Approximate counting in bounded arithmetic

Emil Jeřábek

Institute of Mathematics
Czech Academy of Sciences
jerabek@math.cas.cz
https://users.math.cas.cz/~jerabek/

Constructive Complexity Theory 2025

STOC TheoryFest workshop, Prague, June 2025

Plan for talk

1. Counting in arithmetic
2. Weak pigeonhole principle
3. Counting with additive error
4. Counting with multiplicative error

Counting in arithmetic

Finite sets and counting

Theories of arithmetic

- official objects: natural numbers
- encode all kinds of stuff: strings, graphs, circuits, ...

We can manipulate finite sets of such objects

The counting problem for a set X:

- Can the theory meaningfully determine the size $\#X$?

Why bother?

Potential applications:

- counting arguments in combinatorics
  e.g., Ramsey's theorem
- probabilistic arguments
- formalization of randomized algorithms,
  randomized (and counting) complexity classes
- translation of arithmetic to propositional proofs
  $\implies$ propositional simulations

Bounded arithmetic cheat sheet

Base theory PV_1:

- objects: natural numbers ("in binary")
- symbols for all poly-time functions and relations
  $|x| = \lceil \log_2(x+1) \rceil$
- induction for quantifier-free formulas φ

$$\varphi(0,\dots) \land \forall x\,(\varphi(x,\dots) \to \varphi(x+1,\dots)) \to \forall x\,\varphi(x,\dots)$$

$\mathsf{T}^i_2 = \mathsf{PV}_1 +$ induction for Σ^b_i formulas $\equiv \mathsf{PV}_{i+1}$

$$\exists \vec{x}_1 \le t_1\, \forall \vec{x}_2 \le t_2\, \dots Q\vec{x}_i \le t_i\, \theta(\vec{x}_1,\dots,\vec{x}_i,x,\dots)$$

$\mathsf{S}^i_2 = \mathsf{PV}_1 +$ length-induction for Σ^b_i formulas

$$\varphi(0,\dots) \land \forall x\,(\varphi(x,\dots) \to \varphi(x+1,\dots)) \to \forall x\,\varphi(|x|,\dots)$$

$$\mathsf{T}^0_2 \equiv \mathsf{PV}_1 \subseteq \mathsf{S}^1_2 \ \subseteq\ \mathsf{T}^1_2 \subseteq \mathsf{S}^2_2 \ \subseteq\ \mathsf{T}^2_2 \subseteq \mathsf{S}^3_2 \ \subseteq \cdots$$

One-sorted vs. two-sorted theories

Two-sorted bounded arithmetic:

- (unary/small/auxiliary) natural numbers
- finite sets (= binary strings = binary/large numbers)

one-sorted	two-sorted
numbers	sets (binary numbers)
logarithmic numbers	numbers (unary)
PV_1	VPV
S^i_2, T^i_2	V^i, TV^i
length induction	induction
induction	string induction

We are talking about finite sets of the "binary objects"

$\Longrightarrow$ prefer one-sorted theories to prevent confusion

Representation of sets

Sets (sequences) $X = \{x_i : i < n\}$ encoded by numbers

- can only have logarithmic size
- straightforward to count by a PV-function

Bounded definable sets $X = \{x < a : \varphi(x, p)\}$

- X specified inside the theory by a, p (φ is fixed)
- φ low complexity
  e.g.: φ quantifier-free $\implies$ P/poly sets
  alternatively: given by a circuit $C\colon \{0,1\}^n \to \{0,1\}$, $a = 2^n$
- how to count them???

Approximate counting

$X = \{x < 2^n : C(x) = 1\}$ given by a circuit $C\colon \{0,1\}^n \to \{0,1\}$

Computing $s = \#X$ exactly is #P-complete

- Toda's theorem [T'89]: $\mathrm{PH} \subseteq \mathrm{P}^{\#\mathrm{P}}$
- if we could define $\#X$ by a Σ^b_i formula
  $\implies \#\mathrm{P} \subseteq \mathrm{FPH}$ and PH collapses

But we can approximate $\#X$ in PH !

