

Collapsing Modular Counting in Bounded Arithmetic and Constant Depth Propositional Proofs

Sam Buss

joint work with L. Kołodziejczyk and K. Zdanowski

Limits of Theorem Proving
Rome, Italy
September 2012

$AC^0[p]$ circuits and proofs

Fix an integer $m \geq 2$.

Definition. $AC^0[m]$ **circuits** are *constant depth* circuits for computing boolean functions.

- Literals and constants: x , \bar{x} , \top (*True*), \perp (*False*).
- Unbounded fanin connectives: \wedge , \vee , \oplus_m .
- $\oplus_m(\vec{x}) = \text{True}$ iff the number of true arguments is $\equiv 0 \pmod m$.

Definition. An $AC^0[m]$ **proof** is a propositional proof in which each line is an $AC^0[m]$ -circuit (equivalently: is an $AC^0[m]$ -formula).

We are mostly interested in the case where m is equal to a prime p . We call $AC^0[p]$ proofs **constant depth** PK_{\oplus_p} **proofs**.

Theorem [Beigel-Tarui'94, see also Allender-Hertrampf'94, Toda'91, Yao, ...]. Let $m \geq 2$. An $AC^0[m]$ circuit of size S can be converted into a quasipolynomial-size depth three formula consisting of a symmetric gate applied to \oplus_m -gates applied to polylogarithmic size conjunctions of literals.

This expressive power of quasipolynomial size, depth three formulas, carries over to the power of propositional proofs:

Theorem [Maciel-Pitassi'98] Let $m = p^k$, a prime power. $AC^0[p^k]$ proofs can be quasipolynomially simulated by proofs in which every line is a depth three formula formed from a threshold gate applied to \oplus_m gates of polylogarithmic size conjunctions.

Recall: The “quasipolynomial simulation” means there is a $2^{\log^{O(1)} n}$ -time procedure to thusly convert $AC^0[p^k]$ proofs.

Open: Does this hold for composite m as well?

Collapsing $AC^0[p]$ without threshold gates

Theorem. [Allender-Hertrampf'94] Let p be prime. An $AC^0[p]$ circuit of size S can be converted into a quasipolynomial-size depth 4 circuit formed as disjunctions (\vee 's) of conjunctions (\wedge 's) of \oplus_p -gates of polylogarithmic conjunctions of literals.

Theorem A. [BKZ - this talk]

Let p be prime. Constant depth PK_{\oplus_p} proofs ($AC^0[p]$ proofs) can be quasipolynomially simulated by Tait-style propositional proofs in which each formula is a (subformula of a) depth three formula formed as a conjunction applied to \oplus_p gates of polylogarithmic size conjunctions.

Remainder of the talk will outline the proof.

- Constant-depth proof systems with \oplus_p connectives.
- Fragments of bounded arithmetic with approximate counting.
- Valiant-Vazirani and Toda theorems.
- A collapse of bounded arithmetic with modular counting.
- Paris-Wilkie translations.
- Reflection principles.
- Precise statement of simulation results.
- Concluding questions.

Tait-style systems $PK_{\oplus p}$

The lines of a Tait-calculus proof are **cedents**, sets of formulas which are interpreted as their disjunction. The system, PK , with connectives \wedge and \vee has the rules of inference:

$$\frac{\Gamma}{\Gamma, \Delta} \text{ Weakening}$$

$$\frac{\Gamma, \varphi \quad \Gamma, \bar{\varphi}}{\Gamma} \text{ Cut}$$

$$\frac{\Gamma, \varphi_{i_0}}{\Gamma, \bigvee_{i \in I} \varphi_i} \vee$$

$$\frac{\Gamma, \varphi_i \quad \text{for all } i \in I}{\Gamma, \bigwedge_{i \in I} \varphi_i} \wedge$$

where $i_0 \in I$.

- $\bar{\varphi}$ is the De-Morgan negation of φ .

For counting mod p , we use $2p$ many connectives \oplus_p^k and $\bar{\oplus}_p^k$ which are true (resp.) false when the number of true inputs is $\equiv k \pmod p$.

