An Unexpected Separation Result in Linearly Bounded Arithmetic

Arnold Beckmann[*]
Institute of Discrete Mathematics and Geometry
Vienna University of Technology

Jan Johannsen[†]
Institut für Informatik
Ludwig-Maximilians-Universität München

September 16, 2004

Abstract

The theories $S_1^i(\alpha)$ and $T_1^i(\alpha)$ are the analogues of Buss' relativized bounded arithmetic theories in the language where every term is bounded by a polynomial, and thus all definable functions grow linearly in length.

For every i, a $\Sigma_{i+1}^b(\alpha)$-formula $\mathrm{TOP}^i(a)$, which expresses a form of the total ordering principle, is exhibited that is provable in $S_1^{i+1}(\alpha)$, but unprovable in $T_1^i(\alpha)$. This is in contrast with the classical situation, where S_2^{i+1} is conservative over T_2^i w.r.t. Σ_{i+1}^b-sentences.

The independence results are proved by translations into propositional logic, and using lower bounds for corresponding propositional proof systems.

Keywords: Bounded Arithmetic; Propositional Proof Systems

MSC: Primary 03F20; Secondary 03F07, 68Q15, 68R99.

1 Introduction

Fragments of bounded arithmetic are logical theories that have a strong link to computational complexity classes. They are formulated in a first order language L_2 of arithmetic, whose non-logical symbols and their intended meaning are:

- 0, 1 (constants),

- $+, \cdot$ (addition and multiplication),

- $x \dot{-} y$ (arithmetical subtraction, $x \dot{-} y = \max(x - y, 0)$),

- $x \operatorname{div} y$, $x \operatorname{mod} y$ (integer division and remainder)

- $|x|$ (binary length, $|x| = \lceil \log_2(x + 1) \rceil$),

- $2^{\min(|x|, y)}$ (bounded exponentiation),

- $x \# y$ (smash, $x \# y = 2^{|x| \cdot |y|}$),

- $\leq, =$ (predicates less than or equal, and equality).

As usual, the bounded quantifiers $(\forall x \leq t)\varphi$ and $(\exists x \leq t)\varphi$ abbreviate $(\forall x)x \leq t \to \varphi$ and $(\exists x)x \leq t \wedge \varphi$, respectively. If the bounding term t is of the form $|s|$ for some term s, then the quantifier is called *sharply bounded*. A formula is (sharply) bounded if all quantifiers in it are (sharply) bounded. The set of sharply bounded formulas is denoted by Σ_0^b. Bounded formulas are classified into a hierarchy analogous to the arithmetical hierarchy as follows:

$$
\begin{array}{rcl}
\Sigma_1^b & : & \text{formulas of the form } (\exists x \leq s)(\forall y \leq |t|)\varphi, \\
& & \text{where } \varphi \text{ is a boolean combination of atomic formulas;} \\
\Sigma_{i+1}^b & : & \text{formulas of the form } (\exists x \leq s)\varphi \text{ with } \varphi \in \Pi_i^b; \\
\Pi_i^b & : & \text{prenex forms of negations of } \Sigma_i^b\text{-formulas.}
\end{array}
$$

This formulation of the formula classes Σ_i^b and Π_i^b consists only of so-called "strict" Σ_i^b-formulas. Buss [7] originally considered a different formulation where each class is also closed under sharply bounded quantification. For sake of readability, we just write Σ_i^b for the strict versions of the formula classes, since the original, more general classes will not appear in this work. The strict classes have been studied in several places, e.g. by Pollett [19] and Beckmann [3], and in particular by Impagliazzo and Krajíček [11], on which the present paper builds.

Let BASIC be a suitable set of quantifier-free axioms for the non-logical symbols (e.g. see [7]). The different theories of bounded arithmetic are specified by BASIC plus the amount of induction they are allowed to use. There are two ways of restricting induction: First, we can restrict the set of formulas for which induction is allowed. Second, we can weaken the formulation of induction. The usual schema of induction for formulas in Φ, denoted Φ-IND, is given by all formulas of the form

$$
\varphi(0) \wedge (\forall x < t)(\varphi(x) \to \varphi(x+1)) \to \varphi(t)
$$

for $\varphi \in \Phi$.

The schema of logarithmic induction, denoted Φ-LIND, is given by all formulas of the form

$$
\varphi(0) \wedge (\forall x < |t|)(\varphi(x) \to \varphi(x+1)) \to \varphi(|t|)
$$

for $\varphi \in \Phi$. As exponentiation is not provable total in bounded arithmetic theories, the logarithmic induction principle is potentially weaker than the usual induction principle.

The following theories in the language L_2 are defined by Buss [7]

$$
\begin{array}{rcl}
S_2^i & = & \text{BASIC} + \Sigma_i^b\text{-LIND} \\
T_2^i & = & \text{BASIC} + \Sigma_i^b\text{-IND} .
\end{array}
$$

These theories are related to computational complexity classes through the following notion: a number-theoretic function f is Σ_i^b-definable in a theory T, if there is a Σ_i^b-definition $\varphi_f(\vec{x}, y)$ of the graph of f such that T proves $(\forall \vec{x})(\exists y)\varphi_f(\vec{x}, y)$.