- with additive error: $s - \varepsilon 2^n \le \#X \le s + \varepsilon 2^n$
- with multiplicative error: $(1-\varepsilon)s \le \#X \le (1+\varepsilon)s$
- error parameter $\varepsilon = 1/m$, m given in unary (e.g., $\varepsilon = n^{-c}$)

Goal: do that in bounded arithmetic

Weak pigeonhole principle

PHP and WPHP

Pigeonhole principle:
a pigeonholes cannot accommodate b pigeons, $b > a$

Formal statement for relations (multifunctions) R:

$$mPHP_a^b(R) = \forall y < b \, \exists x < a \, R(y,x) \\ \to \exists y < y' < b \, \exists x < a \, \bigl(R(y,x) \wedge R(y',x)\bigr)$$

$mPHP_a^{a+1}$ is an exact counting principle
not available in bounded arithmetic

Weak PHP (WPHP): b "much" larger than a
often $b = a^2$, $b = 2a$; we take $b = a(1 + 1/n)$, n unary

Theorem [PWW'88, MPW'02]: $T^2_2 \vdash \mathrm{mWPHP}(\Sigma^b_1)$

Variants of (W)PHP

Special cases where R or R^{-1} is a function:

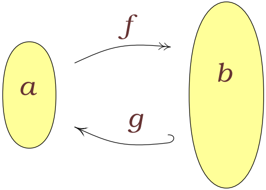

- injective (W)PHP

$$iPHP^b_a(g) = \forall y < b\ g(y) < a \\ \to \exists y < y' < b\ g(y) = g(y')$$

- surjective ("dual") (W)PHP

$$sPHP^b_a(f) = \exists y < b\, \forall x < a\ f(x) \neq y$$

- retraction-pair (W)PHP

$$rPHP^b_a(f, g) = \forall y < b\ g(y) < a \to \exists y < b\ f(g(y)) \neq y$$

For some reasons, our preferred variant is sWPHP (or rWPHP)

Counting with WPHP

Basic idea: witness that $|X| \le a$ by exhibiting

$$\text{a surjection } f\colon [a] \twoheadrightarrow X \text{ (for sWPHP)}$$

or

$$\text{an injection } f\colon X \hookrightarrow [a] \text{ (for iWPHP)}$$

Trouble: Where shall we get these functions from?

Ostensibly, WPHP is a passive counting principle:
it says something is impossible, it does not supply any counting functions

Ad hoc counting arguments using WPHP:

- [PWW'88]: T^∞_2 proves the existence of ∞ many primes
- [Pud'90]: T^∞_2 proves Ramsey's theorem

Can we develop a more general method?

Two general setups

Approximate counting with additive error

- estimate the size of $X \subseteq [2^n]$ within error $2^n/m$
  $=$ estimate $\Pr_{x<a}[x \in X]$ within error $1/m$
- Δ^b_1 sets can be counted in
  $\mathrm{APC}_1 := \mathrm{PV}_1 + \mathrm{sWPHP}(\mathrm{PV}) \subseteq \mathrm{T}^2_2$
- based on pseudorandom generators

Approximate counting with multiplicative error

- estimate the size of $X \subseteq 2^n$ within error $|X|/m$
- Σ^b_1 sets can be counted in
  $\mathrm{APC}_2 := \mathrm{T}^1_2 + \mathrm{sWPHP}(\mathrm{FP}^{\mathrm{NP}}) \subseteq \mathrm{T}^3_2$
- based on linear hashing

Counting with additive error

Main strategy

Goal: $X \subseteq [2^n]$ given by a circuit
$\rightsquigarrow$ approximate $|X|$ with error $\varepsilon 2^n$ (think $\varepsilon = 1/n^c$)

- estimate $\Pr_{x<2^n}[x \in X]$ with error ε by drawing $O(1/\varepsilon)$ independent random samples
  $\implies$ randomized poly-time algorithm
- derandomize using the Nisan–Wigderson generator

How does the theory know that the result is not just a meaningless number?

- analysis of the generator can be carried out in PV_1
  $\implies$ explicit "counting functions" for X !