The system has the same rules of inference as PK, plus the initial cedents:

$$\begin{array}{ll} \varphi, \oplus_p^0\{\varphi\} & \bar{\varphi}, \oplus_p^1\{\varphi\} \\ \oplus_p^k\Phi, \bar{\oplus}_p^k\Phi & \bar{\oplus}_p^k\Phi, \bar{\oplus}_p^\ell\Phi, \text{ for } k \neq \ell \\ \bar{\oplus}_p^k\Phi, \bar{\oplus}_p^\ell\Psi, \oplus_p^{k+\ell}(\Phi \cup \Psi) & \end{array}$$

By convention, " $\oplus_p^{k+\ell}$ " means " $\oplus_p^{(k+\ell) \bmod p}$ ".

Constant depth PK_{\oplus_p} and PCK_p^i

Constant depth PK_{\oplus_p} proofs allow \wedge 's, \vee 's, and \oplus_p gates to appear at any level. **PCK_p^i proofs** are depth $i + 1\frac{1}{2}$ proofs in which the inputs to \oplus_p gates are restricted to be conjunctions of literals:

Definition. For $i \geq 0$, a **PCK_p^i proof** contains literals, conjunctions of literals, disjunctions of literals, and formulas that have $\leq i$ alternating levels of conjunctions and disjunctions above \oplus_p^j and $\bar{\oplus}_p^j$ gates, which are applied only to “small” conjunctions of literals.

Remark: It is also possible to work over finite fields of characteristic p , and replace the \oplus_p gates with low-degree polynomials. The resulting system corresponding to PCK_p^i is denoted $\text{PCK}_{\mathbb{F}_p}^i$. Similar results hold for $\text{PCK}_{\mathbb{F}_p}^i$ as for PCK_p^i ; this talk discusses only the propositional systems PCK_p^i however.

Recall that Theorem A gives a translation from constant depth PK_{\oplus_p} proofs into PCK_p^1 proofs.

For this, the size of PCK_p^i proofs is measured in terms of “ Σ -size”.

Definition. The Σ -size of a PCK_p^i proof P is $\leq S$ provided there are $\leq S$ formulas in P , each of size $\leq S$ symbols, and every conjunction or disjunction of literals in P has size $\leq \log S$.

Bounded arithmetic: subtheories of Peano arithmetic.

Bounded arithmetic theories have close connections with low-level complexity classes, but more relevantly for this talk, proofs in bounded arithmetic can be viewed as **uniform versions of constant depth proofs** via the Paris-Wilkie translation.

Function symbols: all polynomial time functions and relations.

Bounded quantifiers: $(\forall x \leq t)$ and $(\exists x \leq t)$.

Sharply bounded quantifiers: $(\forall x \leq |t|)$ and $(\exists x \leq |t|)$

Classes Σ_i^b and Π_i^b are the formulas containing $\leq i$ alternating blocks of bounded quantifiers, ignoring sharply bounded quantifiers.

The **strict** classes, $\hat{\Sigma}_i^b$ and $\hat{\Pi}_i^b$, require prenex form, and disallow sharply bounded quantifiers outside of bounded quantifiers.

\mathbf{T}_2^i is axiomatized with induction for all Σ_i^b - and Π_i^b -formulas.

Modular counting quantifiers

Fix a prime p . The syntax of bounded arithmetic is now augmented with **modular counting quantifiers** C_p^k , for $0 \leq k < p$. The meaning of

$$(C_p^k x \leq t)A(x)$$

is that the number of $x \leq t$ such that $A(x)$ is $\equiv k \pmod p$. The axioms for the C_p^k quantifiers are:

$$\begin{aligned} A(0) &\rightarrow (C_p^1 x \leq 0)A(x) & \neg A(0) &\rightarrow (C_p^0 x \leq 0)A(x) \\ A(t+1) \wedge (C_p^k x \leq t)A(x) &\rightarrow (C_p^{k+1} x \leq t+1)A(x) \\ \neg A(t+1) \wedge (C_p^k x \leq t)A(x) &\rightarrow (C_p^k x \leq t+1)A(x) \\ \neg[(C_p^k x \leq t)A(x) \wedge (C_p^\ell x \leq t)A(x)] & \quad \text{for } k \neq \ell \pmod p. \end{aligned}$$

The $(C_p^k \leq \dots)$ quantifiers above are considered bounded and may appear in bounded formulas. Notation: $\sum_{\infty}^{b, \oplus_p} P(\oplus_p)$.

