Random Resolution Refutations

Pavel Pudlák and Neil Thapen[1]

Mathematical Institute, Academy of Sciences, Prague

Riga, 6.7.17

[1] the authors are supported by the ERC grant *FEALORA*

History and Motivation

Stefan Dantchev [unpublished]

Buss, Kołodziejczyk, Thapen [2014]

History and Motivation

Stefan Dantchev [unpublished]

Buss, Kołodziejczyk, Thapen [2014]

- Separations of fragments of Bounded Arithmetic
- The first randomized version of a proof system
- Developing lower bound methods
- Understanding what can be proved from random tautologies

Question. Random 3-DNFs (of sufficient density)

1. are tautologies,
2. can be easily generated, and
3. seem to be hard for every proof system (for not too high density).

But can we derive from them any *useful* tautology?

Question. Random 3-DNFs (of sufficient density)

1. are tautologies,
2. can be easily generated, and
3. seem to be hard for every proof system (for not too high density).

But can we derive from them any *useful* tautology?

Corollary (of our results)

With high probability for a random 3-DNF Φ (of sufficient density), there is no bounded depth Frege proof of $\Phi \to PHP$ of subexponential size.

Overview

1. equivalent definitions
2. upper bounds
3. lower bounds
4. generalization to bounded depth Frege proofs
5. problems

Definitions

Definition

An ϵ-*random resolution distribution*, or ϵ-*RR distribution*, of F is a probability distribution Δ on pairs $(B_i, \Pi_i)_{i \sim \Delta}$ such that

1. for each $i \in \Delta$, B_i is a CNF in variables $x_1, \ldots, x_n$ and Π_i is a resolution refutation of $F \wedge B_i$
2. for every $\alpha \in \{0,1\}^n$, $\Pr_{i \sim \Delta}[B_i \text{ is satisfied by } \alpha] \geq 1 - \epsilon$.

The *size* and the *width* of Δ are defined respectively as the maximum size and maximum width of the refutations Π_i (if these maxima exist).

- RR is sound as a refutational system, in the sense that if F has an ϵ-RR distribution then F is unsatisfiable.

  Proof: consider any assignment $\alpha \in \{0,1\}^n$. Since $\epsilon < 1$, there is at least one pair (B_i, Π_i) such that α satisfies B_i and Π_i is a resolution refutation of $F \wedge B_i$. So α cannot also satisfy F, by the soundness of resolution.

- RR is complete, since resolution is complete.
- RR is not a propositional proof system in the sense of Cook and Reckhow because it is defined by a semantic condition.
- The error ϵ can be reduced with only moderate increase of the proofs, so we can w.l.o.g. assume $\epsilon = 1/2$.

Definition

Let Δ be a probability distribution on $\{0,1\}^n$. An *(ϵ, Δ)-random resolution refutation*, or *(ϵ, Δ)-RR refutation*, of F is a pair (B, Π) such that

1. B is a CNF in variables $x_1, \ldots, x_n$ and Π is a resolution refutation of $F \wedge B$
2. $\Pr_{\alpha \sim \Delta}[B[\alpha] = 1] \geq 1 - \epsilon$.

Definition

Let Δ be a probability distribution on $\{0,1\}^n$. An *(ϵ, Δ)-random resolution refutation*, or *(ϵ, Δ)-RR refutation*, of F is a pair (B, Π) such that

1. B is a CNF in variables $x_1, \ldots, x_n$ and Π is a resolution refutation of $F \wedge B$
2. $\Pr_{\alpha \sim \Delta}[B[\alpha] = 1] \geq 1 - \epsilon$.

If an (ϵ, Δ)-RR refutation exists for *all* distributions Δ, then this is equivalent to the existence of an ϵ-RR distribution.

Let $\mathcal{P}$ be a class of resolution refutations, e.g., refutations of width w and size s for some w and s.

Proposition

The following are equivalent.

1. *F has an ϵ-RR distribution of refutations from $\mathcal{P}$.*
2. *F has an (ϵ, Δ)-RR refutation from $\mathcal{P}$ for every distribution Δ on $\{0,1\}^n$.*

Proof.

Consider a zero-sum game between two players, Prover and Adversary:

- Prover picks a pair (B, Π) such that $\Pi \in \mathcal{P}$, B is a CNF, and Π is a refutation of $F \wedge B$,
- Adversary picks an assignment α.