Hard Boolean functions

Nisan–Wigderson generator [NW'94]

- $\mathrm{NW}_h \colon \{0,1\}^{c \log n} \to \{0,1\}^n$ computable in time $n^{O(1)}$
- "fools" circuits $C \colon \{0,1\}^n \to \{0,1\}$ of size $n^{O(1)}$

Needs access to a hard Boolean function:

- $h \colon \{0,1\}^k \to \{0,1\}$, $k = d \log n$
- h cannot be approximated by $2^{k/4}$-size circuits on $\geq \frac{1}{2} + 2^{-k/4}$ fraction of inputs

APC_1 proves: (truth tables of) such hard functions exist!

- simple counting argument
  (count circuits and error vectors)
- enumerate "easy functions" by a poly-time function
  apply sWPHP

Size comparison with error

Naïve idea: $|X| \le |Y|$ iff there is a surjection $Y \twoheadrightarrow X$

What we get from analysing NW is more complicated:

Definition: $X, Y \subseteq [2^n]$ definable sets, $\varepsilon \ge 0$

- $X \preceq_\varepsilon Y$ iff there exist $v > 0$ and a circuit

$$C \colon [v] \times (Y \mathbin{\dot\cup} [\varepsilon 2^n]) \twoheadrightarrow [v] \times X$$

- $X \approx_\varepsilon Y$ iff $X \preceq_\varepsilon Y \wedge Y \preceq_\varepsilon X$
- notation: $n \in \mathrm{Log} \iff \exists a\, n = |a|$ (i.e., 2^n exists)

Theorem [J'07]: APC_1 proves: If X is defined by a circuit and $\varepsilon^{-1} \in \mathrm{Log}$, there exists s such that $X \approx_\varepsilon [s]$.

(we also get injections going the other way)

Working with $\preceq_\varepsilon$

Basic toolbox for using approximate counting:

APC_1 proves

- $\preceq_\varepsilon$ behaves well wrt $X \cup Y$, $X \smallsetminus Y$, $X \times Y$, ...
- averaging principle
  ("if $\Pr_{x,y}[A(x,y)] \geq p$, there is x s.t. $\Pr_y[A(x,y)] \geq p$")
- Chernoff–Hoeffding inequality
- inclusion-exclusion principle

Complexity of $\preceq_\varepsilon$

In more involved arguments, we might need to

- use $\preceq_\varepsilon$ inside an induction formula,
- define other sets using $\preceq_\varepsilon$, etc.

$\Longrightarrow$ need a bound on the complexity of $\preceq_\varepsilon$

- as it stands: $X \preceq_\varepsilon Y$ is an unbounded $\exists\Pi^b_2$-formula
- $\varepsilon^{-1} \in \mathrm{Log}$ and X, Y are defined by circuits
  $\Longrightarrow$ it is Σ^b_2 by the Theorem
- for parametric families, it is $P/\,\mathrm{poly}$:
  $\varepsilon^{-1} \in \mathrm{Log}$, family $\{X_u \mid u < a\}$ of subsets of $[2^n]$
  defined by a circuit $C(u,x)\colon [a] \times [2^n] \to [2]$
  $\Longrightarrow$ a circuit $Sz\colon [a] \to [2^n]$ s.t. $X_u \approx_\varepsilon [Sz(u)]$
  $\Longrightarrow$ can appear in induction formulas in APC_1

Relativization

Relativized bounded arithmetic:

- new predicate $\alpha(x)$ w/o any specific axioms represents an arbitrary oracle A
- $\mathsf{PV}(\alpha)$-functions, $\Sigma^b_i(\alpha)$ formulas theories $\mathsf{PV}_1(\alpha)$, $\mathsf{T}^i_2(\alpha)$, $\mathsf{S}^i_2(\alpha)$, ...

The approximate counting machinery relativizes:

- $\mathsf{APC}_1(\alpha)$ can count $\mathsf{PV}(\alpha)$-definable sets ("$\mathsf{P}^A/\mathsf{poly}$")

Specializing α with Σ^b_i formulas:

- $\mathsf{APC}_{i+1} = \mathsf{T}^i_2 + \mathsf{sWPHP}(\mathsf{PV}_{i+1}) \subseteq \mathsf{T}^{i+2}_2$ can count PV_{i+1}-definable sets ("$\Delta^{\mathsf{P}}_{i+1}/\mathsf{poly}$")

