Overview

Proof Complexity

- Resolution
- Resolution lower bounds
- Separation of tree- and dag-like Resolution
- Width-restricted clause learning

Proof systems

A proof system for set A is a binary relation R with

- $R(x, y)$ can be decided in time $poly(|x|, |y|)$.
- there is p s.t. $R(p, a)$ iff $a \in A$

p with $R(p, a)$ is called a proof of a.

Question: What is the size of a minimal proof of $a \in A$, as a function of $|a|$

Motivation: Complexity Theory, Logic, Algorithms

Proof systems for UNSAT

Proof system R is polynomially bounded, if for some d,
for every $a \in A$ there is p with $R(p, a)$ and $|p| \leq O(|a|^d)$

A has a polynomially bounded proof system iff A is in NP.

Let UNSAT be the set of unsatisfiable formulas.

Fact: UNSAT has a polynomially bounded proof system
iff NP = co-NP.

Simulation

Let R and R' be proof systems for A

R' simulates R $\quad (R \leq R')$,

if for every $a \in A$ and p with $R(p, a)$,
there is p' with $R'(p', a)$ with $|p'| \leq O(|p|^d)$.

R and R' are equivalent $\quad (R \equiv R')$,

if $R \leq R'$ and $R' \leq R$.

Proof systems and SAT algorithms

Let A be a complete SAT algorithm.

Proof system R_A with:

proof that F is UNSAT $\triangleq$ transcript of run of A on F

R_A is equivalent to some natural proof system, for many algorithms A.

Resolution

The resolution rule:

from $C \vee a$ and $D \vee \bar{a}$ derive $C \vee D$.

A Resolution derivation of clause C from formula F
is a dag labelled with clauses s.t.

- every node has in-degree 0 or 2
- there is exactly one sink labelled C
- If v has 2 predecessors u and u', then C_v is derived by resolution from C_u and $C_{u'}$.
- if v is a source, then $C_v \in F$

A Resolution refutation of F is a derivation
of the empty clause $\square$ from F.

Tree-like and Regular Resolution

A resolution refutation is tree-like, if the underlying dag is a tree.

Resolution refutation is regular, if no variable is eliminated twice on a path.

Theorem

F has a tree-like, regular Resolution refutation iff F is unsatisfiable.

Theorem

If F has a tree-like Resolution refutation of size s,

then F has a tree-like regular Resolution refutation of size at most s.

SAT algorithms and Resolution

Theorem

If DPLL runs in time t on unsatisfiable formula F, then F has a tree-like regular Resolution refutation of size at most t.

A version of the converse also holds.

Theorem

If CDCL runs in time t on unsatisfiable formula F, then F has a (dag-like) Resolution refutation of size at most t.

Restrictions

A restriction is a partial assignment.

Theorem

Let P be a Resolution refutation of a formula F, and ρ a restriction.

Then there is a Resolution refutation P' of $F\rho$ of size at most $|P'| \leq |P|$.

We denote the refutation P' from the theorem by $P\restriction_\rho$.

- P tree-like $\rightsquigarrow$ $P\restriction_\rho$ tree-like
- P regular $\rightsquigarrow$ $P\restriction_\rho$ regular

The Pigeonhole Principle

The Pigeonhole Principle formula PHP_n^m:

$$\begin{array}{rcll} P_i & := & x_{i,1} \vee \ldots \vee x_{i,n} & i \leq m \\ H_{i,j;k} & := & \bar{x}_{i,k} \vee \bar{x}_{j,k} & i < j \leq m, \quad k \leq n \end{array}$$

We denote the set of pigeon axioms P_i in PHP_n^m by PA_n^m.

Fact: PHP_n^m is unsatisfiable iff $m > n$.

Matching restrictions

A matching ρ from $[m]$ into $[n]$ is a set of pairs

$$\{(i_1, j_1), \ldots, (i_k, j_k)\} \subset [m] \times [n]$$

such that all i_ν and all j_ν are pairwise distinct.

A matching ρ induces a restriction as follows:

$$\rho(x_{i,j}) = \begin{cases} 1 & \text{if } (i,j) \in \rho \\ 0 & \text{if there is } (i,j') \in \rho \text{ with } j \neq j' \\ & \text{or } (i',j) \in \rho \text{ with } i \neq i' \\ \text{undefined} & \text{otherwise.} \end{cases}$$

Property: $PHP_n^m\restriction_\rho \equiv PHP_{n-|\rho|}^{m-|\rho|}$

Critical assignments

A matching ρ also defines a total assignment α_ρ:

$$\alpha_\rho(x_{i,j}) = \begin{cases} 1 & (i,j) \in \rho \\ 0 & \text{otherwise.} \end{cases}$$

α_ρ satisfies all hole clauses $H_{i,j,k}$
and exactly $|\rho|$ of the pigeon clauses P_i.

