Random resolution

Neil Thapen

Institute of Mathematics
Czech Academy of Sciences

Joint work with Pavel Pudlák

Outline

Motivation

What if a resolution proof is allowed to make errors?

[Dantchev]

Motivation

What if a resolution proof is allowed to make errors?

[Dantchev]

There are bounded arithmetic systems that can formalize reasoning with approximate counting or probabilities.

Are there corresponding propositional systems?

[Jeřábek '05, Buss-Kołodziejczyck-T '14, Atserias-T '14]

Motivation

What if a resolution proof is allowed to make errors?

[Dantchev]

There are bounded arithmetic systems that can formalize reasoning with approximate counting or probabilities.

Are there corresponding propositional systems?

[Jeřábek '05, Buss-Kołodziejczyck-T '14, Atserias-T '14]

(Are there CNFs separating constant depth Frege from R(log)?)

Random resolution distribution

Let F be a CNF in variables $x_1, \ldots, x_n$ and let $0 < \epsilon < 1$.

Definition

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

Random resolution distribution

Let F be a CNF in variables $x_1, \ldots, x_n$ and let $0 < \epsilon < 1$.

Definition

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

1. each B_i is a CNF in variables $x_1, \ldots, x_n$

Random resolution distribution

Let F be a CNF in variables $x_1, \ldots, x_n$ and let $0 < \epsilon < 1$.

Definition

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

1. each B_i is a CNF in variables $x_1, \ldots, x_n$
2. each Π_i is a resolution refutation of $B_i \wedge F$

Random resolution distribution

Let F be a CNF in variables $x_1, \ldots, x_n$ and let $0 < \epsilon < 1$.

Definition

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

1. each B_i is a CNF in variables $x_1, \ldots, x_n$
2. each Π_i is a resolution refutation of $B_i \wedge F$
3. for all assignments $\alpha \in \{0,1\}^n$,

$$\Pr_{i \sim \mathcal{D}}[B_i \text{ is satisfied by } \alpha] \geq 1 - \epsilon.$$

Observations

Observations

The system is sound and complete.

Observations

The system is sound and complete.

Being a proof is “semantic”, and not checkable in polynomial time.

Observations

The system is sound and complete.

Being a proof is "semantic", and not checkable in polynomial time.

We usually take $\epsilon = 1/2$. The exact value is not important.

Observations

The system is sound and complete.

Being a proof is "semantic", and not checkable in polynomial time.

We usually take $\epsilon = 1/2$. The exact value is not important.

We define *size* and *width* as the maximum over the refutations Π_i.

Observations

The system is sound and complete.

Being a proof is "semantic", and not checkable in polynomial time.

We usually take $\epsilon = 1/2$. The exact value is not important.

We define *size* and *width* as the maximum over the refutations Π_i.

We may assume the number of pairs (B_i, Π_i) in the sample space is small, roughly n/ϵ.

Observations

The system is sound and complete.

Being a proof is "semantic", and not checkable in polynomial time.

We usually take $\epsilon = 1/2$. The exact value is not important.

We define *size* and *width* as the maximum over the refutations Π_i.

We may assume the number of pairs (B_i, Π_i) in the sample space is small, roughly n/ϵ.

I will write "RR" instead of "random resolution."

Definition

Upper bounds

Lower bounds

Fixing lemma

Example 1 Random 3-CNFs

Let F be a randomly chosen 3-CNF with $64n$ clauses.
For most F, every assignment falsifies at least $1/16$ of the clauses.

Example 1 Random 3-CNFs

Let F be a randomly chosen 3-CNF with $64n$ clauses.
For most F, every assignment falsifies at least $1/16$ of the clauses.

Pick a clause $C = \{x, y, z\}$ from F at random.
Let B be the CNF $\{\neg x\} \wedge \{\neg y\} \wedge \{\neg z\}$.

Example 1 Random 3-CNFs

Let F be a randomly chosen 3-CNF with $64n$ clauses.
For most F, every assignment falsifies at least $1/16$ of the clauses.