Applications

- formalization of randomized complexity classes (e.g., TFRP, BPP, APP, MA, AM)
- formalization of specific randomized algorithms (e.g., Rabin–Miller primality testing algorithm, [LC'12] Edmonds and Mulmuley–Vazirani–Vazirani perfect matching algorithms)
- as a tool for working with probabilities (e.g., [Pich'15] formalization of the PCP theorem)

Counting with multiplicative error

Overview

Counting $X \subseteq [2^n]$ with error $\varepsilon|X|$ rather than $\varepsilon 2^n$

- witness that $|X| \leq s$ using linear hash functions (Sipser's coding lemma)
- equivalent to existence of suitable surjective "counting functions"
- asymmetric: no witness for $|X| \geq s$!
- can count "sparse" sets
  $\Longrightarrow$ useful for inductive counting arguments

Linear hashing

Let $X \subseteq [2^n] = \mathbb{F}_2^n$, $|X| = s$

- $A \in \mathbb{F}_2^{t \times n}$ random matrix $\implies$ $\mathrm{E}\,\#\{\{x,y\} \in \binom{X}{2} : Ax = Ay\} = 2^{-t}\binom{s}{2}$
- $2^t \ge s^2 \implies$ random A gives an injection $X \hookrightarrow [2^t] = \mathbb{F}_2^t$ $\implies$ we can distinguish sets of size $\le \sqrt{s}$ from size $> s$

Smaller gap: Sipser's coding lemma [Sip'83]

- $A \in \mathbb{F}_2^{t \times n}$ separates x from X if $Ax \ne Ay$ for all $y \in X \smallsetminus \{x\}$
- $\{A_i\}_{i<k}$ isolates X if each $x \in X$ sep. from X by some A_i

Theorem: $t = k = \lceil \log s \rceil + 1 \implies$ random $\{A_i\}_{i<k}$ isolates X

- we can distinguish sets of size $\le \frac{1}{4}s$ from size $\gtrsim s \log s$
- to get $s(1 \pm \varepsilon)$: apply to suitable Cartesian power X^d

Formalization

For $X \subseteq [2^n]$ a definable set, $\varepsilon^{-1} \in \mathrm{Log}$, $s \le 2^n$:

$X \precsim_\varepsilon s \iff \exists \{A_i\}_{i<t} \in (\mathbb{F}_2^{t\times n})^t$ that isolates $X^d \subseteq [2^{nd}]$

where $d = 12|\tilde{s}|\lceil \varepsilon^{-1} \rceil^2$, $t = |\tilde{s}^d| + 1$ for some $0 < \tilde{s} \le s$
(or: $s = 0$ and $X = \varnothing$)

Complexity:

- X Σ_1^b-definable (NP/poly) $\implies$ $X \precsim_\varepsilon s$ is Σ_2^b
- we can essentially make it Π_1^b (coNP/ poly) for a parametric family of Σ_1^b sets
  $\implies$ use in induction arguments, ...

Equivalence with surjections

Key result for manipulating $\precsim$:
equivalent to existence of PV_2 ($\mathrm{FP}^{\mathrm{NP}}/\mathrm{poly}$) surjections

Theorem [J'09]: APC_2 proves for Σ^b_1-definable X:

- $X \precsim_\varepsilon s \implies \exists$ PV_2-retraction pair $[\lfloor s(1+\varepsilon)\rfloor^d] \rightleftarrows X^d$
- if $\exists$ PV_2-surjection* $[rs^e] \twoheadrightarrow [r] \times X^e$ for $r > 0$, $e \in \mathrm{Log}$, then $X \precsim_\varepsilon s$, and

$$\Pr\bigl[\{A_i\}_{i<t} \text{ does not isolate } X^d\bigr] \preceq_0^{\Sigma^b_1} 2/3$$

where d, t are as in the definition of $\precsim$

$\preceq_0^{\Sigma^b_1}$ = counting with additive error relativized with a Σ^b_1 oracle

*with some technical condition

Properties of approximate counting

APC_2 proves (for Σ^b_1 sets):