The relationship is established by the following classic result of Buss [7]:

Theorem 1. *For $i > 0$, the functions Σ_i^b-definable in S_2^i are exactly those in the class $FP^{\Sigma_{i-1}^P}$ of functions computable in polynomial time with an oracle for a set in Σ_{i-1}^P, the $(i-1)^{st}$ level of the polynomial time hierarchy.*

It is easy to see that $S_2^i \subseteq T_2^i \subseteq S_2^{i+1}$ for every $i > 0$. In addition we know that some of these extensions are conservative. If theory T is an extension of theory S,

then we call this extension $\forall \Sigma_i^{\mathrm{b}}$-conservative, written $S \preceq_{\forall \Sigma_i^{\mathrm{b}}} T$, iff S and T prove the same $\forall \Sigma_i^{\mathrm{b}}$-sentences. The following conservativity result was obtained by Buss [8]:

Theorem 2. *For $i \geq 1$, S_2^{i+1} is a $\forall \Sigma_{i+1}^{\mathrm{b}}$-conservative extension of T_2^i. With the right definition of T_2^0 (i.e. as Cook's universal theory PV_1) this also holds for $i = 0$.*

Hence, we have the following picture of the relation between the fragments of bounded arithmetic:

$$\mathrm{S}_2^1 \subseteq \mathrm{T}_2^1 \preceq_{\forall \Sigma_2^{\mathrm{b}}} \mathrm{S}_2^2 \subseteq \mathrm{T}_2^2 \preceq_{\forall \Sigma_3^{\mathrm{b}}} \mathrm{S}_2^3 \subseteq \dots$$

The major open problem concerning fragments of bounded arithmetic is to separate the theories, or equivalently, show that the union $\mathrm{S}_2 := \bigcup_i \mathrm{S}_2^i$ is not finitely axiomatizable. A conditional separation is given by the following result due to Krajíček et al. [14]:

Theorem 3. *If $\mathrm{T}_2^i = \mathrm{S}_2^{i+1}$, then the polynomial time hierarchy collapses to $\Sigma_{i+2}^P = \Pi_{i+2}^P$.*

Relativized theories

The language $L_2(\alpha)$ is obtained from L_2 by adding an unary predicate symbol α. The formula classes $\Sigma_i^{\mathrm{b}}(\alpha)$ and $\Pi_i^{\mathrm{b}}(\alpha)$ are defined exactly as Σ_i^{b} and Π_i^{b}, but in the language $L_2(\alpha)$.

Likewise, the relativized theories $\mathrm{S}_2^i(\alpha)$ and $\mathrm{T}_2^i(\alpha)$ are defined like S_2^i and T_2^i, only that the underlying language contains the predicate symbol α, which may also appear in induction axioms, in other words:

$$
\begin{aligned}
\mathrm{S}_2^i(\alpha) &= \mathrm{BASIC} + \Sigma_i^{\mathrm{b}}(\alpha)\text{-LIND} \\
\mathrm{T}_2^i(\alpha) &= \mathrm{BASIC} + \Sigma_i^{\mathrm{b}}(\alpha)\text{-IND} \ .
\end{aligned}
$$

All the results mentioned above also hold in the relativized case, in particular, for every $i \leq 1$, we have the relationships

$$\mathrm{S}_2^i(\alpha) \subseteq \mathrm{T}_2^i(\alpha) \preceq_{\forall \Sigma_{i+1}^{\mathrm{b}}(\alpha)} \mathrm{S}_2^{i+1}(\alpha) \ .$$

From the relativization of Theorem 3 above, and the existence of oracles that separate the levels of the polynomial time hierarchy (due to Yao [21]), we thus obtain that $\mathrm{T}_2^i(\alpha) \subsetneq \mathrm{S}_2^{i+1}(\alpha)$ for every $i \geq 1$. Due to the conservativity relation, the separating sentence must have quantifier complexity higher than $\Sigma_{i+1}^{\mathrm{b}}(\alpha)$. On the other hand, $\mathrm{S}_2^{i+1}(\alpha)$ is axiomatizable by a set of $\Sigma_0^{\mathrm{b}}(\Sigma_{i+1}^{\mathrm{b}}(\alpha))$-sentences, which gives a close upper bound on the complexity of the separating sentences. Thus, the relationship between $\mathrm{T}_2^i(\alpha)$ and $\mathrm{S}_2^{i+1}(\alpha)$ is fairly well understood.

The status of the inclusion between $\mathrm{S}_2^i(\alpha)$ and $\mathrm{T}_2^i(\alpha)$ is far less clear. The best known general result is due to Buss and Krajíček [9]:

Theorem 4. *For every $i \geq 1$, there is a $\forall \Sigma_i^{\mathrm{b}}(\alpha)$-sentence that separates $\mathrm{S}_2^i(\alpha)$ from $\mathrm{T}_2^i(\alpha)$.*

For the case $i = 2$, a little more is known: Chiari and Krajíček [10] have shown that there is a $\Sigma_1^{\mathrm{b}}(\alpha)$-sentence that separates $\mathrm{S}_2^2(\alpha)$ from $\mathrm{T}_2^2(\alpha)$. It is conjectured that the theories $\mathrm{S}_2^i(\alpha)$ and $\mathrm{T}_2^i(\alpha)$ for every $i \geq 2$ can be separated by sentences of lower complexity, at most by $\Sigma_{i-1}^{\mathrm{b}}(\alpha)$ sentences. Many researchers even believe in the truth of the following:

Conjecture 1. *For every $i \geq 1$, the theories $\mathrm{S}_2^i(\alpha)$ and $\mathrm{T}_2^i(\alpha)$ are separated by a $\forall \Sigma_0^{\mathrm{b}}(\alpha)$-sentence.*

The connection with proof complexity

One method to obtain separations of relativized bounded arithmetic theories arises from the connection between these and certain proof systems for propositional logic. There is a translation * of closed bounded $L_2(\alpha)$-formulas into propositional logic, which in the context of Bounded Arithmetic is called the *Paris-Wilkie-translation* [16]. It is a very natural translation which is well known in proof theory: a similar translation is used to translate first order formulas of arithmetic to an infinitary propositional language [18].

For each relativized theory T defined above, there is an associated propositional proof system P_T with the following property:

> For any bounded $L_2(\alpha)$-formula $\varphi(x)$: if $T \vdash (\forall x)\varphi(x)$, then the tautologies $\varphi(n)^*$ have P_T-proofs of quasi-polynomial size $2^{|n|^{O(1)}}$.

Thus, a lower bound for the proof system P_T implies an independence result for the theory T.

The proof systems P_T and the translation * will be defined in Section 3 below. The authors have elsewhere [5] given proofs of most of the known separations mentioned above following this methodology.

Linearly bounded theories

The language L_1 is defined as L_2 without the function symbol $\#$. Whereas all terms in L_2 are of polynomial length growth rate, $|t(n)| \leq |n|^{O(1)}$, the terms in L_1 are bounded by polynomials, and thus grow only linearly in length, i.e., $|t(n)| \leq O(|n|)$.