Pick a clause $C = \{x, y, z\}$ from F at random.
Let B be the CNF $\{\neg x\} \wedge \{\neg y\} \wedge \{\neg z\}$.
Let Π be the resolution refutation of $B \wedge F$ (it has three steps).

Example 1 Random 3-CNFs

Let F be a randomly chosen 3-CNF with $64n$ clauses.
For most F, every assignment falsifies at least $1/16$ of the clauses.

Pick a clause $C = \{x, y, z\}$ from F at random.
Let B be the CNF $\{\neg x\} \wedge \{\neg y\} \wedge \{\neg z\}$.
Let Π be the resolution refutation of $B \wedge F$ (it has three steps).

For every assignment α, $\Pr[\alpha \models B] > 1/16$.

Thus the random pair (B, Π) is a $15/16$-RR distribution for F.

Example 1 Random 3-CNFs

To improve the error ϵ, pick two clauses C, D from F at random.

Let B be a CNF semantically equivalent to $\neg C \vee \neg D$:

$$B = \bigwedge \{\neg p \vee \neg q : p \in C, q \in D\}.$$

Example 1 Random 3-CNFs

To improve the error ϵ, pick two clauses C, D from F at random.

Let B be a CNF semantically equivalent to $\neg C \vee \neg D$:

$$B = \bigwedge \{\neg p \vee \neg q : p \in C, q \in D\}.$$

For every assignment α, $\Pr[\alpha \vDash B] > 1 - (15/16)^2$.
There is a short resolution refutation of $B \wedge F$.

Hence (B, Π) is a $(15/16)^2$-RR distribution for F.

Example 1 Random 3-CNFs

To improve the error ϵ, pick two clauses C, D from F at random.

Let B be a CNF semantically equivalent to $\neg C \vee \neg D$:

$$B = \bigwedge \{\neg p \vee \neg q : p \in C, q \in D\}.$$

For every assignment α, $\Pr[\alpha \vDash B] > 1 - (15/16)^2$.
There is a short resolution refutation of $B \wedge F$.

Hence (B, Π) is a $(15/16)^2$-RR distribution for F.

Theorem
Random 3-CNFs have 1/2-RR distributions of constant size.

Observation

Theorem

Unless P=NP, 1/2-RR is not a Cook-Reckhow proof system (and cannot be polynomially simulated by one).

Observation

Theorem

Unless P=NP, 1/2-RR is not a Cook-Reckhow proof system (and cannot be polynomially simulated by one).

Proof

By the construction in Example 1 and the PCP theorem.

Example 2 Retraction WPHP

rWPHP$_n$ is a CNF with variables for functions f and g, encoded by their bit-graphs. It asserts that

1. $f : 2n \to n$ and $g : n \to 2n$
2. For all $x < 2n$, $g(f(x)) = x$.

Example 2 Retraction WPHP

rWPHP_n is a CNF with variables for functions f and g, encoded by their bit-graphs. It asserts that

1. $f : 2n \to n$ and $g : n \to 2n$
2. For all $x < 2n$, $g(f(x)) = x$.

Theorem

rWPHP_n has polynomial size, polylog width 1/2-RR distributions.

Proposition

rWPHP_n needs width n and exponential size to refute in resolution.

Another characterization - semantic resolution

For $\mathcal{A} \subseteq \{0,1\}^n$, a *semantic resolution refutation over* $\mathcal{A}$ can use any binary rules, as long as they are sound over assignments in $\mathcal{A}$.

Another characterization - semantic resolution

For $\mathcal{A} \subseteq \{0,1\}^n$, a *semantic resolution refutation over* $\mathcal{A}$ can use any binary rules, as long as they are sound over assignments in $\mathcal{A}$.

Let Δ be a probability distribution on $\{0,1\}^n$.

Definition

An (ϵ, Δ)-semantic resolution refutation of F is a pair $(\mathcal{A}, \Pi)$ such that

Another characterization - semantic resolution

For $\mathcal{A} \subseteq \{0,1\}^n$, a *semantic resolution refutation over* $\mathcal{A}$ can use any binary rules, as long as they are sound over assignments in $\mathcal{A}$.