Definition. The theory $\mathbf{T}_2(\oplus_p)$, equivalent to $\mathbf{S}_2(\oplus_p)$, has induction for *all* bounded formulas; i.e., allowing quantifiers $(C_p^k x \leq t)$ to appear in front of arbitrary bounded formulas.

Definition. A $\oplus_p \mathbf{P}$ -**formula** is atomic, or of the form $(C_p^k x \leq t)A(x)$ where A is a sharply bounded (so A polynomial time computable).

The $\hat{\Sigma}_i^{b, \oplus_p \mathbf{P}}$ and $\hat{\Pi}_i^{b, \oplus_p \mathbf{P}}$ **formulas** are defined by counting alternations of bounded \exists/\forall quantifiers acting on $\oplus_p \mathbf{P}$ formulas, ignoring sharply bounded quantifiers.

Definition. The theory $\mathbf{T}_2^{i, \oplus_p \mathbf{P}}$ is axiomatized with the C_p^k axioms for $\oplus_p \mathbf{P}$ formulas and with induction for $\Sigma_i^{b, \oplus_p \mathbf{P}}$ formulas.

The C_p^k quantifiers can be syntactically restricted to appear only in $\oplus_p \mathbf{P}$ formulas.

In essence, C_p^k quantifiers appear only in $\oplus_p P$ formulas for $T_2^{i, \oplus_p P}$.

This condition can be relaxed somewhat:

Theorem. We have

$$\oplus_p P^{\oplus_p P} = \oplus_p P.$$

In fact, any formula composed of sharply bounded quantifiers, bounded C_p^k quantifiers, boolean operations, and polynomial time predicates is equivalent to a $\oplus_p P$ predicate.

Thus, the theories $T_2^{i, \oplus_p P}$ are robustly defined.

Theorem B.

A lot more than that is true, however: The hierarchy of modular counting theories of bounded arithmetic collapses to the third level:

Theorem B. $T_2(\oplus_p)$ is conservative over $T_2^{3,\oplus_p P}$. In fact, any $\Sigma_\infty^b(\oplus_p)$ formula (i.e., any bounded formula) is provably equivalent to a $\Sigma_2^{b,\oplus_p P}$ formula.

Theorem B is one of the main ingredients for the proof of Theorem A. For its proof we introduce Jeřábek's bounded arithmetic theories for approximate counting.

Jeřábek's theories for approximate counting were axiomatized with the following, surjective, weak pigeonhole principle, **sWPHP**(**f**):

$$(\forall x)(\forall y)[x > 0 \rightarrow (\exists v \leq x(|y|+1))(\forall u \leq x|y|)(f(u) \neq v)]$$

Then, define (in the notation of [BKT]),

$$\text{APC}_1 := \text{PV}_1 + \text{sWPHP}(\text{PV}_1)$$

and

$$\text{APC}_2 := T_2^1 + \text{sWPHP}(\text{PV}_2).$$

PV_1 is the set of polynomial time functions.

PV_2 is the set of functions polynomial time relative to NP.

“APC” = “**AP**roximate **C**ounting”

We work with versions of APC_1 and APC_2 extended to include $\oplus_p P$ predicates and functions:

$$APC_1^{\oplus_p P} := PV_1^{\oplus_p P} + sWPHP(PV_1^{\oplus_p P})$$

and

$$APC_2^{\oplus_p P} := T_2^{1, \oplus_p P} + sWPHP(PV_2^{\oplus_p P}).$$

$PV_1^{\oplus_p P}$: functions polynomial time relative to $\oplus_p P$.

$PV_2^{\oplus_p P}$: functions that are polynomial time relative to $NP^{\oplus_p P}$.

Jeřábek showed that APC_1 can count the size of polynomial time sets to within a constant fraction ϵ .

Namely, let $X, Y \subseteq 2^n$ be defined by Boolean circuits. Roughly speaking the “size of X ” can be defined to within an error tolerance of $\epsilon 2^n$.

Specifically, the relation $\mathbf{X} \preceq_\epsilon \mathbf{Y}$ can be defined expressing

“there exists a nonzero $v \in \text{Log}$ and a circuit G such that G computes a surjection $v \times (Y \sqcup \epsilon 2^n) \rightarrow v \times X$.”

where \sqcup is disjoint union, where Log is the set of lengths, and where $\epsilon = 0$ or $\epsilon \in \text{Log}$.