The theories S_1^i, T_1^i and the relativized theories $S_1^i(\alpha)$, $T_1^i(\alpha)$ are defined like the corresponding theories with lower index 2, only over the language L_1 instead of L_2.

Let us state some basic properties of linearly bounded arithmetic and some basic relationships between linearly bounded arithmetic theories. Using only properties from BASIC we can see that the formula classes Σ_i^b and Π_i^b are closed under conjunctions and disjunctions. Additionally, Σ_i^b is closed under bounded existential, and Π_i^b under bounded universal quantification.

In theories in a language that includes $\#$ that contain sufficiently strong induction, esp. in theories that comprise at least S_2^1, these classes are also provably closed under sharply bounded quantification. It is unknown whether closure under sharply bounded quantification holds for the classes Σ_i^b and Π_i^b in linearly bounded arithmetic theories.

The same proofs as given by Buss [7] also show that the linearly bounded arithmetic theories form an increasing hierarchy:

$$S_1^1 \subseteq T_1^1 \subseteq S_1^2 \subseteq T_1^2 \subseteq S_1^3 \subseteq \dots$$

The linearly bounded arithmetic theories are mainly studied for the reason that the relationships between them seem to reflect those between the corresponding L_2-theories, but independence results are often easier to obtain. This can be explained by the connection to proof complexity: an independence result for an L_1-theory generally requires smaller lower bounds than that for the corresponding L_2-theory. Thus, the independence of the pigeonhole principle

$$
\begin{aligned}
&(\forall x < a + 1)\,(\exists y < a)\,\alpha(x, y) \\
&\quad \rightarrow (\exists y < a)\,(\exists x_1 < x_2 < a + 1)\,\alpha(x_1, y) \wedge \alpha(x_2, y)
\end{aligned}
$$

from $S_1(\alpha)$ follows from Ajtai's [1] superpolynomial lower bound, whereas the independence from $S_2(\alpha)$ requires the larger lower bounds obtained later [15, 17].

A weak version of Conjecture 1 for the fragments of linearly bounded arithmetic was recently obtained by Impagliazzo and Krajíček [11]:

Theorem 5. *For every $i \geq 1$, there is a $\forall\Sigma_2^{\mathrm{b}}(\alpha)$-sentence that separates $\mathrm{T}_1^i(\alpha)$ from $\mathrm{T}_1^{i+1}(\alpha)$.*

Impagliazzo and Krajíček state that the stronger result with $\mathrm{S}_1^{i+1}(\alpha)$ instead of $\mathrm{T}_1^i(\alpha)$ would require the $\forall\Sigma_{i+1}^{\mathrm{b}}(\alpha)$-conservativity between these theories, which, in contrast to the case of $L_2(\alpha)$-theories, is not known.

In this paper we prove the surprising result that these conservativity relations do in fact not hold for the $L_1(\alpha)$-theories. This shows that the fragments of linearly bounded arithmetic can have different behavior than their polynomially bounded cousins w.r.t. conservativity relations.

Results

The motivation for our results stems from the first author's work on so-called dynamic ordinal analysis. It has been shown in [3] that adding smash functions of higher growth rate (i.e., the functions $\#_k$ for $k > 2$ given by $\#_2 = \#$ and $x \#_{k+1} y = 2^{|x| \#_k |y|}$) to bounded arithmetic theories results in longer chains of theories having all the same dynamic ordinal. It has been conjectured at the same place that therefore such theories are conservative over each other (see [3] for more explanations):

Conjecture 2 ([3]). *For $k \geq 4$ it holds,*

$$\mathrm{T}_k^1 \preceq_{\forall\Sigma_2^{\mathrm{b}}(L_k)} \mathrm{S}_k^2 \preceq_{\forall\Sigma_2^{\mathrm{b}}(L_k)} \cdots \preceq_{\forall\Sigma_2^{\mathrm{b}}(L_k)} \Sigma_k^{\mathrm{b}}\text{-}\mathrm{L}^{k-1}\mathrm{IND}(L_k)$$

The subscript k denotes that the smash functions $\#_2, \ldots, \#_k$ are present in the language.

This conjecture is proven true for $k = 2, 3$. Thus, driving this conjecture into the other direction, i.e. removing $\#_2$ from the language, resulted in the following conjecture:

Conjecture 3. $\mathrm{T}_1^1(\alpha) \not\preceq_{\forall\Sigma_2^{\mathrm{b}}(\alpha)} \mathrm{S}_1^2(\alpha)$.

Our main result is a proof of this Conjecture and extensions thereof, thereby obtaining the mentioned separation between fragments of linearly bounded arithmetic:

Theorem 6. *For $i \geq 1$, the theories $\mathrm{T}_1^i(\alpha)$ and $\mathrm{S}_1^{i+1}(\alpha)$ are separated by a $\forall\Sigma_{i+1}^{\mathrm{b}}(\alpha)$-sentence.*

The separating sentence will be a formulation of the Total Ordering Principle. The propositional proof complexity of it has been studied in several places [20, 6, 2], we will follow Beckmann and Buss [4]. We will also make use of the results obtained there on the proof-complexity of the Total Ordering Principle.

To formulate the Total Ordering Principle and also for later use we choose some form of sequence coding accessible in linearly bounded arithmetic. For every $k \in \mathbb{N}$ let $\langle a_1, \ldots, a_k \rangle_n$ be the L_1-term in the free variables $a_1, \ldots, a_k$ and n which has value $\sum_{i=1}^k n^{i-1} \cdot a_i$.

Let $[n]$ denote the set $\{0, \ldots, n-1\}$. The function $a_1, \ldots, a_k \mapsto \langle a_1, \ldots, a_k \rangle_n$ is a bijection between $[n]^k$ and $[n^k]$. It can be seen, using div and mod, that this property is already provable in S_1^1. Likewise, a code $\langle a_1, \ldots, a_k \rangle_n$ can be effectively decoded using the functions div and mod.

