$$

is reducible to the j-GI_k sentence written above, provably in PV, by a reduction in which all functions $f_0, \ldots, f_j$ are projections. $\Box$

3 A witnessing theorem

Theorem 4 can be seen as a kind of witnessing theorem, since in some sense it gives a mechanical way to witness a provable $\hat{\Sigma}_k^b$ sentence, by reducing it to an instance of game induction. Furthermore it works over the relatively weak theory PV. We can use this, together with the fact that k-GI_k can be witnessed by a $\Box_{k+1}^p$ machine using binary search, to give a strengthening of one direction of Buss's witnessing theorem about the $\forall\hat{\Sigma}_k^b$ consequences of S_2^k.

In its original form in [3], this was the following result: if ϕ is a $\hat{\Pi}_k^b$ formula and $S_2^{k+1} \vdash \forall x \, \exists y \, \phi(x, y)$, then there is a $\Box_{k+1}^p$ function f such that $\mathbb{N} \models \forall x \, \phi(x, f(x))$. In [5] Buss strengthened this by showing that, under the same assumptions, the sentence $\forall x \, \phi(x, f(x))$ is actually provable in T_2^k, for a natural way of formalizing the function f. We show below that, for the right choice of f, this witnessing is provable even in PV. However we do not show that PV proves that f is a total function (and we do not expect this to be provable); rather we prove in PV that, on any input x, if there is any correct computation w of f with output y, then $\phi(x, y)$.

For a $\Box_{k+1}^p$ machine M, that is, a polynomial time Turing machine with an oracle for a $\hat{\Sigma}_k^b$ formula $\exists x < t \, \Theta(q, x)$, where Θ is some complete $\hat{\Pi}_{k-1}^b$ formula, let $\mathrm{Comp}_M(x, y, w)$ express that w is a correct history of a computation of machine M on input x giving output y. In detail, it expresses that the initial configuration of the work tape contains x, that the final configuration contains y, that for each j, going from configuration j to configuration $j+1$ obeys the transition rules, and that oracle queries are replied to correctly as follows: for each pair of a query and reply q_j and r_j recorded in w, either r_j witnesses that the oracle answer is "yes" (r_j is a number in $[0, t)$ and $\Theta(q_j, r_j)$ is true) or r_j correctly records that the oracle answer is "no" ($r_j =$"no" and $\forall x < t \, \neg\Theta(q_j, x)$). In this way we can write $\mathrm{Comp}_M(x, y, w)$ as a $\hat{\Pi}_k^b$ formula.

Theorem 7 *For $k \geq 0$, suppose $S_2^{k+1} \vdash \forall u \, \exists v \, \chi(u,v)$, where χ is a $\hat{\Pi}_k^b$ formula. Then there is a $\square_{k+1}^p$ machine M such that*

$$\mathrm{PV} \vdash \forall u, v, w, \, \mathrm{Comp}_M(u,v,w) \to \chi(u,v).$$

Proof We may suppose $k \geq 1$, since the case $k = 0$ already follows from [5]. Suppose we have

$$S_2^{k+1} \vdash \forall u \, \exists v \, \forall z \, \phi(u,v,z)$$

for ϕ a $\hat{\Sigma}_{k-1}^b$ formula, which we assume contains some implicit bound t on the variable z. Then by the witnessing theorem of [5] there is a $\square_{k+1}^p$ machine P such that $\forall u \, \forall z \, \phi(u, P(u), z)$, provably in T_2^k. We may write this as

$$T_2^k \vdash \forall u, z, w, v, \, \neg \mathrm{Comp}_P(u,v,w) \vee \phi(u,v,z).$$

The right hand side is equivalent to a $\forall \hat{\Sigma}_k^b$ sentence and is provable in T_2^k, so by Theorem 4 it is reducible, provably in PV, to an instance of k-GI_k taking parameters u, z, w, v. Let us write this instance as a $\hat{\Sigma}_k^b$ sentence $\exists x \, H(u,z,w,v,x)$. We do not need the full strength of reducibility, but only the consequence that the existence of a solution x implies the above $\hat{\Sigma}_k^b$ formula. That is,

$$\mathrm{PV} \vdash \forall u, z, w, v \, [\exists x \, H(u,z,w,v,x) \to \neg \mathrm{Comp}_P(u,v,w) \vee \phi(u,v,z)]. \quad (*)$$