Let Δ be a probability distribution on $\{0,1\}^n$.

Definition

An (ϵ, Δ)-semantic resolution refutation of F is a pair $(\mathcal{A}, \Pi)$ such that

1. $\mathcal{A} \subseteq \{0,1\}^n$ and $\mathrm{Pr}_{\alpha \sim \Delta}[\alpha \in \mathcal{A}] \geq 1 - \epsilon$

Another characterization - semantic resolution

For $\mathcal{A} \subseteq \{0,1\}^n$, a *semantic resolution refutation over* $\mathcal{A}$ can use any binary rules, as long as they are sound over assignments in $\mathcal{A}$.

Let Δ be a probability distribution on $\{0,1\}^n$.

Definition

An (ϵ, Δ)-semantic resolution refutation of F is a pair $(\mathcal{A}, \Pi)$ such that

1. $\mathcal{A} \subseteq \{0,1\}^n$ and $\Pr_{\alpha \sim \Delta}[\alpha \in \mathcal{A}] \geq 1 - \epsilon$
2. Π is a semantic refutation of F over $\mathcal{A}$.

Another characterization - semantic resolution

For $\mathcal{A} \subseteq \{0,1\}^n$, a *semantic resolution refutation over* $\mathcal{A}$ can use any binary rules, as long as they are sound over assignments in $\mathcal{A}$.

Let Δ be a probability distribution on $\{0,1\}^n$.

Definition

An (ϵ, Δ)-semantic resolution refutation of F is a pair $(\mathcal{A}, \Pi)$ such that

1. $\mathcal{A} \subseteq \{0,1\}^n$ and $\Pr_{\alpha \sim \Delta}[\alpha \in \mathcal{A}] \geq 1 - \epsilon$
2. Π is a semantic refutation of F over $\mathcal{A}$.

Lemma

The following are equivalent.

1. F has a small ϵ-RR distribution.
2. F has small (ϵ, Δ)-semantic resolution refutations for *every* Δ.

Definition

Upper bounds

Lower bounds

Fixing lemma

Main results - width lower bounds

Main results - width lower bounds

CPLS_n is a family of CNF contradictions.

It has polylogarithmic initial width.

Theorem

- CPLS_n has polynomial size resolution refutations.
- CPLS_n has no $1/2$-RR distribution of width $n^{1/20}$.

This answers the main question about the system from [BKT '14].

Main results - size lower bounds

Main results - size lower bounds

CPLS_n^2 is a family of CNF contradictions.

It has polylogarithmic initial width.

Theorem

- CPLS_n^2 has polynomial size Res(2) refutations.
- CPLS_n^2 has no 1/2-RR distribution of subexponential size.

Main results - size lower bounds

CPLS_n^2 is a family of CNF contradictions.

It has polylogarithmic initial width.

Theorem

- CPLS_n^2 has polynomial size Res(2) refutations.
- CPLS_n^2 has no 1/2-RR distribution of subexponential size.

This is proved by a reduction to the width lower bound for CPLS_n.

Main results - size lower bounds

CPLS^2_n is a family of CNF contradictions.

It has polylogarithmic initial width.

Theorem

- CPLS^2_n has polynomial size Res(2) refutations.
- CPLS^2_n has no 1/2-RR distribution of subexponential size.

This is proved by a reduction to the width lower bound for CPLS_n.

Theorem

- PHP_n has no 1/2-RR distribution of subexponential size.

The CNF contradiction $\mathrm{CPLS}_{a,b,c}$

$\mathrm{CPLS}_{a,b,c}$ is a combinatorial principle as hard as the reflection principle for resolution.

The CNF contradiction $\mathrm{CPLS}_{a,b,c}$

$\mathrm{CPLS}_{a,b,c}$ is a combinatorial principle as hard as the reflection principle for resolution.

Pairs $(i, x) \in [a] \times [b]$ represent nodes on level i.
Numbers $y < c$ represent colours.
Each node has some, maybe empty, set of colours.

The CNF contradiction $\mathrm{CPLS}_{a,b,c}$

$\mathrm{CPLS}_{a,b,c}$ is a combinatorial principle as hard as the reflection principle for resolution.