The payoff is $B[\alpha]$, i.e., Prover gets 1 if α satisfies B and 0 otherwise.

Then

- definition 1 says: Prover has a mixed strategy to achieve a payoff of at least $1 - \epsilon$, and
- definition 2 says: Adversary does not have a mixed strategy to achieve a payoff less than $1 - \epsilon$.

By the minimax theorem these statements are equivalent. □

Definition

Let $A \subseteq \{0,1\}^n$ be a nonempty set of truth assignments. We say that a formula C is a *semantic consequence over* A of formulas $C_1, \ldots, C_r$, if every assignment in A that satisfies $C_1, \ldots, C_r$ also satisfies C.

A *semantic resolution refutation of* F *over* A is a sequence Π of *clauses*, ending with the empty clause, in which every clause either belongs to F or is a semantic consequence over A of at most two earlier clauses.

Definition

Let $A \subseteq \{0,1\}^n$ be a nonempty set of truth assignments. We say that a formula C is a *semantic consequence over* A of formulas $C_1, \ldots, C_r$, if every assignment in A that satisfies $C_1, \ldots, C_r$ also satisfies C.

A *semantic resolution refutation of* F *over* A is a sequence Π of *clauses*, ending with the empty clause, in which every clause either belongs to F or is a semantic consequence over A of at most two earlier clauses.

Definition

Let Δ be a probability distribution on $\{0,1\}^n$. An *(ϵ, Δ)-semantic refutation* of F is a pair (A, Π) such that

1. Π is a semantic refutation of F over A, and
2. $\Pr_{\alpha \sim \Delta}[\alpha \in A] \geq 1 - \epsilon$.

Definition

Let $A \subseteq \{0,1\}^n$ be a nonempty set of truth assignments. We say that a formula C is a *semantic consequence over* A of formulas $C_1, \ldots, C_r$, if every assignment in A that satisfies $C_1, \ldots, C_r$ also satisfies C.

A *semantic resolution refutation of* F *over* A is a sequence Π of *clauses*, ending with the empty clause, in which every clause either belongs to F or is a semantic consequence over A of at most two earlier clauses.

Definition

Let Δ be a probability distribution on $\{0,1\}^n$. An *(ϵ, Δ)-semantic refutation* of F is a pair (A, Π) such that

1. Π is a semantic refutation of F over A, and
2. $\Pr_{\alpha \sim \Delta}[\alpha \in A] \geq 1 - \epsilon$.

Note: *no auxiliary formulas!*

Proposition

1. *If F has an (ϵ, Δ)-RR refutation of width w and size s, then it also has an (ϵ, Δ)-semantic resolution refutation of width $\leq w$ and size $\leq s$.*
2. *If F has has an (ϵ, Δ)-semantic refutation of width w and size s, then it also has an (ϵ, Δ)-RR refutation of width $O(w)$ and size at most $O(sw^2)$.*

The strength of RR

- A random 3-CNF with n variables and $64n$ clauses has a $1/2$-RR distribution of constant width and constant size with probability exponentially close to 1.

The strength of RR

- A random 3-CNF with n variables and $64n$ clauses has a $1/2$-RR distribution of constant width and constant size with probability exponentially close to 1.
- The *retraction weak pigeonhole principle* that asserts that there is no pair of functions $f : [2n] \to [n]$ and $g : [n] \to [2n]$ such that $g(f(x)) = x$ for all $x < n$ has a narrow $1/2$-RR distribution.

The strength of RR

- A random 3-CNF with n variables and $64n$ clauses has a $1/2$-RR distribution of constant width and constant size with probability exponentially close to 1.
- The *retraction weak pigeonhole principle* that asserts that there is no pair of functions $f : [2n] \to [n]$ and $g : [n] \to [2n]$ such that $g(f(x)) = x$ for all $x < n$ has a narrow $1/2$-RR distribution.
- If $\mathbf{P} \neq \mathbf{NP}$, then $1/2$-RR cannot be polynomially simulated by any Cook-Reckhow refutation system. In particular, $1/2$-RR is not itself a Cook-Reckhow refutation system if $\mathbf{P} \neq \mathbf{NP}$.

Lemma

Let $F := C_1 \wedge \cdots \wedge C_m$ be a k-CNF formula such that for every assignment α the number of clauses that are satisfied by α is $\leq \delta m$ for some constant $0 < \delta < 1$. Then F has a δ-RR distribution of size $2k$ which can be constructed in polynomial time.

