Bounded Arithmetic and Bounded Depth Propositional Proofs

Sam Buss
Department of Mathematics
U.C. San Diego

Talk outline

- Normal forms for bounded arithmetic proofs.
- Propositional LK proofs, of constant Σ'-depth.
- The Paris-Wilkie translation.
- New proofs of the Main Witnessing Theorems for S^1_2 and T^1_2.
  Function classes:
  $\mathrm{FP} = \Box^p_1$ (Polynomial time computable functions)
  PLS (Polynomial Local Search).

Bounded Arithmetic Theories

Language: $0,, S, +, \cdot, \#, \lfloor \frac{1}{2}x \rfloor, |x|, \leq$, Gödel β function, $\langle\rangle$, $*$ (last three for sequence coding operations).

Bounded Quantifiers: $(\forall x \leq t), (\exists x \leq t)$.

Sharply Bounded Quantifiers: $(\forall x \leq |t|), (\exists x \leq |t|)$.

Σ_i^b- **and** Π_i^b-**formulas** . Count alternations of bounded quantifiers, ignoring sharply bounded quantifiers.

Σ_i^b-**PIND induction axioms** For S_2^i:

$$A(0) \wedge (\forall x)(A(\lfloor \tfrac{1}{2}x \rfloor) \rightarrow A(x)) \rightarrow (\forall x)A(x).$$

Σ_i^b-**IND induction axioms** For T_2^i:

$$A(0) \wedge (\forall x)(A(x) \rightarrow A(x+1)) \rightarrow (\forall x)A(x).$$

Restricted Bounded Arithmetic

We define some notions similar to Takeuti's *"pure i-form"* and *"strictly i-normal proof"*.

A formula is *form restricted* Σ_i^b if it is:

$$(\exists y_1 \leq t_1)(\forall y_2 \leq t_2) \cdots (Qy_i \leq t_i)(\overline{Q}z \leq |r|)B,$$

where B is quantifier-free. Quantifiers alternate between $\exists$ and $\forall$. Every Σ_i^b-formula is equivalent to a form restricted one: this fact can be proved in S_2^i using induction on only form restricted Σ_i^b-formulas.

Therefore, by free-cut elimination, bounded arithmetic may be equivalently formulated with induction allowed on only form restricted Σ_i^b-formulas.

Restricted by parameter variables. Let P be a proof. The free variables in the endsequent, $\vec{a}$, are called *parameter variables*. A quantifier $(Qx \leq t)$ is *restricted by parameter variables* iff t uses only parameter variables.

A proof is *restricted by parameter variables* iff (a) every quantifier is restricted by parameter variables and (b) every sequent which contains a non-parameter b contains a formula $b \leq t(\vec{a})$ in its antecedent.

Theorem 1. *Let R be S^i_2 or T^i_2, $i \geq 1$. If A is a form restricted Σ^b_i-formula and $R \vdash A$, then there is an R-proof of A which is restricted by parameter variables and in which every formula is form restricted Σ^b_i.*

Such proofs are called *restricted-Σ^b_i*. These proofs are conveniently formed for translation into propositional logic.

Constant depth propositional LK proofs

Syntax: Tait-style calculus. Variables: p. Literals: $p, \overline{p}$.

Unbounded fanin OR's and AND's: $\bigvee$ and $\bigwedge$.

Cedent Γ is set of formulas; intended meaning is the disjunction $\bigvee \Gamma$.

Axioms: *Neg:* $p, \overline{p}$ *Taut:* Γ , where Γ is a tautology.

Rules of inference:

$\bigvee$: $\dfrac{\Gamma, \varphi_{i_0}}{\Gamma, \bigvee_{i \in \mathcal{I}} \varphi_i}$, where $i_0 \in \mathcal{I}$ $\bigwedge$: $\dfrac{\Gamma, \varphi_i \quad \text{for all } i \in \mathcal{I}}{\Gamma, \bigwedge_{i \in \mathcal{I}} \varphi_i}$

Weakening: $\dfrac{\Gamma}{\Gamma, \Delta}$ *Cut:* $\dfrac{\Gamma, \varphi \quad \Gamma, \overline{\varphi}}{\Gamma}$

Σ'-depth of LK formulas and proofs

Definition Let S be a proof size parameter (size upper bound). The formulas that have Σ'-depth d *with respect to* S are inductively defined as follows:

a. If φ has size $\leq \log S$, then φ has Σ'-depth 0.

b. If φ has Σ'-depth d, then it has Σ'-depth d' for all $d' > d$.

c. If each φ_i has Σ'-depth d, then $\bigvee_{i\in\mathcal{I}} \varphi_i$ and $\bigwedge_{i\in\mathcal{I}} \varphi_i$ have Σ'-depth $(d+1)$.