Pairs $(i, x) \in [a] \times [b]$ represent nodes on level i.
Numbers $y < c$ represent colours.
Each node has some, maybe empty, set of colours.

We have variables $G_i(x, y)$ for "colour y is present at node (i, x)" and variables for the bit-graphs of functions $f_i : b \to b$ and $u : b \to c$.

The CNF contradiction $\mathrm{CPLS}_{a,b,c}$

Ranges of indices are $i < a$, $x, x' < b$, $y < c$. Axioms are

1. $\bigwedge_{y} \neg G_0(0, y)$
2. $\bigwedge_{i,y,x,x'} f_i(x) = x' \wedge G_{i+1}(x', y) \to G_i(x, y)$
3. $\bigwedge_{x,y} u(x) = y \to G_{a-1}(x, y).$

The CNF contradiction $\mathrm{CPLS}_{a,b,c}$

Ranges of indices are $i < a$, $x, x' < b$, $y < c$. Axioms are

1. $\bigwedge_{y} \neg G_0(0, y)$
2. $\bigwedge_{i,y,x,x'} f_i(x) = x' \wedge G_{i+1}(x', y) \to G_i(x, y)$
3. $\bigwedge_{x,y} u(x) = y \to G_{a-1}(x, y).$

1. Node $(0, 0)$ has no colours

The CNF contradiction $\mathrm{CPLS}_{a,b,c}$

Ranges of indices are $i < a$, $x, x' < b$, $y < c$. Axioms are

1. $\bigwedge_{y} \neg G_0(0, y)$
2. $\bigwedge_{i,y,x,x'} f_i(x) = x' \wedge G_{i+1}(x', y) \to G_i(x, y)$
3. $\bigwedge_{x,y} u(x) = y \to G_{a-1}(x, y).$

1. Node $(0, 0)$ has no colours
2. Every colour y present at node $(i + 1, f_i(x))$ is also present at node (i, x)

The CNF contradiction $\mathrm{CPLS}_{a,b,c}$

Ranges of indices are $i < a$, $x, x' < b$, $y < c$. Axioms are

1. $\bigwedge_{y} \neg G_0(0, y)$
2. $\bigwedge_{i,y,x,x'} f_i(x) = x' \wedge G_{i+1}(x', y) \to G_i(x, y)$
3. $\bigwedge_{x,y} u(x) = y \to G_{a-1}(x, y)$.

1. Node $(0, 0)$ has no colours
2. Every colour y present at node $(i + 1, f_i(x))$ is also present at node (i, x)
3. Some colour $u(x)$ is present at node $(a - 1, x)$.

Width lower bound for CPLS

What tools do we have to prove lower bounds?

Width lower bound for CPLS

What tools do we have to prove lower bounds?

We will need one more characterization of RR.

Random resolution refutations

Let $\mathcal{R}$ be a distribution on partial assignments $\{0, 1, *\}^n$.

Definition

An $(\epsilon, \mathcal{R})$-RR refutation of F is a pair $(\mathcal{B}, \Pi)$ such that

Random resolution refutations

Let $\mathcal{R}$ be a distribution on partial assignments $\{0, 1, *\}^n$.

Definition

An $(\epsilon, \mathcal{R})$-RR refutation of F is a pair (B, Π) such that

1. B is a CNF in variables $x_1, \ldots, x_n$

Random resolution refutations

Let $\mathcal{R}$ be a distribution on partial assignments $\{0, 1, *\}^n$.

Definition

An $(\epsilon, \mathcal{R})$-RR refutation of F is a pair (B, Π) such that

1. B is a CNF in variables $x_1, \ldots, x_n$
2. Π is a resolution refutation of $B \wedge F$

Random resolution refutations

Let $\mathcal{R}$ be a distribution on partial assignments $\{0, 1, *\}^n$.

Definition

An $(\epsilon, \mathcal{R})$-RR refutation of F is a pair (B, Π) such that

1. B is a CNF in variables $x_1, \ldots, x_n$
2. Π is a resolution refutation of $B \wedge F$
3. $\Pr_{\rho \sim \mathcal{R}}[B \text{ is falsified by } \rho] \leq \epsilon$.