Fix a finite set $[a]$ and let $\prec$ be the binary relation given by $x \prec y \iff \alpha(\langle x, y \rangle_a)$. The *Total Ordering Principle* $\mathrm{TOP}(a, \prec)$ states that if $\prec$ is a total,

5

transitive and irreflexive relation on $[a]$, i.e., a total ordering, then $\prec$ has a minimal element on $[a]$:

$$
\begin{aligned}
(\forall x, y &< a)\,(x \prec y \vee x = y \vee y \prec x) \\
\wedge\ (\forall x, y, z &< a)\,(x \prec y \wedge y \prec z \to x \prec z) \\
\wedge\ (\forall x &< a)\,(\neg x \prec x) \\
\to\ (\exists x &< a)\,(\forall y < a)\,(\neg y \prec x)\ .
\end{aligned}
$$

The separating sentence between $\mathrm{T}_1^1(\alpha)$ and $\mathrm{S}_1^2(\alpha)$ is then given by the sentence $(\forall a)\,\mathrm{TOP}(|a|^3, \prec)$. The results for higher levels are obtained using the lifting techniques first described in [12] and improved in [4]. We will replace $\prec$ by a suitable Sipser function in α to obtain $\mathrm{TOP}^i(a, \alpha)$, and for the lower bounds we will utilize cut reduction by switching.

2 Upper bounds

In this section we will prove one part of Theorem 6, viz. that the Total Ordering Principles are provable in relativized linearly bounded arithmetic. The next proposition shows this for the base case.

Proposition 7. *The sentences $(\forall a)\,\mathrm{TOP}(|a|^k, \prec)$ are consequences of $\mathrm{S}_1^2(\alpha)$, for all $k \in \mathbb{N}$.*

A first idea to prove this is by induction on x in $\mathrm{TOP}(x, \prec)$. But as each term $t(a)$ in $\mathcal{L}_1$ has linear growth rate, one logarithmic induction can only access a part linear in $|a|$. Thus, we will need k nested logarithmic inductions to reach $|a|^k$.

Let $n = |a|$ and assume that $\prec$ is a total ordering on $[n^k]$. From now on we identify $[n^k]$ and $[n]^k$ which we are allowed using the previously defined effective coding and decoding functions $(\mathrm{mod}\,n)$. Hence we view $\prec$ as an ordering on $[n]^k$.

Let $\vec{x} \times [b] \times [n]^{k-i}$ denote the $(k-i+1)$-dimensional cylinder given by

$$
\left\{ \langle \vec{x}, y, \vec{z} \rangle_n \,;\, y \in [b], \vec{z} \in [n]^{k-i} \right\}\ .
$$

The intermediate induction assertions express that for $\vec{x} \in [n]^{i-1}$ and $0 < b \leq n$ there exists a $\prec$-minimal element in the cylinder $\vec{x} \times [b] \times [n]^{k-i}$. They will be shown shown by (main) induction on $i = k, \ldots, 1$ (on the "meta" level) and (side) induction on $b = 1, \ldots, n$ (inside $\mathrm{S}_1^2(\alpha)$). Formally, the intermediate induction assertions are given by the formulas $M^i(n, x_1, \ldots, x_{i-1}, b)$ defined as

$$
\begin{aligned}
(\exists z_i \in [b])\,(\exists z_{i+1}, \ldots, z_k \in [n])\,(\forall y_i \in [b])\,(\forall y_{i+1}, \ldots, y_k \in [n]) \\
(\neg\,\langle x_1, \ldots, x_{i-1}, y_i, y_{i+1}, \ldots, y_k \rangle_n \prec \langle x_1, \ldots, x_{i-1}, z_i, z_{i+1}, \ldots, z_k \rangle_n)
\end{aligned}
$$

By coding succeeding quantifiers of the same type into one we obtain that this is equivalent to a $\Sigma_2^{\mathrm{b}}(\alpha)$-formula. Hence we can use the formulas M^i as inductive assertions in $\mathrm{S}_1^2(\alpha)$.

We now argue inside $\mathrm{S}_1^2(\alpha)$ to prove the above Proposition. As said before, let $n = |a|$ and assume that $\prec$ is total ordering on $[n^k]$. We will show by meta induction on $i = k, \ldots, 1$ that the formulas

$$
(\forall x_1, \ldots, x_{i-1} \in [n]) \quad M^i(n, x_1, \ldots, x_{i-1}, n) \tag{1}
$$

are consequences of $\mathrm{S}_1^2(\alpha)$. For $i = 1$ this implies that $\prec$ has a minimal element on $[n]^k$, which proves the assertion.

Let us now consider (1) for $i = k$, which serves as the induction base of our meta induction. Let $\vec{x}$ denote $x_1, \ldots, x_k$, then we prove $M^k(n, \vec{x}, n)$ by logarithmic induction on $b = 1, \ldots, n$ in $M^k(n, \vec{x}, b)$ inside $\mathrm{S}_1^2(\alpha)$. $M^k(n, \vec{x}, 1)$ holds as it

expresses that there is a $\prec$-minimal element in the cylinder $\vec{x} \times [1]$. This cylinder consists only of one element $\langle \vec{x}, 0 \rangle_n$, which is $\prec$-minimal on the cylinder as $\prec$ is irreflexive by assumption.

For the induction step from b to $b+1$ with $b+1 \in [n]$ we have by induction hypothesis $M^k(n, \vec{x}, b)$. I.e., there is some $a \in [b]$ such that $\langle \vec{x}, a \rangle_n$ is $\prec$-minimal on the cylinder $\vec{x} \times [b]$. We then find a minimal element on the cylinder $\vec{x} \times [b+1] = (\vec{x} \times [b]) \cup \{\langle \vec{x}, b \rangle_n\}$ simply by comparing $\langle \vec{x}, a \rangle_n$ with $\langle \vec{x}, b \rangle_n$. If $\neg \langle \vec{x}, b \rangle_n \prec \langle \vec{x}, a \rangle_n$, then $\langle \vec{x}, a \rangle_n$ is also $\prec$-minimal on the cylinder $\vec{x} \times [b+1]$, otherwise $\langle \vec{x}, b \rangle_n$ does the job. Hence $M^k(n, \vec{x}, b+1)$ follows. In turn (1) follows for $i = k$.