We will now describe a $\square_{k+1}^p$ machine Q that solves H, given the parameters as input. Furthermore, this will be provable in PV, in the sense that

$$\mathrm{PV} \vdash \forall u, z, w, v, x, s \, [\mathrm{Comp}_Q((u,z,w,v), x, s) \to H(u,z,w,v,x)].$$

Machine Q works as follows. It first makes the oracle query "can player B always win G_0?". If the answer is "no", then by binary search it looks for a witness that condition 2 of the definition of k-GI_k is false, that is, a first move x_1 for player A that puts A into a winning position in G_0. It then makes an oracle query to check that x_1 really has this property. Provably in PV, there are three possibilities: either x_1 is such a witness, or the final oracle reply was incorrect, or the binary search breaks down at some point because the oracle asserts that there is a witness within some interval but that there is no witness in either half of the interval (and so one of these three oracle replies must be incorrect). Note that induction for polynomial time predicates is enough to show that such a breakdown in the binary search exists, if neither of the first two possibilities holds.

Machine Q then queries "can player A always win G_{a-1}?". Again if the answer is "no", then it is easy to compute either a witness that condition 4 of the definition of k-GI_k is false, or an incorrect oracle reply, or a small set of inconsistent replies, provably in PV.

If both answers are "yes", then by binary search the machine finds i such that player B can always win G_i but player A can always win G_{i+1}, and from this computes a witness to condition 3 being false. Again PV is enough to prove that if this does not work, some oracle reply must be incorrect.

The machine M needed for the theorem now works as follows. On input u, it first simulates P, obtaining strings v and w for the output and computation of P. It then uses an oracle query to find out whether $\forall z < t\, \phi(u, v, z)$ (where t is the implicit bound on z in ϕ). If the answer is "yes", M halts and outputs v. If it is "no", M uses binary search to find a counterexample z, then simulates Q on input (u, z, w, v), then halts.

We claim that M witnesses $\forall u\, \exists v\, \forall z < t\, \phi(u, v, z)$, provably in PV. In a model of PV, let s be the history of a correct computation of M on some input u. First observe that s contains a correct computation w of P on input u, with output v. If the reply to the query "$\forall z < t\, \phi(u, v, z)$" was "yes" then, by correctness, the output v of M is the desired witness. If the reply was "no", then either there was an inconsistency in the binary search or s contains a computation of Q on (u, z, w, v) for some counterexample z. By correctness, this implies that s contains some ouput x of Q such that $H(u, z, w, v, x)$. But this now contradicts $(*)$, since we have $H(u, z, w, v, x)$, $\mathrm{Comp}_P(u, v, w)$ and $\neg\phi(u, v, z)$. $\square$

4 A uniform collapse

In [26] we strengthened the "no gap" theorem of [10] and showed in particular that if, in a relativized world, $T^k_2(\alpha) \vdash \mathrm{GI}_{k+1}(\alpha)$ for some $k \in \mathbb{N}$ then $T^k_2(\alpha) \vdash \mathrm{GI}_i(\alpha)$ for all $i \in \mathbb{N}$ with $i \geq k$. The purpose of this section is to show that the constructions used to show this result are uniform enough that it can be extended up to non-constant values of i. The argument is difficult to do in a purely first-order way since this would involve talking about formulas of non-standard quantifier depth, so instead we use a mixture of propositional and first-order logic, using bounded arithmetic as a tool to argue about families of propositional proofs. Unfortunately the presentation becomes rather technical, but the only really important thing happening is the analysis of the growth rate of the objects involved.

$\overline{\mathrm{GI}}_m(a)$ is a propositional contradiction, defined below. It is a straightforward translation of the first order sentence "GI_m fails for games, strategies and reductions G, U, V, W at a". We want $\overline{\mathrm{GI}}_m(a)$ to be a narrow CNF, that is, one in which every disjunction has size polynomial in $|a|$, so we will translate functions as bit-graphs rather than graphs.