Random resolution refutations

Let $\mathcal{R}$ be a distribution on partial assignments $\{0, 1, *\}^n$.

Definition

An $(\epsilon, \mathcal{R})$-RR refutation of F is a pair (B, Π) such that

1. B is a CNF in variables $x_1, \ldots, x_n$
2. Π is a resolution refutation of $B \wedge F$
3. $\Pr_{\rho \sim \mathcal{R}}[B \text{ is falsified by } \rho] \leq \epsilon$.

Lemma

The following are equivalent.

1. F has an ϵ-RR distribution of width w and size s.
2. F has an $(\epsilon, \mathcal{R})$-RR refutation of width w and size s for *every* distribution $\mathcal{R}$.

Width lower bound for CPLS

Fix parameters $a = b = n$, $c = \lfloor n^{1/7} \rfloor$ for CPLS. Let $w = \lfloor n^{1/8} \rfloor$.

Suppose that CPLS has a $1/2$-RR distribution of width w.

Width lower bound for CPLS

Fix parameters $a = b = n$, $c = \lfloor n^{1/7} \rfloor$ for CPLS. Let $w = \lfloor n^{1/8} \rfloor$.

Suppose that CPLS has a $1/2$-RR distribution of width w.

Define a distribution $\mathcal{R}$ of random restrictions ρ:

Width lower bound for CPLS

Fix parameters $a = b = n$, $c = \lfloor n^{1/7} \rfloor$ for CPLS. Let $w = \lfloor n^{1/8} \rfloor$.

Suppose that CPLS has a $1/2$-RR distribution of width w.

Define a distribution $\mathcal{R}$ of random restrictions ρ:

1. Set $f_i(x)$ at almost all nodes (i, x) in the style of PHP, so that on each row f_i is a partial permutation.

   This decomposes the table into many disjoint f-paths.

Width lower bound for CPLS

Fix parameters $a = b = n$, $c = \lfloor n^{1/7} \rfloor$ for CPLS. Let $w = \lfloor n^{1/8} \rfloor$.

Suppose that CPLS has a $1/2$-RR distribution of width w.

Define a distribution $\mathcal{R}$ of random restrictions ρ:

1. Set $f_i(x)$ at almost all nodes (i, x) in the style of PHP, so that on each row f_i is a partial permutation.

   This decomposes the table into many disjoint f-paths.
2. Set all colours on the path touching $(0, 0)$ to 0.

Width lower bound for CPLS

Fix parameters $a = b = n$, $c = \lfloor n^{1/7} \rfloor$ for CPLS. Let $w = \lfloor n^{1/8} \rfloor$.

Suppose that CPLS has a $1/2$-RR distribution of width w.

Define a distribution $\mathcal{R}$ of random restrictions ρ:

1. Set $f_i(x)$ at almost all nodes (i, x) in the style of PHP, so that on each row f_i is a partial permutation.

   This decomposes the table into many disjoint f-paths.
2. Set all colours on the path touching $(0, 0)$ to 0.
3. For almost every other path π, assign to all nodes on π a unique colour y_π. Otherwise leave all colours on π as $*$. Restrict u consistently with this.

Width lower bound for CPLS

Fix parameters $a = b = n$, $c = \lfloor n^{1/7} \rfloor$ for CPLS. Let $w = \lfloor n^{1/8} \rfloor$.

Suppose that CPLS has a $1/2$-RR distribution of width w.

Define a distribution $\mathcal{R}$ of random restrictions ρ:

1. Set $f_i(x)$ at almost all nodes (i, x) in the style of PHP, so that on each row f_i is a partial permutation.

   This decomposes the table into many disjoint f-paths.
2. Set all colours on the path touching $(0, 0)$ to 0.
3. For almost every other path π, assign to all nodes on π a unique colour y_π. Otherwise leave all colours on π as $*$. Restrict u consistently with this.

By the lemma, CPLS has a $(1/2, \mathcal{R})$-RR refutation of width w.