Proof.

The distribution is defined by:

1. pick $i \in [m]$ randomly
2. let B_i (the auxiliary formula) be $\neg C_i$ and Π_i the proof of $\bot$ from B_i and C_i.

□

Lemma

Let $F := C_1 \wedge \cdots \wedge C_m$ be a k-CNF formula such that for every assignment α the number of clauses that are satisfied by α is $\leq \delta m$ for some constant $0 < \delta < 1$. Then F has a δ-RR distribution of size $2k$ which can be constructed in polynomial time.

Proof.

The distribution is defined by:

1. pick $i \in [m]$ randomly
2. let B_i (the auxiliary formula) be $\neg C_i$ and Π_i the proof of $\bot$ from B_i and C_i.

□

For random 3-CNFs, $\delta = 7/8$. To get $\epsilon \leq 1/2$, take random sixtuples $i_1, \ldots, i_6 \in [m]$ and the CNFs equivalent to

$$\neg C_{i_1} \vee \cdots \vee \neg C_{i_6}$$

The weakness of RR

Theorem
PHP_n has no $1/2$-RR distribution of size $O(2^{n^{1/12}})$.

Theorem
The formula $CPLS_n^2$ (will be defined later) does not have a $1/2$-RR distribution of size $O(2^{n^{1/17}})$.

CPLS_n^2 has polynomial size Res(2) proofs.

Bounded depth Frege *refutation* systems

1. clauses, $\bigvee$-formulas – Resolution $=$ 1-Frege system
2. DNFs – $\bigvee\bigwedge$-formulas – 2-Frege system
3. etc.

Problem

Does there exist a CNF contradiction refutable in some d-Frege system, $d > 2$, by quasipolynomial size refutation that does not have such a 2-Frege refutation.[2]

[2] Open even for 1.5-Frege ($=$Res(log)).

Bounded depth Frege *refutation* systems

1. clauses, $\bigvee$-formulas – Resolution $=$ 1-Frege system
2. DNFs – $\bigvee\bigwedge$-formulas – 2-Frege system
3. etc.

Problem

Does there exist a CNF contradiction refutable in some d-Frege system, $d > 2$, by quasipolynomial size refutation that does not have such a 2-Frege refutation.[2]

WPHP separates Resolution from 2-Frege.

[2] Open even for 1.5-Frege (=Res(log)).

Bounded depth Frege *refutation* systems

1. clauses, $\bigvee$-formulas – Resolution $=$ 1-Frege system
2. DNFs – $\bigvee\bigwedge$-formulas – 2-Frege system
3. etc.

Problem

Does there exist a CNF contradiction refutable in some d-Frege system, $d > 2$, by quasipolynomial size refutation that does not have such a 2-Frege refutation.[2]

WPHP separates Resolution from 2-Frege.

Our result:

CPLS^2 separates RR from 2-Frege.

[2] Open even for 1.5-Frege (=Res(log)).

Definition

The formula $CPLS_{a,b,c}$ consists of the following three sets of clauses:

1. For each $y < c$, the clause $\neg G_0(0, y)$
2. For each $i < a - 1$, each pair $x, x' < b$ and each $y < c$, the clause

$$f_i(x) = x' \wedge G_{i+1}(x', y) \rightarrow G_i(x, y)$$

3. For each $x < b$ and each $y < c$, the clause

$$u(x) = y \rightarrow G_{a-1}(x, y).$$

Definition

The formula $CPLS_{a,b,c}$ consists of the following three sets of clauses:

1. For each $y < c$, the clause $\neg G_0(0, y)$
2. For each $i < a - 1$, each pair $x, x' < b$ and each $y < c$, the clause

$$f_i(x) = x' \wedge G_{i+1}(x', y) \rightarrow G_i(x, y)$$

3. For each $x < b$ and each $y < c$, the clause

$$u(x) = y \rightarrow G_{a-1}(x, y).$$

The formula $CPLS^2$ is a variant of $CPLS$ where for each i, x, y, instead of the single variable $G_i(x, y)$ it has two variables $G_i^0(x, y)$ and $G_i^1(x, y)$. To express that colour y is present at node (i, x) we now use the conjunction $G_i^0(x, y) \wedge G_i^1(x, y)$.