For simplicity we will restrict ourselves to powers of 2 for a, so a is 2^n for some n. We also only consider $\overline{\mathrm{GI}}_m(a)$ for values of m less than $|a|$. The propositional variables in $\overline{\mathrm{GI}}_m(a)$ are then:

1. $G_{ix_1\ldots x_m}$ for all $i, x_1, \ldots, x_m < a$, expressing whether Player B wins game G_i with the play $x_1, \ldots, x_m$;

2. $U^r_{jx_1x_3\ldots x_{j-1}}$ for all even $1 \leq j \leq m$, all $x_1, x_3, \ldots, x_{j-1} < a$ and all $r < n$, expressing the rth bit of the move played at turn j by player B in strategy U, in response to player A playing $x_1, x_3, \ldots, x_{j-1}$ so far;

3. $V^r_{jx_2x_4\ldots x_{j-1}}$ for all odd $1 \leq j \leq m$, all $x_2, x_4, \ldots, x_{j-1} < a$ and all $r < n$, expressing the rth bit of the move played at turn j by player A in strategy V, in response to player B playing $x_2, x_4, \ldots, x_{j-1}$ so far;

4. $W^r_{ijz_1\ldots z_j}$ for all $i < a-1$, all $1 \leq j \leq m$ and all $z_1, \ldots, z_j$, expressing the rth bit of the jth function in the game-reduction W_i, on inputs $z_1, \ldots, z_j$.

We will call these respectively variables *in* G, U, V or W.

For readability, in the next definition we will write clauses as implications rather than disjunctions. For variables expressing the bit graphs of functions we will write, for example, $(U_{2x_1} = y)$ as shorthand for $\bigwedge_{r<n} U^r_{2x_1} = \delta_r$ where δ_r is 0 or 1 depending on the rth bit of y.

Definition 8 *For even m, $\overline{\mathrm{GI}}_m(a)$ is the CNF consisting of the following three groups of clauses.*

1. *For each $x_1, \ldots, x_m < a$, the clause*
$$(U_{2x_1} = x_2) \wedge (U_{4x_1x_3} = x_4) \wedge \ldots \wedge (U_{mx_1\ldots x_{m-1}} = x_m) \rightarrow G_{0x_1\ldots x_m}.$$
These express that U is a winning strategy for player B in G_0.

2. *For each $x_1, \ldots, x_m < a$, the clause*
$$(V_1 = x_1) \wedge (V_{3x_2} = x_3) \wedge \ldots \wedge (V_{(m-1)x_2\ldots x_{m-2}} = x_{m-1}) \rightarrow \neg G_{(a-1)x_1\ldots x_m}.$$
These express that V is a winning strategy for player A in G_{a-1}.

3. *For each $x_1, \ldots, x_m, y_1, \ldots, y_m < a$ and each $i < a-1$, the clause*
$$(W_{i1y_1} = x_1) \wedge (W_{i2y_1x_2} = y_2) \wedge \ldots \wedge (W_{imy_1x_2\ldots x_m} = y_m) \\ \wedge G_{ix_1\ldots x_m} \rightarrow G_{(i+1)y_1\ldots y_m}.$$
These express that W_i is a reduction of G_{i+1} to G_i.

For odd m the formula is similar, but the first two groups of clauses are changed to reflect that A now has the final move in all games, and the clauses in the third group become

$$(W_{i1y_1} = x_1) \wedge (W_{i2y_1x_2} = y_2) \wedge \ldots \wedge (W_{imy_1x_2\ldots y_m} = x_m) \\ \wedge G_{ix_1\ldots x_m} \rightarrow G_{(i+1)y_1\ldots y_m}.$$

Observe that there are no more than a^{2m+1} clauses and that the maximum size of a clause is $nm + 2$.

Definition 9 *$\overline{\mathrm{GI}}_{4,m}(a)$ is the set of cedents obtained by taking $\overline{\mathrm{GI}}_4(a)$ and replacing, for all $i, x_1, \ldots, x_4 < a$, each occurrence of the literal $G_{ix_1\ldots x_4}$ with the formula*

$$\bigwedge_{y_1} \bigvee_{y_2} \ldots G'_{ix_1\ldots x_4 y_1\ldots y_m}$$

and each occurrence of the literal $\neg G_{ix_1\ldots x_4}$ with the formula

$$\bigvee_{y_1} \bigwedge_{y_2} \ldots \neg G'_{ix_1\ldots x_4 y_1\ldots y_m}$$

where the connectives range over $[0, a)$ and we are using a new set of propositional variables $G'_{ix_1\ldots x_4 y_1\ldots y_m}$ for $i, x_1, \ldots, x_4, y_1, \ldots, y_m < a$. $\overline{\mathrm{GI}}_{4,m}(a)$ is a propositional contradiction, since $\overline{\mathrm{GI}}_4(a)$ is.