For the meta-induction step from $i+1$ to i we can inductively assume that (1) holds for $i+1$. I.e., letting $\vec{x}$ denote $x_1, \ldots, x_{i-1}$, we have

$$(\forall \vec{x}, x \in [n]) \quad M^{i+1}(n, \vec{x}, x, n) \tag{2}$$

Fix $\vec{x} \in [n]$. We show $M^i(n, \vec{x}, n)$ by logarithmic induction on $b = 1, \ldots, n$ in $M^i(n, \vec{x}, b)$ inside $S_1^2(\alpha)$. $M^i(n, \vec{x}, 1)$ expresses that there is a $\prec$-minimal element on

$$\vec{x} \times [1] \times [n]^{k-i} \quad = \quad \vec{x}, 0 \times [n] \times [n]^{k-(i+1)} .$$

Hence $M^i(n, \vec{x}, 1)$ is equivalent to $M^{i+1}(n, \vec{x}, 0, n)$, and the latter follows from our meta induction hypothesis (2). This serves as the induction base.

For the induction step from b to $b+1$ with $b+1 \in [n]$ we have by induction hypothesis $M^i(n, \vec{x}, b)$. I.e. there are $a \in [b]$ and $\vec{a} \in [n]^{k-i}$ such that $\langle \vec{x}, a, \vec{a} \rangle_n$ is $\prec$-minimal on the cylinder $\vec{x} \times [b] \times [n]^{k-i}$. Let A denote $\langle \vec{x}, a, \vec{a} \rangle_n$. By (2) we have $M^{i+1}(n, \vec{x}, b, n)$, i.e. there exists $\vec{c} \in [n]^{k-i}$ such that $\langle \vec{x}, b, \vec{c} \rangle_n$ (let us denote this by B) is $\prec$-minimal on the cylinder $\vec{x}, b \times [n]^{k-i}$. Now we compare A with B to find a $\prec$-minimal element on the cylinder $\vec{x} \times [b+1] \times [n]^{k-i}$.

If $\neg B \prec A$, then by totality $A \prec B$, and hence A is $\prec$-minimal on $\vec{x} \times [b+1] \times [n]^{k-i}$. In the other case, B is the $\prec$-minimal element.

Altogether we have shown $M^i(n, \vec{x}, b+1)$ and the assertion follows inductively. This finishes the proof of Proposition 7.

We now describe how the base case can be lifted to all levels of linearly bounded arithmetic. In particular we define $\mathrm{TOP}^d(a, \alpha)$ for $d > 0$.

For $d \in \mathbb{N}$, the Sipser function $S_{d,a,\alpha}(x, y)$ is defined by the $\Sigma_d^b(\alpha)$-formula

$$(\exists z_1 \leq a) (\forall z_2 \leq a) \ldots (Q^d z_d \leq a) \, \alpha(\langle x, y, z_1, \ldots, z_d \rangle_a)$$

where Q^d is either $\exists$ or $\forall$, depending on whether d is odd or even, respectively. The separating formula for higher levels of relativized linearly bounded arithmetic is then defined by substituting an appropriate Sipser function for $\prec$ in $\mathrm{TOP}(a, \prec)$. I.e., for $d > 0$ let $\mathrm{TOP}^d(a, \alpha)$ be defined by $\mathrm{TOP}(a, S_{d,a,\alpha})$, that is $\mathrm{TOP}(a, \prec)$ in which every occurrence of the formula $u \prec v$ is replaced by $S_{d,a,\alpha}(u, v)$.

The formula $\mathrm{TOP}^d(a, \alpha)$ is equivalent to a $\Sigma_{d+1}^b(\alpha)$-formula in the theory BASIC. Using the proof of Proposition 7 we see that $\mathrm{TOP}^d(a, \alpha)$ is provable in $S_1^{d+1}(\alpha)$. This is true, because the Sipserized intermediate induction assertions from that proof are of the form

$$(\exists z \in [b]) \, (\exists \vec{z} \in [n]^{k-i}) \, (\forall y \in [b]) \, (\forall \vec{y} \in [n]^{k-i}) \, \neg S_{d,a}(\langle \vec{x}, y, \vec{y} \rangle_n, \langle \vec{x}, z, \vec{z} \rangle_n)$$

By coding succeeding quantifiers of the same type into one we obtain that this is equivalent to a $\Sigma_{d+1}^b(\alpha)$-formula. Hence we can use it as an inductive assertion in $S_1^{d+1}(\alpha)$. Hence we obtain:

Proposition 8. *The sentences $(\forall a) \, \mathrm{TOP}^d(|a|^k, \alpha)$ are consequences of $S_1^{d+1}(\alpha)$, for all $k \in \mathbb{N}$.*

7

3 Lower bounds

In this section we will show that $\mathrm{T}_1^d(\alpha)$ does not prove $(\forall a)\,\mathrm{TOP}^d(|a|^3, \alpha)$. The strategy will be as described in the introduction, i.e. we will translate potential $\mathrm{T}_1^d(\alpha)$-proofs into corresponding propositional proof systems and then utilize known lower bounds for the Total Ordering Principle in the propositional proof system.

In order to explain the relevant lower bounds for propositional proof systems and the translation therein we have to introduce some notions. First, we want to explain the propositional proof systems. We will follow Beckmann and Johannsen [5], which is based on Beckmann and Buss [4]. The proof system LK is a form of Gentzen's propositional LK. Formulas are build up from connectives $\bigvee, \bigwedge$ of unbounded, but finite fanin, and propositional variables and negated propositional variables. Negation for arbitrary formulas is defined as a syntactic operation according to the de Morgan rules.