A critical assignment is a total assignment $\alpha = \alpha_\rho$,
where ρ is a maximal matching of size $|\rho| = n$.

The Monotone Calculus

The monotone calculus is a proof system to refute pigeonhole axioms.

Lines in a proof are positive clauses.

Let $P_{I,J} = \bigvee_{i \in I} \bigvee_{j \in J} x_{i,j}$, and $P_{I,j} = P_{I,\{j\}}$

The only inference rule:

$$\frac{C \vee P_{I_0,j} \qquad D \vee P_{I_1,j}}{C \vee D}$$

where I_0 and I_1 are disjoint subsets of $[m]$.

The monotone calculus is correct w.r.t. critical assignments.

The Monotone Calculus and Resolution

Proposition

If PA_n^m has a monotone calculus refutation of size s, then there is a Resolution refutation of PHP_n^m of size at most $m^2 \cdot s$.

Theorem

If PHP_n^m has a Resolution refutation of size s, then there is a monotone calculus refutation of PA_n^m of size at most s.

Lower bound for *PHP*

Theorem

If P is a monotone calculus refutation of PA^n_{n-1},
then P is of size $|P| \geq 2^{n/20}$.

Strategy of proof:

1. Convert a short refutation P to a refutation P' of $PA^{n'}_{n'-1}$ for $n' < n$, such that P' contains only narrow clauses.
2. Show that any refutation of PA^n_{n-1} contains a wide clause.

Goal 1: Removing wide clauses

Lemma

If P is a monotone refutation of PA^n_{n-1} of size $|P| < 2^{n/20}$,
then there is a matching restriction ρ with

- $|\rho| \leq 0.329n$
- $P\restriction_\rho$ *contains no large clause C of width $w(C) \geq n^2/10$.*

Greedy algorithm to find ρ:

$\rho := \emptyset$
while there is a large clause in $P\restriction_\rho$
    pick (i,j) s.t. $x_{i,j}$ occurs in most large clauses
    $\rho := \rho \cup \{(i,j)\}$

Goal 2: Width lower bound

Lemma

If P is a monotone calculus refutation of PA^n_{n-1},
then there is a clause C in P with $w(C) \geq 2n^2/9$.

Proof strategy:

- Define a measure $\mu(C) \leq n$ on clauses in P.
- Show that there is C with $n/3 \leq \mu(C) \leq 2n/3$.
- Show that $w(C) \geq \mu(C)(n - \mu(C))$.

The measure μ

$F \models_{cr} C$ if $\alpha \models C$ for every critical $\alpha \models F$.

$\mu(C) := \min\{\, |F| ;\ F \subseteq PA^n_{n-1} \text{ and } F \models_{cr} C \,\}$

$\mu(\Box) = n$ and $\mu(P_i) = 1$.

If D follows from C and C' by resolution, then $\mu(D) \leq \mu(C) + \mu(C')$.

$\rightsquigarrow$ there is C in P with $n/3 \leq \mu(C) \leq 2n/3$.

Lemma

$w(C) \geq \mu(C)(n - \mu(C))$.

Upper bound

Theorem

The clauses PA_n^{n+1} have a monotone calculus refutation of size $O(n2^n)$.

Proof: By induction on k, derive all clauses

$$P_{I,\{k,\ldots,n\}}$$

for every set $I \subseteq [n+1]$ of size $|I| = k$.

Corollary

The clauses PHP_n^{n+1} have Resolution refutations of size $O(n^3 2^n)$.

Separation of tree- and dag-like Resolution

Theorem

If P is a tree-like Resolution refutation of PHP_n^{n+1}, then the size of P is at least $2^{\Omega(n \log n)}$.

Proof strategy: Construct a tree T_P with

- every vertex in T_P is a vertex in P,
  children of v in T are descendants of v in P
- the depth of T_P is $n/2$
- every vertex in T_P has either 1 or $n/4$ children
- on every path in T_P at least $n/4$ vertices have $n/4$ children
- therefore $|P| \geq |T_P| \geq (n/4)^{(n/4)}$

Some definitions

For each node v labelled C_v define restriction ρ_v with $C_v\rho_v = 0$:

- for the root r set $\rho_r = \emptyset$
- $C_v = D \vee D'$ inferred from $C_{v0} = D \vee x_v$ and $C_{v1} = D' \vee \bar{x}_v$ set $\rho_{v0} = \rho_v \cup [x_v := 0]$ and $\rho_{v1} = \rho_v \cup [x_v := 1]$

Let ρ be a restriction. Variable $x_{i,j}$ is

- consistent with ρ, if there is no i', j' such that $\rho(x_{i',j}) = 1$ or $\rho(x_{i,j'}) = 1$,
- active for ρ, if consistent with ρ and $\rho(x_{i,j})$ is undefined,
- bad for ρ, if consistent with ρ, and $\rho(x_{i,j}) = 0$.