We need to argue about exponentially large (in $|a|$) propositional formulas, derivations and assignments. To do this, it is convenient to think of these things as coded by second-order objects (in the form of exponentially long strings of bits) and to allow second-order constants and variables to appear in our bounded arithmetic formulas. So long as we only use universal quantification over these variables, and avoid any second-order quantifiers in induction hypotheses, we may treat these new objects exactly like oracles (except that unlike oracles, they have a size bound). So from now on assume that our theories are relativized with as many oracles $X, Y, \ldots$ as we need, to be used in this way. We will also make use of one conventional oracle α, coding a family of PK^0_1 refutations (see the paragraph before Lemma 11). We will continue to write the theories as, for example, PV rather than $\mathrm{PV}(\alpha, X, Y, \ldots)$.

We will say that a second-order object Y is given by a polynomial time machine $A(\bar{X}, a, \bar{p})$, where a is a size parameter and $\bar{X}$ stands for a tuple of second-order variables or oracles, if there is a function $f(\bar{X}, a, \bar{p}, j)$ which takes the parameters $a, \bar{p}, j$ as inputs, has oracle access to $\bar{X}$, runs in time polynomial in $|a|$, and outputs the jth bit of Y.

Below, propositional formulas and derivations are formalized as in [26], except that the functions and relations involved will now sometimes be coded by second-order objects. The *size* of a propositional derivation means the size of the second-order object coding it; in particular this is a bound on both the number of cedents in the derivation and on the number of names for formulas occuring in it. Similarly the size of a CNF is a bound on the number of clauses and the number of literals in it.

Lemma 10 *$\overline{\mathrm{GI}}_{4,m}(a)$ is shortly derivable from $\overline{\mathrm{GI}}_{m+4}(a)$ in PK^0_{m+1} (with a natural renaming of variables from G' to G, which we will not say any more about). In fact, there is a polynomial time machine F such that provably in* PV*, for all a and all $m < |a|$, $F(a,m)$ is a PK^0_{m+1} derivation of $\overline{\mathrm{GI}}_{4,m}(a)$ from $\overline{\mathrm{GI}}_{m+4}(a)$ of size quasipolynomial in a.* $\square$

Suppose that for some $c \in \mathbb{N}$ there is a family of PK^0_1 refutations of $\overline{\mathrm{GI}}_4(a)$ of size $2^{|a|^c}$. Let $I(\alpha, a)$ be a machine that recovers a sequence of second-order objects that have been coded into an oracle α, and let T be the theory

$$\mathrm{PV} + \forall a[I(\alpha, a) \text{ is a } \mathrm{PK}^0_1 \text{ refutation of } \overline{\mathrm{GI}}_4(a) \text{ of size } 2^{|a|^c}].$$

We will not use the assumption about the existence of a family of refutations until the end of this section, but we are stating it now so that we have a suitable exponent c available for the definition of T.

Lemma 11 *There is a polynomial time machine A such that provably in T, for all a and all $m < |a|$, $A(\alpha, a, m)$ is a PK^0_{m+1} refutation of $\overline{\mathrm{GI}}_{m+4}(a)$ of size quasipolynomial in a.*

Proof Let Π be the quasipolynomial size PK^0_1 refutation of $\overline{\mathrm{GI}}_4(a)$ guaranteed to exist by T. The first step is to change Π into a PK^0_{m+1} refutation of $\overline{\mathrm{GI}}_{4,m}(a)$, as follows.

For each formula ϕ appearing in a cedent in Π, if ϕ is a literal in U, V or W, leave it unchanged. If ϕ is a literal of the form $G_{ix_1\dots x_4}$ or $\neg G_{ix_1\dots x_4}$, replace ϕ with a level m conjunction or disjunction respectively, as in Definition 9. Now suppose that ϕ is a conjunction of literals $l_1, \dots, l_m$. Replace ϕ with a level $m+1$ conjunction, defined as follows (recall that in a PK^0 proof all formulas in a conjunction must be disjunctions of the same level): if l_j is a literal in U, V or W, simply make l_j into a level m disjunction by padding. If l_j is a literal of the form $\neg G_{ix_1\dots x_4}$, replace l_j with the level m disjunction from Definition 9. If l_j is a literal of the form $G_{ix_1\dots x_4}$, replace l_j with the set of conjuncts

$$\{\bigvee_{y_2}\bigwedge_{y_3}\dots\{G'_{ix_1\dots x_4 y_1\dots y_m}\} : y_1 < a\},$$

where the curly brackets around $\{G'_{ix_1\ldots x_4 y_1\ldots y_m}\}$ are meant to indicate that the literal has been padded up by one level so that each formula in this set is a level m disjunction.