Definition Let S be a size parameter. An LK-proof P is a Σ'-depth d proof of size S provided:

a. P has $\leq S$ symbols,

b. Every formula in P has Σ'-depth d,

c. Every *Taut* axiom has size at most $\log S$. (Only small tautologies allowed).

Similar definitions: Krajíček['94] of Σ-depth; Beckmann-Buss['03] of Θ-depth.

Conversion from S_2^i, T_2^i to LK

Let $d \geq 1$ and R be one of S_2^d or T_2^d. Suppose $A(x)$ is form restricted Σ_d^b and $R \vdash A$. We describe how to transform a restricted proof of A into a Σ'-depth d LK proof. W.l.o.g., x is the only parameter variable.

First step: choose an arbitrary value $n \in \mathbb{N}$. The translation $[\![A]\!]_n$ is a propositional formula stating that $A(x)$ is true for all x such that $|x| \leq n$. The free variables of $[\![A]\!]_n$ are variables $p_{x,i}$ representing the i-th bit of the binary representation of x.

For quantifier-free formulas φ, the formula $[\![\varphi]\!]$ is defined with any polynomial size formula that expresses the value of ϕ. (All function and relation symbols are describable with polynomial size formulas.) Because we have the *Taut* axioms, the choice of translation formula $[\![\phi]\!]$ is unimportant. Note $[\![\phi]\!]$ has size only $m^{O(1)}$ if the free variables of φ are integers of length $\leq m$. We will have $m = n^{O(1)}$ and $m < \log S(n)$.

Consider a sharply bounded formula $(\forall y \leq |s|)B$ or $(\exists y \leq |s|)B$. Because the term s contains only parameter variables as variables, and since the parameter variables have at most n bits, we can find a bound $n_y = n^{O(1)}$ such that $|s| \leq n_y$. Then,

$$[\![(\forall y \leq |s|)B]\!] = \bigwedge_{i=0}^{n_y} [\![y \leq |s| \rightarrow B]\!]/(y \mapsto i). \tag{1}$$

The notation "$\psi/(y \mapsto i)$" means replace each $p_{y,j}$ by the (constant) jth bit of the integer i. $[\![(\forall y \leq |s|)B]\!]$ has size only $n^{O(1)}$. Thus, it has Σ'-depth 0 for suitable $S(n) = 2^{n^{O(1)}}$.

General bounded quantifiers translated by exactly the same construction, but have bigger size: $2^{n^{O(1)}}$.

A Σ^b_d-formula becomes a Σ'-depth d formula for suitable $S(n) = 2^{n^{O(1)}}$.

To translate a cedent Γ, view it as a Tait-style cedent by moving all formulas to right of the sequent. All non-parameter variable $y_1, \ldots, y_k$ are restricted by parameter variables. So $|y_j| \leq n_j$ for some $n_j = n^{O(1)}$.

Γ is translated into a set of cedents, one cedent for each choice of $i_1, \ldots, i_k$ with each $|i_j| < n_j$. The cedents are just

$$[\![\Gamma]\!]/(y_1 \mapsto i_1, \ldots, y_k \mapsto i_k),$$

where the translation is applied individually to each formula. Note: the only variables left are $p_{x,i}$.

As the next theorem states, the translated cedents Γ can be pieced together into a valid proof.

Paris-Wilkie translation theorem

Theorem 2. *Let $i \geq 1$. Suppose $A(x) \in \Sigma^b_i$. Let $[\![A]\!]_n$ denote the propositional translation of A; $[\![A]\!]_n$ has free variables $p_{x,i}$, for $i < n$.*

a. *Suppose $S^i_2 \vdash A$. Then there is a function $S(n) = 2^{n^{O(1)}}$ such that, for all n, $[\![A]\!]_n$ has a Σ'-depth i proof of size $S(n)$. This proof*

- **i.** *has height $O(\log \log S(n))$, and*
- **ii.** *contains only $O(1)$ many formulas in each cedent.*

b. *Suppose $T^i_2 \vdash A$. Then there is a function $S(n) = 2^{n^{O(1)}}$ such that, for all n, $[\![A]\!]_n$ has a Σ'-depth i proof of size $S(n)$. This proof*

- **i.** *has height $O(\log S(n))$, and*
- **ii.** *contains only $O(1)$ many formulas per cedent.*

Similar theorems apply to $S^i_2(\alpha)$ and $T^i_2(\alpha)$.