Width lower bound for CPLS

CPLS has a $(1/2, \mathcal{R})$-RR refutation of width w.

Width lower bound for CPLS

CPLS has a $(1/2, \mathcal{R})$-RR refutation of width w.

So there exist

1. a w-CNF B such that $\Pr[B \restriction \rho = 0] \leq 1/2$
2. a width w resolution refutation Π of $B \wedge \text{CPLS}$.

Width lower bound for CPLS

CPLS has a $(1/2, \mathcal{R})$-RR refutation of width w.

So there exist

1. a w-CNF B such that $\Pr[B \restriction \rho = 0] \leq 1/2$
2. a width w resolution refutation Π of $B \wedge$ CPLS.

Definition

A *legal* restriction is one in which no f-path is coloured in a way immediately inconsistent with CPLS (plus other nice properties).

Width lower bound for CPLS

CPLS has a $(1/2, \mathcal{R})$-RR refutation of width w.

So there exist

1. a w-CNF B such that $\Pr[B \restriction \rho = 0] \leq 1/2$
2. a width w resolution refutation Π of $B \wedge$ CPLS.

Definition

A *legal* restriction is one in which no f-path is coloured in a way immediately inconsistent with CPLS (plus other nice properties).

A legal restriction, if it is not too big, represents a safe move for the Adversary in the Prover-Adversary game on CPLS.

Width lower bound for CPLS

There exist

1. a w-CNF B such that $\Pr[B \restriction \rho = 0] \leq 1/2$
2. a width w resolution refutation Π of $B \wedge \text{CPLS}$.

Width lower bound for CPLS

There exist

1. a w-CNF B such that $\Pr[B \restriction \rho = 0] \leq 1/2$
2. a width w resolution refutation Π of $B \wedge \text{CPLS}$.

Fixing Lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

That is, with polynomially high probability, $B \restriction \rho$ is forced to be either false or unfalsifiable, over legal restrictions $\sigma \supseteq \rho$.

Width lower bound for CPLS

There exist

1. a w-CNF B such that $\Pr[B \restriction \rho = 0] \leq 1/2$
2. a width w resolution refutation Π of $B \wedge \text{CPLS}$.

Fixing Lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

That is, with polynomially high probability, $B \restriction \rho$ is forced to be either false or unfalsifiable, over legal restrictions $\sigma \supseteq \rho$.

From the lemma and 1., $\Pr[\text{no legal } \sigma \supseteq \rho \text{ falsifies } B] > 1/4$.

Fix a "good" ρ with this property, where "good" means it has roughly the expected distribution of stars.

Width lower bound for CPLS

We have

1. a width w resolution refutation Π of $B \wedge \text{CPLS}$
2. a good restriction ρ such that no legal $\sigma \supseteq \rho$ falsifies B.

Width lower bound for CPLS

We have

1. a width w resolution refutation Π of $B \wedge \text{CPLS}$
2. a good restriction ρ such that no legal $\sigma \supseteq \rho$ falsifies B.

Inductively find a path through Π from the empty clause up to an initial clause, such that for every clause C on the path, there is a legal $\sigma \supseteq \rho$ such that

- σ falsifies C
- σ extends ρ by at most w bits.

Width lower bound for CPLS

We have

1. a width w resolution refutation Π of $B \wedge \text{CPLS}$
2. a good restriction ρ such that no legal $\sigma \supseteq \rho$ falsifies B.

Inductively find a path through Π from the empty clause up to an initial clause, such that for every clause C on the path, there is a legal $\sigma \supseteq \rho$ such that

- σ falsifies C
- σ extends ρ by at most w bits.

This is a Prover-Adversary argument. With only w bits, the Prover cannot remember long paths; or the colours present on many paths; or that all colours at a node are 0.

Width lower bound for CPLS

We have

1. a width w resolution refutation Π of $B \wedge \text{CPLS}$
2. a good restriction ρ such that no legal $\sigma \supseteq \rho$ falsifies B.