Call this new object Π'. Π' is something like a PK^0_{m+1} refutation of $\overline{\mathrm{GI}}_{4,m}(a)$, except that the cedents do not follow from each other by valid PK^0_{m+1} rules. But we can add in quasipolynomially many new cedents to make it a valid PK^0_{m+1} refutation. For example, a resolution step

$$\frac{\Gamma,\ G_{ix_1\ldots x_4} \qquad \Gamma,\ \neg G_{ix_1\ldots x_4}}{\Gamma}$$

in Π will turn into this in Π':

$$\frac{\Gamma,\ \bigwedge_{y_1}\bigvee_{y_2}\ldots G'_{ix_1\ldots x_4 y_1\ldots y_m} \qquad \Gamma,\ \bigvee_{y_1}\bigwedge_{y_2}\ldots \neg G'_{ix_1\ldots x_4 y_1\ldots y_m}}{\Gamma}.$$

This looks like an application of a Π_m-cut rule, which is not available in PK^0_{m+1}. However, using the methods of the proof of Theorem 21 of [26], we can simulate this rule in PK^0_{m+1} by adding at most $O(a^m)$ new cedents to Π'.

Our new refutation is defined locally in a simple way using the local properties of Π, and in particular can be defined in polynomial time from the oracle α and the parameters. We combine it with the derivation from Lemma 10 to get the desired refutation of $\overline{\mathrm{GI}}_{m+4}(a)$. □

Definition 12 *For $m < |a|$, $\overline{1\mathrm{-Ref}(\mathrm{PK}^0_m)}(a)$ is a propositional contradiction, of size quasipolynomial in a, expressing that there is a narrow CNF formula which is both satisfiable and refutable in PK^0_m. Formally, it has seven sets of propositional variables F, A, Q, R, S, T and f and states that*

1. *F codes a CNF of size $< a$ in which each clause has size at most $|a|$;*
2. *(Q, R, S, T, f) code a PK^0_m refutation of F, of size a;*
3. *A is a satisfying assignment to F.*

This is a propositional translation of the negation of the $1\mathrm{-Ref}(\mathrm{PK}^0_k)$ principle of [26], except that here we give explicit bounds to the size of the clauses and of the refutation in terms of a, so that we have one fixed quasipolynomial bound on the size of the propositional formula.

Lemma 13 *There is a polynomial time machine B such that provably in* PV*, for all a, all $m < |a|$ and all second-order objects X, if X is a satisfying assignment to $\overline{1\mathrm{-Ref}(\mathrm{PK}^0_m)}(a)$ then $B(X, a, m)$ is a satisfying assignment to $\overline{\mathrm{GI}}_{m+2}(a)$.*

Proof This is shown for constant $m \in \mathbb{N}$ in the proof of Theorem 4 of [26]. The same construction works for general $m < |a|$. $\square$

Lemma 14 *There is a polynomial time machine C and a constant $d \in \mathbb{N}$ such that provably in T, for all a, all $m < |a|$ and all second-order variables X, if X is a satisfying assignment to $\overline{\mathrm{GI}}_{m+4}(a)$ then $C(X, a, m)$ is a satisfying assignment to $\overline{\mathrm{GI}}_{m+3}(2^{|a|^d})$.*

Proof By Lemma 11 there is $d \in \mathbb{N}$ such that $A(\alpha, a, m)$ is a PK^0_{m+1} refutation of $\overline{\mathrm{GI}}_{m+4}(a)$ of size $2^{|a|^d}$. We also have a satisfying assignment X to $\overline{\mathrm{GI}}_{m+4}(a)$, and we may assume that $\overline{\mathrm{GI}}_{m+4}(a)$ is of size $< 2^{|a|^d}$ and that its clauses are of size $< |a|^d$. This is exactly what we need to define from α and X a satisfying assignment to $\overline{1\mathrm{-Ref}(\mathrm{PK}^0_{m+1})}(2^{|a|^d})$, and from this by Lemma 13 we can define a satisfying assignment to $\overline{\mathrm{GI}}_{m+3}(2^{|a|^d})$. $\square$

Recall that T^3_3 is the theory T^3_2 together with the axiom that $2^{2^{||x||^2}}$ exists for all x [3].