Lower bound for the pigeonhole principle in RR

PHP_n has variables p_{ij} for $i \in U$ and $j \in V$ and consists of clauses

1. $\bigvee_{j \in V} p_{ij}$ for all $i \in U$
2. $\neg p_{ij} \vee \neg p_{i'j}$ for all $i, i' \in U$ with $i \neq i'$ and all $j \in V$.

1. width reduction

For a clause C, we define $w_{\mathrm{ec}}(C)$, the *edge covering width* or *ec-width* of C, to be the smallest size of a set $W \subseteq U \cup V$ that intersects all pairs $\{i, j\}$ mentioned in C. Formally,

$$w_{\mathrm{ec}}(C) := \min\{|W| \mid \forall i \in U, j \in V, (p_{ij} \in C \vee \neg p_{ij} \in C) \to (i \in W \vee j \in W)\}.$$

If Φ is a CNF formula, then $w_{\mathrm{ec}}(\Phi)$ is the maximum of the ec-widths of its clauses.

We will denote by $\mathcal{R}_m$ the set of partial matchings of size $n - m$ equipped with the uniform distribution.

1. width reduction

For a clause C, we define $w_{\mathrm{ec}}(C)$, the *edge covering width* or *ec-width* of C, to be the smallest size of a set $W \subseteq U \cup V$ that intersects all pairs $\{i, j\}$ mentioned in C. Formally,

$$w_{\mathrm{ec}}(C) := \min\{|W| \mid \forall i \in U, j \in V, (p_{ij} \in C \vee \neg p_{ij} \in C) \to (i \in W \vee j \in W)\}.$$

If Φ is a CNF formula, then $w_{\mathrm{ec}}(\Phi)$ is the maximum of the ec-widths of its clauses.

We will denote by $\mathcal{R}_m$ the set of partial matchings of size $n - m$ equipped with the uniform distribution.

Lemma

There exist constants $c > 0$ and $0 < d < 1$ such that for every clause C and every $1 \leq \ell \leq n^{1/2}$,

$$\Pr[w_{\mathrm{ec}}(C^\rho) > \ell] \leq d^\ell,$$

where the probability is over $\rho \sim \mathcal{R}_{\lfloor cn^{1/4} \rfloor}$.

2. the fixing lemma

Ideally, we would like to eliminate the auxiliary formula B by finding $\rho \in \mathcal{R}_m$ such that $B^\rho \equiv 1$. This is not always possible.

2. the fixing lemma

Ideally, we would like to eliminate the auxiliary formula B by finding $\rho \in \mathcal{R}_m$ such that $B^\rho \equiv 1$. This is not always possible.

Lemma

Let B be a CNF formula such that $w_{\mathrm{ec}}(B) \leq \ell$ and

$$\Pr_{\rho \sim \mathcal{R}_m}[B^\rho = 0] \leq 1/2$$

where $\ell < m < n$. Suppose that

$$\frac{\ell m(m-1)}{n-m+1} < \frac{1}{2}.$$

Then there exists a $\rho \in \mathcal{R}_m$ such that there is no extension $\sigma \supseteq \rho$ to a partial matching such that $B^\sigma = 0$.

Random Bounded Depth Frege proofs

Definition

Let $d \geq 3$ be constant. A *random d-Frege proof* of a DNF tautology F is a depth d Frege proof of

$$B \to F,$$

where B is a depth 3 formula $\bigvee_i B_i$, where B_is are CNFs, such that for every assignment α at least $1/2$ of B_is are satisfied.

Random Bounded Depth Frege proofs

Definition

Let $d \geq 3$ be constant. A *random d-Frege proof* of a DNF tautology F is a depth d Frege proof of

$$B \to F,$$

where B is a depth 3 formula $\bigvee_i B_i$, where B_is are CNFs, such that for every assignment α at least $1/2$ of B_is are satisfied.

Theorem

For every d, there is $\epsilon_d > 0$ *such that every random d-Frege proof of* PHP_n *has size* $\geq 2^{n^{\epsilon_d}}$.

Corollary

With high probability for a random 3-DNF Φ with n variables and $64n$ clauses, there is no bounded depth Frege proof of $\Phi \to PHP$ of subexponential size.

Problems

- find natural random version for more proof systems
- cutting planes, ResLin, and, in particular, for the Nullstellensatz system
- randomized SAT algorithms and random resolution

THANK YOU