For APC_2 , Jeřábek showed that the size of X can be approximated to within an error tolerance of $\epsilon|X|$.

Furthermore, APC_1 and APC_2 can prove many properties about approximate counting, including facts about union, intersection, some versions of exclusion/inclusion, Chebyshev inequalities, randomized computation, BPP, AM, MA, and so forth.

These constructions all relativize to $\text{APC}_1^{\oplus_p P}$ and $\text{APC}_2^{\oplus_p P}$.

Formalized Valiant-Vazirani theorem

Theorem C. (in APC_2) There is a PV_1 function which, given a CNF formula φ over the propositional variables $\vec{q} = \langle q_1, \dots, q_n \rangle$ and a (randomly chosen) value r of length $(n+3)n + |n|$, outputs a CNF formula φ_r with the same variables \vec{q} such that

$$\begin{aligned} \varphi \in \text{SAT} &\implies \Pr[\neg \exists^1 b, b \models \varphi_r] \leq 1 - \frac{1}{2^{|n|} \cdot 65}, \\ \varphi \notin \text{SAT} &\implies \varphi_r \notin \text{SAT}. \end{aligned}$$

“ $\exists^1 \mathbf{b}$ ” means there exists a **unique** b .

“ $\mathbf{b} \models \varphi_r$ ” means that b codes a satisfying assignment for φ_r as a string of n bits.

The constant $1/65$ is not as good as in VV'86 due to the need to formalize the result in APC_2 .

Definition. (in $\text{APC}_1^{\oplus p^P}$ or $\text{APC}_2^{\oplus p^P}$). $\oplus_p^k \text{SAT}$ is the set of propositional formulas φ such that the number of satisfying assignments of φ is congruent to $k \bmod p$.

A language L is in $\mathbf{BP} \cdot \oplus_p \mathbf{P}$ if there exist PV_1 functions f and u such that for all x ,

$$x \in L \iff \Pr_{r < u(x)} [f(x, r) \notin \oplus_p^1 \text{SAT}] \preceq_0 1/4,$$

$$x \notin L \iff \Pr_{r < u(x)} [f(x, r) \notin \oplus_p^0 \text{SAT}] \preceq_0 1/4.$$

Amplification allows the constant $1/4$ to be improved to 2^{-n^c} .

Formalized Toda theorem in $\text{APC}_2^{\oplus_p P}$

Theorem D. $T_2(\oplus_p)$ proves that any $\Sigma_\infty^b(\oplus_p)$ formula defines a property in $\text{BP} \cdot \oplus_p P$. Furthermore, these equivalences can be essentially expressed and proved in $\text{APC}_2^{\oplus_p P}$.

Proof idea: The necessary probabilistic arguments can be carried out (with difficulty) in $\text{APC}_2^{\oplus_p P}$.

Corollary. $T_2(\oplus_p)$, and in essence $\text{APC}_2^{\oplus_p P}$, can prove the uniform analogue of the Allender-Hertrampf theorem about the collapse of $\text{AC}^0[p]$.

Corollary E. $T_2(\oplus_p)$ is conservative over $\text{APC}_2^{\oplus_p P}$.

Since $T_2^{3, \oplus_p P} \models \text{APC}_2^{\oplus_p P}$, we also get

Corollary F. $T_2(\oplus_p)$ is conservative over $T_2^{3, \oplus_p P}$.

Paris-Wilkie translation

Let $\varphi(\vec{x})$ be a bounded formula involving an oracle α .

The Paris-Wilkie translation of $\varphi(\vec{x})$ gives an infinite family of constant depth propositional formulas $\llbracket \varphi \rrbracket_{\vec{n}}$.

The propositional variables of $\llbracket \varphi \rrbracket_{\vec{n}}$ are x_i 's indicating that $\alpha(i)$ is true.

The integer values \vec{n} assign values to free variables \vec{x} of φ .

Sharply bounded subformulas of φ become small depth (depth polylogarithmic in \vec{n}) decision trees in $\llbracket \varphi \rrbracket_{\vec{n}}$; expressed as a disjunction of small conjunctions.

$\oplus_p P$ subformulas of φ become a \oplus_p gate applied to small conjunctions.

Bounded quantifiers ($\exists x \leq t$) and ($\forall x \leq t$) in φ become big (quasipolynomial size in n) disjunctions or conjunctions $\llbracket \varphi \rrbracket_{\vec{n}}$.