Lemma 15 *Provably in the theory*

$$T^3_3 + \forall a[I(\alpha, a) \textit{ is a } \mathrm{PK}^0_1 \textit{ refutation of } \overline{\mathrm{GI}}_4(a) \textit{ of size } 2^{|a|^c}],$$

for all a and all second-order X, X is not a satisfying assignment to $\overline{\mathrm{GI}}_{|||a|||}(a)$.

Proof We will write γ for $|||a|||$. Suppose X satisfies $\overline{\mathrm{GI}}_\gamma(a)$. Then we can apply Lemma 14 to get

$$C(X, a, \gamma - 4) \text{ satisfies } \overline{\mathrm{GI}}_{\gamma-1}(2^{|a|^d}),$$

and then again to get

$$C(C(X, a, \gamma - 4), 2^{|a|^d}, \gamma - 5) \text{ satisfies } \overline{\mathrm{GI}}_{\gamma-2}(2^{|a|^{d^2}}),$$

and so on. If we can formalize repeating this step $\gamma - 3$ times as an induction, we will have shown a contradiction, since GI_3 is provable in T^3_3.

Let $M = 2^{|a|^{d^\gamma}}$. Then M is a bound on the largest parameters we will need in the induction, and since $|a|^{d^\gamma} < |a|^{2^{d\gamma}} = |a|^{||a||^d}$, M is guaranteed to exist in T^3_3. Now let D be the machine which iterates C, that is, such that $D(X, a, 0) = X$ and $D(X, a, i+1) = C(D(X, a, i), 2^{|a|^{d^i}}, \gamma - 4 - i)$. We want to estimate the time bound on D.

Let $f(Y, b, m, j)$ be the polynomial time function, with time bound $|b|^e$ for $e \in \mathbb{N}$, which calculates the jth bit of $C(Y, b, m)$. In our induction the parameter b will always be less than M, so the maximum time to calculate

f is $|M|^e$. Calculating a bit of $D(X,a,i)$ requires calling f recursively, once for each node of a tree of depth i and fan-out $< |M|^e$, so for $i < \gamma$ we can bound the time taken by $|M|^{e\gamma} < |a|^{e\gamma ||a||^d} < |a|^{||a||^{d+1}}$. Hence the function to calculate bits of D is definable in our theory.

Therefore we can write our inductive hypothesis

$$D(X,a,i) \text{ satisfies } \overline{\mathrm{GI}}_{\gamma-i}(2^{|a|^{d^i}})$$

as a $\hat{\Pi}^b_1$ formula. Induction on i up to $\gamma - 3$ completes the proof. □

Theorem 16 *Suppose that for some $c \in \mathbb{N}$ there is a family of PK^0_1 refutation of $\overline{\mathrm{GI}}_4(a)$ of size $2^{|a|^c}$. Then for some $s \in \mathbb{N}$, there is a family of PK^0_1 refutations of $\overline{\mathrm{GI}}_{|||a|||}(a)$ of size $2^{2^{||a||^s}}$.*

Proof By Lemma 15 and Parikh's theorem, there is a term t (with a $2^{2^{||a||^{O(1)}}}$ growth rate) such that

$$\begin{aligned} T^3_3 \vdash \forall X, \forall b{<}t(a)\, (I(\alpha,b) \text{ is a } \mathrm{PK}^0_1 \text{ refutation of } \overline{\mathrm{GI}}_4(b) \text{ of size } 2^{|b|^c}) \\ \to (X \text{ is not a satisfying assignment to } \overline{\mathrm{GI}}_{|||a|||}(a)). \end{aligned}$$

Hence by doing some rearrangement and using the Paris-Wilkie translation of first-order into propositional proofs (in the form of Theorem 21 of [26]), for some $s \in \mathbb{N}$ there is a family π_a of $2^{2^{||a||^s}}$-size PK^0_1 refutations of the set of cedents $E_a \cup F_a$, where E_a is the propositional translation of

$$\forall b{<}t(a)\, (I(\alpha,b) \text{ is a } \mathrm{PK}^0_1 \text{ refutation of } \overline{\mathrm{GI}}_4(b) \text{ of size } 2^{|b|^c})$$

and F_a is the translation of

$$(X \text{ is a satisfying assignment to } \overline{\mathrm{GI}}_{|||a|||}(a))$$

(both of these are $\hat{\Pi}^b_1$). Here E_a has propositional atoms translating the bits of the oracle α. F_a has atoms translating the second-order variable X, and we may set up the translation so that it is isomorphic to $\overline{\mathrm{GI}}_{|||a|||}(a)$.