One case of the proof: translation of $\wedge$:right inference

An $\wedge$*:right* inference

$$\frac{\Gamma, \varphi \quad \Gamma, \psi}{\Gamma, \varphi \wedge \psi}$$

translates to

$$\dfrac{\dfrac{[\![\Gamma]\!], [\![\psi]\!] \qquad [\![\Gamma]\!], [\![\psi \wedge \phi]\!], \overline{[\![\phi]\!]}, \overline{[\![\psi]\!]}}{[\![\Gamma]\!], [\![\phi]\!] \qquad [\![\Gamma]\!], [\![\psi \wedge \phi]\!], \overline{[\![\phi]\!]}}\ \textit{Cut} \qquad \dfrac{[\![\psi \wedge \phi]\!], \overline{[\![\phi]\!]}, \overline{[\![\psi]\!]}}{[\![\Gamma]\!], [\![\psi \wedge \phi]\!], \overline{[\![\phi]\!]}, \overline{[\![\psi]\!]}}\ \textit{Weakening}}{[\![\Gamma]\!], [\![\psi \wedge \phi]\!]}\ \textit{Cut}$$

Note that the upper right sequent is a *Taut* axiom.

Another case of the proof: induction rule

Consider an induction inference in P. This translates into m *Cut* inferences in the LK proof, where m is the "length" of the induction. By balancing the tree of cuts, the height (maximal number of cedents along any branch) is only $O(\log m)$. (The induction bound t involves only parameter variables.)

If R is S^i_2, the induction inference translates into $|t| = n^{O(1)}$ many cuts, so the height is $O(\log n)$.

If R is T^i_2, the induction inference translates into $t = 2^{n^{O(1)}}$ many cuts, so the height is $O(n^{O(1)})$. □

Important fact: The LK-proofs given by Theorem 2 are polynomial time uniform.

Main Theorem for S^1_2

Theorem 3. (Buss ['85]) *Suppose* $A(x,y) \in \Sigma^b_1$ *and that* S^1_2 *proves* $(\forall x)(\exists y)A(x,y)$. *Then there is a polynomial time function* $f(x) = y$ *such that for all* $x \in \mathbb{N}$, $A(x, f(x))$ *holds.*

Proof By Parikh, $S^1_2 \vdash (\exists y \leq s(x))A(x,y)$. x is the parameter variable. Applying Theorem 2(a) yields a Σ'-depth 1 proof; adding a *Cut* to the end of this proof turns the proof into a refutation R of

$$[\![\forall y \leq s(x))\neg A(x,y)]\!]. \tag{2}$$

We give a polynomial time procedure that is has as input a particular value for x, and traverses the refutation R until it arrives at a false initial cedent. Of necessity this false initial cedent is the cedent (2), and when it is reached, the procedure will know a value y that falsifies the cedent. This value for y will be the value of $f(x)$.

The polynomial time procedure acts as follows: it starts at the root of the proof and traverses the proof upward, backtracking as needed as described below. At each stage, the procedure is at some cedent Γ in the proof that it believes to be false. In particular, every Σ'-depth 0 formula in Γ is *False*. (Recall that the variables $p_{x,i}$ are the only variables in R, and the procedure has values for these.) Furthermore, for any formula in Γ which is a conjunction of Σ'-depth 0 formulas, a particular conjunct is known to be false. For the formulas which are a disjunction of Σ'-depth 1 formulas, the procedure does not know for sure that they are false, it merely tentatively assumes they are false.

At the beginning, the procedure is at the endsequent of R, which is the empty cedent.

Other cases are *Cut* inference, $\bigwedge$ inference, and $\bigvee$ inference....

If the procedure is at the lower cedent of a cut inference

$$\frac{\Gamma, \overline{\varphi} \quad \Gamma, \varphi}{\Gamma}$$

If φ is Σ'-depth 0, then it can be evaluated as being either *True* or *False*. If it is true, the procedure proceeds to the left upper cedent, otherwise, it proceeds to the right upper cedent. Otherwise, φ is w.l.o.g. a disjunction, and the algorithm proceeds to the right upper cedent.

If the procedure is at the lower cedent of a $\bigwedge$-inference:

$$\frac{\Gamma, \psi_i \qquad \text{, for } i \in \mathcal{I}}{\Gamma, \bigwedge_{i \in \mathcal{I}} \psi_i}$$

the algorithm acts as follows. By assumption, the procedure knows a value i_0 such that the conjunct ψ_{i_0} is false. The algorithm proceeds to the upper cedent Γ, ψ_{i_0} where $i = i_0$.