- $\precsim_\varepsilon$ agrees with exact counting and $\preceq_\varepsilon$ where possible
- $\precsim_\varepsilon$ behaves well wrt $X \cup Y$, $X \times Y$ (even for logarithmically many terms)
- averaging principles
- approximate increasing enumeration: There are t, s s.t. $s \le t \le \lfloor s(1+\varepsilon) \rfloor$, and non-decreasing PV_2-retraction pairs

$$[t] \underset{f'}{\overset{f}{\rightleftarrows}} X \underset{g'}{\overset{g}{\rightleftarrows}} [s]$$

s.t. f, g are almost 1-to-1, and

$$\left\lfloor \tfrac{s}{t} x \right\rfloor \le g(f(x)) \le \left\lceil \tfrac{s}{t} x \right\rceil$$

Relativization

Again, everything relativizes:

- $\mathrm{APC}_2(\alpha)$ can count $\Sigma^b_1(\alpha)$-definable sets (“$\mathrm{NP}^A/\,\mathrm{poly}$”)

Specializing α with Σ^b_{i-1} formulas ($i \ge 1$):

- $\mathrm{APC}_{i+1} = \mathrm{T}^i_2 + \mathrm{sWPHP}(\mathrm{PV}_{i+1}) \subseteq \mathrm{T}^{i+2}_2$ can count Σ^b_i-definable sets ($\Sigma^{\mathrm{P}}_i/\,\mathrm{poly}$)
- (recall: with additive error, APC_{i+1} can do $\Delta^{\mathrm{P}}_{i+1}/\,\mathrm{poly}$)

Applications

- formalization of combinatorial counting arguments (e.g., Ramsey's theorem, tournament principle)
- improved collapse of hierarchies: if $T^i_2 = S^{i+1}_2$, then $T^i_2 = T_2$ proves $\Sigma^P_{i+1} \subseteq \Delta^P_{i+1}/\operatorname{poly}$ and $\mathrm{PH} = \mathrm{Bool}(\Sigma^P_{i+1})$
- [BKT'14] APC_2 proves the ordering principle
- [BKZ'15] collapse of constant-depth proofs with modular-counting gates (formalization of Valiant–Vazirani and Toda's theorem in ${\mathrm{APC}_2}^{\oplus_p \mathrm{P}}$)

References

General bounded arithmetic

- J. Krajíček: Proof complexity, Cambridge University Press, 2019
- ________: Bounded arithmetic, propositional logic, and complexity theory, Cambridge University Press, 1995
- S. R. Buss: Bounded arithmetic, Bibliopolis, Naples, 1986
- S. A. Cook, P. Nguyen: Logical foundations of proof complexity, Cambridge University Press, 2010

Approximate counting machinery

- E. Jeřábek: Dual weak pigeonhole principle, Boolean complexity, and derandomization, APAL 129 (2004), 1–37
- ________: Approximate counting in bounded arithmetic, JSL 72 (2007), 959–993
- ________: Approximate counting by hashing in bounded arithmetic, JSL 74 (2009), 829–860

References (other)

- S. R. Buss, L. A. Kołodziejczyk, N. Thapen: Fragments of approximate counting, JSL 49 (2014), 496–525
- S. R. Buss, L. A. Kołodziejczyk, K. Zdanowski: Collapsing modular counting in bounded arithmetic and constant depth propositional proofs, Trans. AMS 367 (2015), 7517–7563
- Đ. T. M. Lê and S. A. Cook: Formalizing randomized matching algorithms, LMCS 8 (2012), art. #5
- A. Maciel, T. Pitassi, A. Woods: A new proof of the weak pigeonhole principle, JCSS 64 (2002), 843–872
- N. Nisan, A. Wigderson: Hardness vs. randomness, JCSS 49 (1994), 149–167
- J. B. Paris, A. J. Wilkie, A. R. Woods: Provability of the pigeonhole principle and the existence of infinitely many primes, JSL 53 (1988), 1235–1244

References (cont'd)

- J. Pich: Logical strength of complexity theory and a formalization of the PCP theorem in bounded arithmetic, LMCS 11 (2015), art. #8
- P. Pudlák: Ramsey's theorem in bounded arithmetic, Proc. CSL'90, Springer, 1991, 308–317
- M. Sipser: A complexity theoretic approach to randomness, Proc. 15th STOC (1983), 330–335
- S. Toda: On the computational power of PP and $\oplus$P, Proc. 30th FOCS (1989), 514–519