Paris-Wilkie Theorem: (one of several forms)

Let $\varphi(x) \in \widehat{\Sigma}_i^{b, \oplus_p P}(\alpha)$ be a formula of the form

$$\varphi(x) := (\exists y \leq t(x)) (\forall z \leq s(x)) \xi(x, y, z).$$

so $\xi(x, y, z) \in \widehat{\Sigma}_{i-2}^{b, \oplus_p P}(\alpha)$.

Express $\neg\varphi(x)$ as the set Ξ_n of $t(n) + 1$ many cedents

$$\overline{\llbracket \xi \rrbracket_{n,m,0}}, \overline{\llbracket \xi \rrbracket_{n,m,1}}, \dots, \overline{\llbracket \xi \rrbracket_{n,m,s(n)}}$$

where $0 \leq m \leq t(n)$. Each cedent in Ξ_n has $s(n) + 1$ formulas.

Suppose $T_2^{i, \oplus_p P}(\alpha) \vdash (\forall x)\varphi(x)$, $i \geq 2$. Then

The set of cedents Ξ_n has a dag-like PCK_p^{i-2} refutation P such that the Σ -size of P is quasipolynomial in n .

A reflection theorem

The Reflection Principle, $j\text{-Ref}(d\text{-PK}_{\oplus_p})(\alpha, \beta, \gamma)$:

If β codes a $\Sigma_j(\oplus_p^-)$ propositional formula φ , and α codes a depth d PK_{\oplus_p} proof of φ , then the truth assignment coded by γ satisfies φ .

Note that α, β, γ are second-order, hence code exponentially large objects.

$j\text{-Ref}(d\text{-PK}_{\oplus_p})$ is a $\forall \Sigma_j^{b, \oplus_p P}(\alpha, \beta, \gamma)$ formula.

Reflection Theorem: $T_2(\oplus_p)(\alpha, \beta, \gamma) \vdash j\text{-Ref}(d\text{-PK}_{\oplus_p})$.

Corollary: $T_2^{3, \oplus_p P}(\alpha, \beta, \gamma) \vdash j\text{-Ref}(d\text{-PK}_{\oplus_p})$.

Proof idea: $T_2(\oplus_p)(\alpha, \beta, \gamma)$ can give a truth definition for the depth d formulas in the proof coded by α .

Theorem A'. Let φ be $\bigvee_{k < K} \bigwedge_{l < L_k} \psi_{k,l}$, where the $\psi_{k,l}$'s are $\Sigma_1(\oplus_p^-)$ (that is, PCK_p^1 -formulas). Suppose φ has a depth d PK_{\oplus_p} proof of size $\leq S$. Then the set of cedents

$$\{\overline{\psi_{k,0}}, \dots, \overline{\psi_{k,L_k-1}}\}_{k < K}$$

has a PCK_p^1 refutation of Σ -size $S^{\log^e S}$, where $e \in \mathbb{N}$ is a constant depending only on d .

Proof idea. Form the Paris-Wilkie translation of the reflection principle as provable in $T_2^{3, \oplus_p P}(\alpha, \beta, \gamma)$.

This gives a PCK_p^1 -proof Q .

In Q , substitute for the propositional variables describing values of α and β , the actual description of the proof P and the formula ψ . Using the fact that β really codes a valid proof P of ψ , the resulting restriction simplifies the proof Q into the desired refutation of the above cedents.

QED

Remark. It is possible to work with $APC_2^{\oplus pP}(\alpha, \beta, \gamma)$ instead $T_2^{3, \oplus pP}(\alpha, \beta, \gamma)$.

This allows replacing the PCK_p^1 proof of Theorem A' with a PCK_p^0 proof, but at the cost of adding additional initial cedents that express the surjective weak pigeonhole principle.

Some questions

1. What does the collapse of PK_{\oplus_p} proofs imply about the possibility of proving superpolynomial lower bounds on the size of PK_{\oplus_p} -proof (that is, $\text{AC}^0[p]$ -proofs)?
2. The theories $\text{APC}_1^{\oplus_p P}$ and $\text{APC}_2^{\oplus_p P}$, as well as APC_1 and APC_2 , deserve more study. For instance, is $\text{APC}_2 \subseteq T_2^2$?

Thank you!