If the procedure is at the lower cedent of a $\bigvee$-inference:

$$\frac{\Gamma, \psi_{i_0}}{\Gamma, \bigvee_{i \in \mathcal{I}} \psi_i}$$

the algorithm acts as follows. If ψ_{i_0} is false, it proceeds to the upper cedent. However, if it is true, the algorithm has discovered a disjunct of $\varphi = \bigvee_{i \in \mathcal{I}} \psi_i$ which is true, contradicting the tentative assumption that φ was false. The procedure then backtracks down the path towards the root until it finds the *Cut* inference where the formula φ was added to the cedent. It then proceeds to the other upper cedent of the *Cut*, and saves the information about which conjunct of $\overline{\varphi}$ is false.

The run time is $O(n^{O(1)})$, because there are only this many *Cut*'s. It thus can terminate only at the cedent (2). When it reaches this, it knows a value for y that falsifies it. This value of y satisfies $A(x, y)$. □

The Main Theorem for T_2^i

PLS = Polynomial Local Search [Johnson, Papadimitriou, Yanakakis, '88].

Theorem 4. $(\mathrm{Buss},\ \mathrm{Krajíček},\ '94])$ *Suppose* $A(x,y) \in \Sigma_1^b$ *and that* T_2^1 *proves* $(\forall x)(\exists y)A(x,y)$. *Then there is a Polynomial Local Search (*PLS*) function* $f(x) = y$ *such that for all* $x \in \mathbb{N}$, $A(x, f(x))$ *holds.*

The proof is identical to before, based on exactly the same procedure. Now the procedure may need $2^{n^{O(1)}}$ steps, instead of $n^{O(1)}$. Use the position in the proof to define a decreasing cost function. □

Theorems 3 and 4 both hold if all true Π_1^b-formulas are added as axioms (no change to proof needed).

Main theorems for $i > 1$

Theorem 5. *Let $i > 1$ and $A(x, y) \in \Sigma_i^b$.*

a. *Suppose $S_2^i \vdash (\forall x)(\exists y)A(x, y)$. Then, there is a function f in $\Box_i^p = \mathrm{FP}^{\Sigma_{i-1}^p}$ such that $A(x, f(x))$ holds for all $x \in \mathbb{N}$. Here FP is the class of polynomial time functions, so $\Box_i^p$ is the class of functions computable in polynomial time relative to a Σ_{i-1}^p-oracle (i.e., relative to a set at the ith level of the polynomial time hierarchy.*

b. *Suppose $T_2^i \vdash (\forall x)(\exists y)A(x, y)$. Then, there is a function f in $\mathrm{PLS}^{\Sigma_{i-1}^p}$ such that $A(x, f(x))$ holds for all $x \in \mathbb{N}$.*

Similar, relativized proofs work.

Transforming constant depth proofs.

Theorem 6. *(Based on [Krajíček,'94; Razborov'94; Beckmann-Buss'03].) Let $d \in \mathbb{N}$, and $\{\mathcal{A}_n\}_n$ be a family of sets of cedents. Then the following conditions (1) and (2) are equivalent:*

(1) *$\mathcal{A}_n$ has a Σ'-depth d LK refutation of sequence-size quasi-polynomial in n, for all n.*

(2) *$\mathcal{A}_n$ has a Σ'-depth $(d+1)$ LK refutation of tree-size quasi-polynomial in n, for all n.*

Furthermore, the following conditions (3) and (4) are equivalent:

(3) *$\mathcal{A}_n$ has Σ'-depth d LK refutation of tree-size quasi-polynomial in n, for all n.*

(4) *$\mathcal{A}_n$ has a Σ'-depth $(d+1)$ LK refutation which simultaneously has tree-size quasi-polynomial in n and height poly-logarithmic in n, for all n.*

Corollary 7. *Let $d \geq 2$. Suppose A is a Σ^b_d-formula and that $T^d_2 \vdash A$. Without loss of much generality, A has the form*

$$(\exists y \leq t(x))(\forall z \leq r(x))C(x, y, z).$$

Let $n_t = n^{O(1)}$ bound $|t(x)|$ for all $x < 2^n$, and $n_r = n^{O(1)}$ bound $|r(x)|$ for all $x < 2^n$ Then the set $\mathcal{A}_n$ of cedents

$$\left\{ [\![y \leq t \rightarrow (z \leq r(x) \wedge \neg C(x, y, z))]\!]_n / (y \mapsto i, z \mapsto j) \, : \, j < 2^{n_r} \right\},$$

for $i < 2^{n_t}$, has a Σ'-depth $(d-2)$ LK-refutation.

Explanation: In effect, $[\![A]\!]$ has a Σ'-depth $(d-2)$ proof.