Fix a set $\mathcal{A}$ of cedents which will serve as additional axioms. In the proof system we consider finite sets of formulas, which are called cedents. LK-derivations from hypotheses $\mathcal{A}$ have the following axioms and inference rules:

$$\text{Logical Axiom: } \frac{}{\Gamma, \neg x, x}\,, \quad \text{for variables } x$$

$$\text{Non-Logical Axiom: } \frac{}{\Gamma, \Phi} \quad \text{and} \quad \frac{}{\Gamma, \bigvee \Phi}\,, \quad \text{for } \Phi \in \mathcal{A}$$

$$\bigvee \frac{\Gamma, \varphi}{\Gamma, \bigvee \Phi}\,, \quad \text{where } \varphi \in \Phi$$

$$\bigwedge \frac{\Gamma, \varphi \qquad \text{for all } \varphi \in \Phi}{\Gamma, \bigwedge \Phi}$$

$$\text{Cut: } \frac{\Gamma, \neg \varphi \qquad \Gamma, \varphi}{\Gamma}$$

The formula φ in the Cut-rule is called *cut-formula*.

The constants 0 and 1 are defined as abbreviations of the empty disjunction resp. empty conjunction. Observe that an introduction rule for 1 is implicit in the $\bigwedge$-rule.

A derivation in LK is a finite tree of cedents such that for every cedent in the tree, the cedent together with its children forms an instance of one of the inference rules. If Γ is the cedent at the root of the tree we say that Γ *has an* LK-*derivation from* $\mathcal{A}$. We say that $\mathcal{A}$ *has an* LK-*refutation* iff the empty cedent has an LK-derivation from $\mathcal{A}$.

There are several ways to measure the complexity of derivations: by their tree-size, defined as the number of occurrences of cedents, i.e., the number of nodes in the tree; their dag-size, defined as the number of different cedents; and their tree-height, defined as the length of the longest path from the root to some leaf, not counting the root. Here we will only be concerned with the height of derivations. A comparison of the different measures has been carried out for example by Beckmann and Buss [4].

Constant depth LK will be defined by restricting all cut-formulas in an LK-derivation to certain sets of constant depth formulas which we will define next. Fix a width parameter w. To calculate the depth of a formula it is common to count the depth of bottom-level connectives of logarithmically small fan-in, i.e., $\bigwedge_{i < \log w} l_i$ and $\bigvee_{i < \log w} l_i$ for literals l_i only by $\frac{1}{2}$. This motivates the following definition:

Definition. *Let w, d be in $\mathbb{N}$. We inductively define $\varphi \in \Theta_d^w$ and $\varphi \in \Theta_{d+0.5}^w$ by the following clauses:*

(1) $\varphi \in \Theta_0^w$ *iff φ is a literal.*

(2) $\varphi \in \Theta_{0.5}^w$ *iff φ is a literal or a $\bigwedge$ or $\bigvee$ of at most $\log w$ many literals.*

(3) $\varphi \in \Theta_{d+1}^w$ *iff φ is in Θ_d^w, or it has the form of a $\bigwedge$ or $\bigvee$ of at most w many formulas from Θ_d^w .*

Let $d \in \frac{1}{2}\mathbb{N} := \{0, \frac{1}{2}, 1, 1\frac{1}{2}, 2, 2\frac{1}{2}, \dots\}$.

Definition. *An* LK-*derivation is called a Θ_d^w-LK-derivation if all occurring* cut-formulas *are in Θ_d^w.*

The previously defined complexity measure "tree-height" only counts cedents in LK derivations. If we speak about Θ_d^w-LK of a certain tree-height this implicitly bounds also the total number of symbols in the proof. If we speak about d-LK we still want some control of the symbol size of the derivation. This is implicit in the next definition.

Definition (of d-LK). *A d-LK-derivation of tree-size s is a Θ_d^s-LK-derivation of tree-size s. A d-LK-derivation of tree-height h is a $\Theta_d^{2^h}$-LK-derivation of tree-height h.*

The relationship between tree-like LK and height-restricted LK is well known. For example, let $\{\mathcal{A}_n\}_n$ be a family of sets of cedents Φ with $\bigvee\Phi \in \Theta_{d+1}^{n^{O(1)}}$. Then $\mathcal{A}_n$ has a d-LK refutation of tree-size polynomial in n, for all n, if and only if $\mathcal{A}_n$ has a $(d+1)$-LK refutation of tree-height logarithmic in n, which at the same time has tree-size polynomial in n, and has $O(1)$ many formulas in each cedent, for all n (cf. [4]).

In terms of tree-height the following separation results are known.

Theorem 9 ([4]). *Fix $d \in \frac{1}{2}\mathbb{N}$. For sufficiently large h there are negations of tautologies of depth $d+2$ which have $(d+2)$-LK-refutations of tree-height h, but every $(d+1.5)$-LK-refutation of them requires tree-height $2^{\Omega(h)}$.*

The separating principle is a form of the (Total) Ordering Principle; the lower bound which is used in this separation will be utilized later to obtain the independence results for linearly bounded arithmetic.

Next, we describe the Paris-Wilkie translation [16] from relativized bounded arithmetic to LK. A translation * between the languages is given as follows. It is defined for all bounded formulas of relativized bounded arithmetic, which are closed, i.e. do not contain free first order variables. For a closed term t, let $t^\mathbb{N}$ denote the value of t in the standard interpretation of the symbols.

(1) Consider the atomic formula $s \le t$. By assumption s and t are closed terms. We define
$$(s \le t)^* \quad := \quad \begin{cases} 1 & : & \text{if } s^\mathbb{N} \le t^\mathbb{N} \\ 0 & : & \text{otherwise}. \end{cases}$$

Similar for $s = t$.

(2) Consider $\alpha(s)$. We define $\quad \alpha(s)^* := p_{s^\mathbb{N}}$.

(3)
$$\begin{aligned} (\neg\varphi)^* &:= \neg\varphi^* \\ (\varphi_0 \wedge \varphi_1)^* &:= \varphi_0^* \wedge \varphi_1^* \\ (\varphi_0 \vee \varphi_1)^* &:= \varphi_0^* \vee \varphi_1^* \end{aligned}$$

(4) Consider $(\forall x \le t)\varphi(x)$. We define

$$((\forall x \le t)\varphi(x))^* \quad := \quad \bigwedge_{i \le t^{\mathbb{N}}} \varphi(i)^* \ .$$

Similar for $(\exists x \le t)\varphi(x)$.