By the assumption that short PK^0_1 refutations of $\overline{\mathrm{GI}}_4(a)$ exist, we know that there is an assignment to the oracle α which satisfies E_a, for every a. Under this assignment each π_a collapses immediately to a refutation of F_a, and hence a refutation of $\overline{\mathrm{GI}}_{|||a|||}(a)$. □

Theorem 17 *Suppose that there is no size $2^{2^{||a||^{O(1)}}}$ constant-depth refutation of $\overline{\mathrm{GI}}_{|||a|||}(a)$. Then the narrow CNFs refutable in polynomial (or quasipolynomial) size and constant depth form a strict hierarchy with depth.*

In particular, for each $k \in \mathbb{N}$ the narrow CNF family $\overline{\mathrm{GI}}_{k+3}$ has polynomial-size refutations in PK_{k+1} but no quasipolynomial-size refutations in PK_k (or in $\mathrm{Res}(\log)$ in the case $k = 0$).

Proof Firstly, by the constructions in Theorem 21 of [26], any PK_k refutation of $\overline{\mathrm{GI}}_j(a)$ can be made into a PK_k^0 refutation that is at most quasipolynomially larger, and vice versa.

Secondly, in Theorem 16, $\overline{\mathrm{GI}}_4(a)$ and PK_1^0 could be replaced with $\overline{\mathrm{GI}}_{k+3}(a)$ and PK_k^0 for any constant $k \in \mathbb{N}$ greater than 1, and the same argument would still go through.

Finally, if there is a quasipolynomial-size Res(log) refutation of $\overline{\mathrm{GI}}_3(a)$ then by Theorem 8 of [26] there is a quasipolynomial-size PK_1^0 refutation of $\overline{\mathrm{GI}}_4(a)$, to which Theorem 16 applies. □

A possible approach to a low-level separation using $\mathrm{GI}_{|||a|||}(a)$ may come from a proposal in [16]. There, Krajíček defines the *isomorphism-chain principle*: let L be a first-order language and let Φ and Ψ be two Σ_1^1 L-sentences that cannot be satisfied simultaneously in any finite L-structure. Then for any numbers m, n and any chain $C_1, \ldots, C_m$ of finite L-structures with the universe $[n]$, it cannot be the case that $C_1 \models \Phi$, $C_m \models \Psi$, and C_i is isomorphic to C_{i+1} for each $i = 1, \ldots, m-1$. Krajíček poses the following question: if this principle has small constant-depth proofs, does it follow that there is a family of small constant-depth circuits that separate L-structures satisfying Φ from those satisfying Ψ? The intuition behind this is that the existence of such circuits would allow a very natural constant-depth proof of the principle, by induction along the chain.

If we take L to consist of a single k-ary relation defining a k-turn game, and take Φ to be "Player B has a winning strategy" and Ψ to be "Player A has a winning strategy", then such a restricted version of the isomorphism-chain principle becomes a special case of GI_k. We can alter the principle to deal with games with a non-constant number of turns by considering three-sorted structures with a number sort, an index sort and a sequence sort, each of an appropriate size (rather than a single-sorted structure on a universe $[n]$), and adding relations to the language and putting suitable axioms into Ψ and Φ to allow us to talk about indexed sequences of numbers. We take the language to include a relation G expressing whether a sequence of numbers is a win for A or B, and take Φ and Ψ to express that respectively B or A has a winning strategy, as above. Then, if there is a small constant-depth proof of $\mathrm{GI}_{|||a|||}(a)$, it follows that there is also such a proof of such an instance of chain-isomorphism. If the answer to the question posed in [16], suitably altered, is "yes", then this implies that there is a small constant-depth circuit which decides whether A or B has a winning strategy. This is impossible as these represent Sipser functions of non-constant depth [27, 14].

References

[1] A. Beckmann and S. Buss. Polynomial local search in the polynomial hierarchy and witnessing in fragments of bounded arithmetic. *Journal of Mathematical Logic*, 9(1):103–138, 2009.