Inductively find a path through Π from the empty clause up to an initial clause, such that for every clause C on the path, there is a legal $\sigma \supseteq \rho$ such that

- σ falsifies C
- σ extends ρ by at most w bits.

This is a Prover-Adversary argument. With only w bits, the Prover cannot remember long paths; or the colours present on many paths; or that all colours at a node are 0.

When C is an initial clause this gives a contradiction, because no legal σ falsifies CPLS, and no legal $\sigma \supseteq \rho$ falsifies B.

□

Definition

Upper bounds

Lower bounds

Fixing lemma

Proof of the fixing lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

Proof of the fixing lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

Proof

Fix B. Let S be the set of ρ for which the lemma fails, that is,

ρ does not falsify B but some legal $\sigma \supseteq \rho$ falsifies B.

Proof of the fixing lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

Proof

Fix B. Let S be the set of ρ for which the lemma fails, that is,

ρ does not falsify B but some legal $\sigma \supseteq \rho$ falsifies B.

Let C be the first clause of B falsified by some legal $\sigma \supseteq \rho$.
Let z be the first variable of C not fixed by ρ.

Proof of the fixing lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

Proof

Fix B. Let S be the set of ρ for which the lemma fails, that is,

ρ does not falsify B but some legal $\sigma \supseteq \rho$ falsifies B.

Let C be the first clause of B falsified by some legal $\sigma \supseteq \rho$.
Let z be the first variable of C not fixed by ρ.

Let σ' be the minimal $\rho \subset \sigma' \subseteq \sigma$ which fixes z.

Proof of the fixing lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

Proof

Fix B. Let S be the set of ρ for which the lemma fails, that is,

ρ does not falsify B but some legal $\sigma \supseteq \rho$ falsifies B.

Let C be the first clause of B falsified by some legal $\sigma \supseteq \rho$.
Let z be the first variable of C not fixed by ρ.

Let σ' be the minimal $\rho \subset \sigma' \subseteq \sigma$ which fixes z.

Define $\theta : S \to \mathcal{R}$ by $\theta : \rho \mapsto (\sigma', \text{position of } z \text{ in } C, \beta < 3)$
where β is extra information guaranteeing that θ is an injection.

Proof of the fixing lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

Proof

Fix B. Let S be the set of ρ for which the lemma fails, that is,

ρ does not falsify B but some legal $\sigma \supseteq \rho$ falsifies B.

Let C be the first clause of B falsified by some legal $\sigma \supseteq \rho$.
Let z be the first variable of C not fixed by ρ.

Let σ' be the minimal $\rho \subset \sigma' \subseteq \sigma$ which fixes z.

Define $\theta : S \to \mathcal{R}$ by $\theta : \rho \mapsto (\sigma', \text{position of } z \text{ in } C, \beta < 3)$
where β is extra information guaranteeing that θ is an injection.

By analyzing $\mathcal{R}$, we can show $\Pr[\sigma'] \geq \frac{1}{2} n^{1/7} \Pr[\rho]$.

Proof of the fixing lemma

For any w-CNF B, for $\rho \sim \mathcal{R}$,

$$\Pr[(B \restriction \rho = 0) \text{ or } (B \restriction \sigma \neq 0 \text{ for every legal } \sigma \supseteq \rho)] > 1 - n^{-1/60}.$$

Proof

Fix B. Let S be the set of ρ for which the lemma fails, that is,

ρ does not falsify B but some legal $\sigma \supseteq \rho$ falsifies B.

Let C be the first clause of B falsified by some legal $\sigma \supseteq \rho$.
Let z be the first variable of C not fixed by ρ.

Let σ' be the minimal $\rho \subset \sigma' \subseteq \sigma$ which fixes z.

Define $\theta : S \to \mathcal{R}$ by $\theta : \rho \mapsto (\sigma', \text{position of } z \text{ in } C, \beta < 3)$
where β is extra information guaranteeing that θ is an injection.

By analyzing $\mathcal{R}$, we can show $\Pr[\sigma'] \geq \frac{1}{2} n^{1/7} \Pr[\rho]$.
It follows that $\Pr[S] \leq 6wn^{1/7}$. □