Every closed $\Sigma_d^{\mathrm{b}}(\alpha)$-formula is translated to a formula equivalent to a $\Theta_{d+0.5}$ formula by possibly merging levels of connectives of the same kind. Furthermore, the missing size-parameter can be chosen polynomial in the parameters of the formula if the underlying language is $\mathcal{L}_1(\alpha)$ (in case of $\mathcal{L}_2(\alpha)$ the size parameter grows quasi-polynomial in the parameters of the formula). I.e., if $\varphi(a) \in \Sigma_d^{\mathrm{b}}(\alpha)$ then $\varphi(n)^* \in \Theta_{d+0.5}^{n^{O(1)}}$.

Having stated this we directly have to correct it: The translation of $\Sigma_d^{\mathrm{b}}(\alpha)$-formulas in general does not produce $\Theta_{d+0.5}$ formulas because of the boolean combination of atomic formulas in the kernel of Σ_d^{b} formulas. In order to be very precise we would have to proceed as follows: First, we redefine in the definition of Σ_d^{b} the class Σ_1^{b} by allowing only conjunctive normal forms instead of arbitrary boolean combinations of atomic formulas. The theories based on this definition of Σ_d^{b} are the same as the ones defined before. Second, we redefine in the definition of Θ_d^w the class $\Theta_{0.5}^w$ by allowing all formulas of size w such that all subformulas are in $\Delta_1^{\log w}$, i.e. are at the same time equivalent to some $\bigwedge$ of $\bigvee$'s of at most $\log w$ literals and to some $\bigvee$ of $\bigwedge$'s of at most $\log w$ literals. Then it is clear that Σ_d^{b} formulas translate to $\Theta_{d+0.5}^w$ formulas of that kind, because formulas of the form $\bigwedge_{i<c\log n} \bigvee_{j<c} \varphi_{ij}$ are equivalent to $\bigvee_{J<c^{c\log n}} \bigwedge_{i<c\log n} \varphi_{i(J)_i}$. The lower bounds for TOP from Beckmann and Buss [4] used later are also true for this alternative definition of Θ-classes, the proofs in [4], esp. the proof of Cut Reduction by Switching, are literally also true for this refined definition of Θ_d^w.

By normalizing proofs (i.e. partial cut-elimination) and unraveling induction we directly obtain the following translation of proofs (see [5, 3] for an exposition.)

Theorem 10. *Let $\varphi(x)$ be a bounded $\mathcal{L}_1(\alpha)$-formula with all free variables shown, and assume that $\mathrm{T}_1^d(\alpha)$ proves $(\forall x)\varphi(x)$ for $d > 0$. Then $\varphi(n)^*$ has $(d+0.5)$-LK-derivations of tree-height $O(\log n)$.*

The idea for a proof of this is that the length of each application of induction is bounded polynomially in n, and hence such an application of induction can be translated by a balanced tree of cuts which in turn has logarithmic height.

A stronger assertion also holds that under the assumptions of the Theorem $\varphi(n)^*$ has $(d-0.5)$-LK-derivations of tree-size polynomial in n (cf. [11], which is based on [13, 12]). This can be obtained by proving Theorem 10 in a slightly stronger form: it can be shown under the assumptions of the Theorem that the LK-derivations have in addition their tree-size polynomial in n, and that there is a constant which bounds the number of formulas in each cedent. This can then be used to transform the obtained derivations into polynomial tree-size $(d-0.5)$-LK-derivations using the relationship of height-restricted LK to tree-like LK as described above. In our exposition this is not needed, because we directly work with lower bounds on the height of $\neg \mathrm{TOP}^d(n)$. (Remark: Actually, the lower bounds on the tree-size as shown by Beckmann and Buss [4] are obtained by first transforming tree-like derivations into height-restricted derivations, and then showing a lower bound on the tree-height. This means that going to tree-like derivations in our situation would be a detour.)

Finally we explain what lower bounds are known for the translation of the Total Ordering Principle. The propositional proof complexity of the (Total) Ordering Principle has been studied at several places, we will follow Beckmann and Buss [4]. Let $\neg \mathrm{TOP}^d(n)$ denote the set of clauses corresponding to $\neg \mathrm{TOP}^d(n, \alpha)^*$ (For

a definition of the set of clauses for the Ordering Principle see [5, 4]; they can easily be modified to the set of clauses for the Total Ordering Principle). We use the formulation of the lower bound as given by Beckmann and Johannsen [5]; it is implicit in the proofs given by Beckmann and Buss [4]. (Attention: the Ordering Principle $\neg\,\mathrm{OP}^d(n)$ as defined in these papers corresponds in the present notation to $\neg\,\mathrm{TOP}^{d+1}(n)$ (modulo totality).)

Theorem 11 ([4]). *Let $d \in \mathbb{N}$ and $0 < \epsilon < \frac{1}{2}$. The tree-height of any $(d+1.5)$-LK-refutation of $\neg\,\mathrm{TOP}^{d+1}(n)$ must be larger than n^ϵ, for sufficiently large n.*

The last two Theorems together plus a direct transformation of derivations to refutations (which, for example, is described by Beckmann and Johannsen [5]) yield the missing part of our Main Theorem.

Proposition 12. *The sentence $(\forall a)\,\mathrm{TOP}^d(|a|^3, \alpha)$ is not a consequence of $\mathrm{T}_1^d(\alpha)$, for $d > 0$.*

Proof. Assume for the sake of contradiction that $(\forall a)\,\mathrm{TOP}^d(|a|^3, \alpha)$ is provable in $\mathrm{T}_1^d(\alpha)$. Translating this proof using Theorem 10 shows that $\mathrm{TOP}^d((\log n)^3, \alpha)^*$ has $(d+0.5)$-LK-derivations of tree-height $O(\log n)$. Now, transforming derivations into refutations increases the tree-height only by a constant factor, which implies that $\neg\,\mathrm{TOP}^d((\log n)^3)$ has $(d+0.5)$-LK-refutations of tree-height $h_n = O(\log n)$.