[2] A. Beckmann and S. Buss. Characterizing definable search problems in bounded arithmetic via proof notations. In R. Schindler, editor, *Ways of Proof Theory*, pages 65–134. ontos verlag, 2010.

[3] S. Buss. *Bounded Arithmetic*. Bibliopolis, 1986.

[4] S. Buss. Polynomial size proofs of the propositional pigeonhole principle. *Journal of Symbolic Logic*, 52(4):916–927, 1987.

[5] S. Buss. Axiomatizations and conservation results for fragments of bounded arithmetic. In *Logic and Computation, Proceedings of a Workshop held at Carnegie Mellon University*, pages 57–84. AMS, 1990.

[6] S. Buss. Relating the bounded arithmetic and polynomial time hierarchies. *Annals of Pure and Applied Logic*, 75(1–2):67–77, 1995.

[7] S. Buss. Chapter 1: An introduction to proof theory & Chapter 2: First-order proof theory of arithmetic. In S. Buss, editor, *Handbook of Proof Theory*. Elsevier, 1998.

[8] S. Buss and J. Krajíček. An application of Boolean complexity to separation problems in bounded arithmetic. *Proceedings of the London Mathematical Society*, 69:1–21, 1994.

[9] M. Chiari and J. Krajíček. Witnessing functions in bounded arithmetic and search problems. *Journal of Symbolic Logic*, 63(3):1095–1115, 1998.

[10] M. Chiari and J. Krajíček. Lifting independence results in bounded arithmetic. *Archive for Mathematical Logic*, 38(2):123–138, 1999.

[11] S. Cook. Feasibly constructive proofs and the propositional calculus. *Proceedings of the 7th Annual ACM Symposium on Theory of computing*, pages 83–97, 1975.

[12] F. Ferreira. What are the $\forall\Sigma^b_1$-consequences of T^1_2 and T^2_2? *Annals of Pure and Applied Logic*, 75(1):79–88, 1995.

[13] J. Hanika. Herbrandizing search problems in bounded arithmetic. *Mathematical Logic Quarterly*, 50(6):577–586, 2004.

[14] J. Håstad. *Computational limitations for small-depth circuits.* MIT Press, 1987.

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

[16] J. Krajíček. A form of feasible interpolation for constant depth frege systems. *Journal of Symbolic Logic*, 75(2):774–784, 2010.

[17] J. Krajíček and P. Pudlák. Quantified propositional calculi and fragments of bounded arithmetic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 36(1):29–46, 1990.

[18] J. Krajíček, P. Pudlák, and G. Takeuti. Bounded arithmetic and the polynomial hierarchy. *Annals of Pure and Applied Logic*, 52:143–153, 1991.

[19] J. Krajíček, P. Pudlák, and A. Woods. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Structures and Algorithms*, 7(1):15–39, 1995.

[20] J. Krajíček, A. Skelley, and N. Thapen. NP search problems in low fragments of bounded arithmetic. *Journal of Symbolic Logic*, 72(2):649–672, 2007.

[21] J. Krajíček and G. Takeuti. On induction-free provability. *Annals of Mathematics and Artificial Intelligence*, 6:107–126, 1992.

[22] J. Paris and A. Wilkie. Counting problems in bounded arithmetic. In *Methods in Mathematical Logic*, number 1130 in Lecture Notes in Mathematics, pages 317–340. Springer, 1985.

[23] T. Pitassi, P. Beame, and R. Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational complexity*, 3:97–220, 1993.

[24] P. Pudlák. Consistency and games - in search of new combinatorial principles. In V. Stoltenberg-Hansen and J. Väänänen, editors, *Logic Colloquium '03*, number 24 in Lecture Notes in Logic, pages 244–281. ASL, 2006.

[25] P. Pudlák. Fragments of bounded arithmetic and the lengths of proofs. *Journal of Symbolic Logic*, 73(4):1389–1406, 2008.

[26] A. Skelley and N. Thapen. The provably total search problems of bounded arithmetic. *Proceedings of the London Mathematical Society*, 2011. doi:10.1112/plms/pdq044.

[27] A. Yao. Separating the polynomial-time hierarchy by oracles. In *Proc. 26th annual symposium on Foundations of computer science*, pages 1–10. IEEE Press, 1985.

[28] D. Zambella. Notes on polynomially bounded arithmetic. *Journal of Symbolic Logic*, 61(3):942–966, 1996.