On the other hand, the lower bound from Theorem 11 shows for $\epsilon = \frac{2}{5}$ that

$$ h_n \quad \geq \quad ((\log n)^3)^{\frac{2}{5}} \quad = \quad (\log n)^{\frac{6}{5}} $$

for large n, which is impossible. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

4 Conclusion

We have shown that the conservativity of $\mathrm{S}_2^{i+1}(\alpha)$ over $\mathrm{T}_2^i(\alpha)$ w.r.t. $\forall\Sigma_{i+1}^{\mathrm{b}}(\alpha)$-sentences does not carry over to $\mathrm{S}_1^{i+1}(\alpha)$ and $\mathrm{T}_1^i(\alpha)$. This shows that the conservativity relations between fragments of linearly bounded arithmetic can be different from the situation in polynomially bounded arithmetic. Thus the result of Impagliazzo and Krajíček (Theorem 5) does not necessarily give much hope towards the solution of problem of the separation between S_2^i and T_2^i.

Our result also demonstrates again the usefulness of the approach to separate bounded arithmetic theories by the method of translation into propositional proof systems. Note that for the fragments of linearly bounded arithmetic, this method is the only one available, since there is no useful computational characterization of the definable functions in these theories.

We conclude with the statement of the main problem left open by this work: for what class of formulas is S_1^{i+1} conservative over T_1^i? Can one prove conservativity w.r.t. $\forall\Sigma_i^{\mathrm{b}}$-sentences? Or can $\mathrm{S}_1^{i+1}(\alpha)$ be separated from $\mathrm{T}_1^i(\alpha)$ by a sentence of lower complexity? By the result of Impagliazzo and Krajíček (Theorem 5) we know that at least one pair of theories $\mathrm{T}_1^i(\alpha), \mathrm{S}_1^{i+1}(\alpha)$ or $\mathrm{S}_1^{i+1}(\alpha), \mathrm{T}_1^{i+1}(\alpha)$ is separated by a $\forall\Sigma_2^{\mathrm{b}}(\alpha)$-sentence, but we do not know which one. We conjecture that no $\forall\Sigma_0^{\mathrm{b}}(\alpha)$-conservativity relation holds between any two of these theories. This conjecture is again inspired by our observations about the dynamic ordinals of these theories (cf. [3]).

References

[1] Miklos Ajtai. The complexity of the pigeonhole principle. *Combinatorica*, 14:417–433, 1994. Preliminary Version in Proc. 29th Symposium on Foundations of Computer Science, 1988.

[2] Michael Alekhnovich, Jan Johannsen, Toniann Pitassi, and Alasdair Urquhart. An exponential separation between regular and general resolution. In *Proc. 34th ACM Symposium on Theory of Computing*, pages 448–456, 2002.

[3] Arnold Beckmann. Dynamic ordinal analysis. *Archive for Mathematical Logic*, 42:303–334, 2003.

[4] Arnold Beckmann and Samuel R. Buss. Separation results for the size of constant-depth propositional proofs. Submitted, 2003.

[5] Arnold Beckmann and Jan Johannsen. *Bounded arithmetic and resolution based proof systems*. Collegium Logicum Vol. 7. Kurt Gödel Society, Vienna, Austria, 2004. In press.

[6] Maria Luisa Bonet and Nicola Galesi. Optimality of size-width tradeoffs for resolution. *Computational Complexity*, 10(4):261–276, 2001.

[7] Samuel R. Buss. *Bounded arithmetic*, volume 3 of *Studies in Proof Theory, Lecture Notes*. Bibliopolis, Naples, 1986.

[8] Samuel R. Buss. Axiomatizations and conservation results for fragments of bounded arithmetic. In Wilfried Sieg, editor, *Logic and Computation*, volume 106 of *Contemporary Mathematics*, pages 57–84. American Mathematical Society, Providence, 1990.

[9] Samuel R. Buss and Jan Krajíček. An application of boolean complexity to separation problems in bounded arithmetic. *Proceedings of the London Mathematical Society*, 69:1–21, 1994.

[10] Mario Chiari and Jan Krajíček. Witnessing functions in bounded arithmetic and search problems. *Journal of Symbolic Logic*, 63:1095–1115, 1998.

[11] Russell Impagliazzo and Jan Krajíček. A note on conservativity relations among bounded arithmetic theories. *Mathematical Logic Quarterly*, 48(3):375–377, 2002.

[12] Jan Krajíček. Lower bounds to the size of constant-depth propositional proofs. *Journal of Symbolic Logic*, 59:73–86, 1994.

[13] Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Cambridge University Press, Heidelberg/New York, 1995.

[14] Jan Krajíček, Pavel Pudlák, and Gaisi Takeuti. Bounded arithmetic and the polynomial hierarchy. *Annals of Pure and Applied Logic*, 52:143–153, 1991.

[15] Jan Krajíček, Pavel Pudlák, and Alan R. Woods. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Structures and Algorithms*, 7:15 ff., 1995.

[16] J. Paris and A. Wilkie. Counting problems in bounded arithmetic. In *Methods in mathematical logic (Caracas, 1983)*, volume 1130 of *Lecture Notes in Math.*, pages 317–340. Springer, Berlin, 1985.

[17] Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3:97–108, 1993.

[18] Wolfram Pohlers. *Proof Theory. An Introduction*. Number 1407 in Lecture Notes in Mathematics. Springer-Verlag, Berlin/Heidelberg/New York, 1989.

[19] Chris Pollett. Structure and definability in general bounded arithmetic theories. *Annals of Pure and Applied Logic*, 100:189–245, 1999.

[20] Nathan Segerlind, Samuel R. Buss, and Russell Impagliazzo. A switching lemma for small restrictions and lower bounds for k-DNF resolution. *SIAM Journal on Computing*, 33(5):1171–1200, 2004.

[21] Andrew C. Yao. Separating the polynomial-time hierarchy by oracles. *Proceedings of th 26th IEEE Symposium on Foundations of Computer Science*, pages 1–10, 1985.