$B(\rho)$ is the number of variables that are bad for ρ.

ρ is dangerous, if there is i s.t. $x_{i,j}$ is bad for ρ for $n/2$ many j.

Defining the tree T_P

The root of T_P is the root of P.

To define the children of a node v, we inductively define sets C_i and nodes v_i:

- $C_0 = \emptyset$ and $v_0 = v$
- if $|C_i| < n/4$ and ρ_{v_i} is not dangerous:
  - $C_{i+1} = C_i \cup \{v_i 1\}$ if x_{v_i} is active for ρ_{v_i},
  - $C_{i+1} = C_i$ otherwise,
  - $v_{i+1} = v_i 0$.
- if $|C_i| = n/4$ and ρ_{v_i} is not dangerous:
  children of v are all vertices in C_i.
- if ρ_{v_i} is dangerous, then $v_{i-1}1$ is the only child of v

Properties of the tree T_P

Lemma

The definition of T_P can be continued until depth $n/2$.

Lemma

Let v' be a child of v in T_P.

- *If v' is the only child of v, then $B(\rho_{v'}) \leq B(\rho_v) - n/4$.*
- *If v' is not the only child of v, then $B(\rho_{v'}) \leq B(\rho_v) + n/4 - 1$.*

Lemma

On every path in T_P, at most $n/4$ vertices have only one child in T_P.

The Ordering Principle

. . . says: An ordering of $[n]$ has a maximum

The formula Ord_n :

- variables $x_{i,j}$ for $i, j \leq n$ and $i \neq j$
- totality clauses $x_{i,j} \vee x_{j,i}$ for all i, j
- asymmetry clauses $\bar{x}_{i,j} \vee \bar{x}_{j,i}$ for all i, j
- transitivity clauses $\bar{x}_{i,j} \vee \bar{x}_{j,k} \vee \bar{x}_{k,i}$ for all i, j, k
- maximum clauses $M_i^{(n)}$ $\bigvee_{j \leq n, j \neq i} x_{i,j}$ for all i

Upper bound for the Ordering Principle

Theorem

There are regular resolution proofs of Ord_n of size $O(n^3)$.

Proof: By induction on k from n downward, derive Ord_k.

It suffices to derive the clauses $M_i^{(k)} = \bigvee_{j \leq k, j \neq i} x_{i,j}$

Ordering restrictions

Ordering restriction: $\sigma = \sigma_{S,\prec}$ defined by $S \subseteq [n]$ and an ordering $\prec$ on S.

$$\sigma(x_{i,j}) = \begin{cases} 1 & \text{if } i,j \in S \text{ and } i \prec j \\ 0 & \text{if } i,j \in S \text{ and } j \prec i \\ x_{i,j} & \text{otherwise,} \end{cases}$$

For $\sigma = \sigma_{S,\prec}$, we denote S by $S(\sigma)$ and $\prec$ by $\prec_\sigma$.

Lower bound for the Ordering Principle

Theorem

Every tree-like resolution refutation P of Ord_n has size $|P| \geq 2^{\Omega(n)}$.

Proof strategy: Construct a sub-tree T_P of P with

- for every vertex v in T_P define ordering restriction σ_v with $C_v\sigma_v = 0$
- every vertex in T_P has either 1 or 2 children
- every path in T_P has at least $n/4$ vertices with 2 children
- therefore $|P| \geq |T_P| \geq 2^{(n/4)}$

Construction of T_P

The root of T_P is the root of P.

Let v be a vertex in T_P, with $|S(\sigma_v)| < n/2$ and variable $x_v = x_{i,j}$.

- if $\{i,j\} \subseteq S(\sigma_v)$, so $\sigma_v(x_{i,j})$ is defined
  - let v' be the child of v in P with $C_{v'}\sigma_v = 0$
  - add v' as the only child of v in T_P, and let $\sigma_{v'} = \sigma_v$
- otherwise $\sigma_v(x_{i,j})$ is undefined
  - add both children of v in P to T_P
  - let $S' = S \cup \{i,j\}$
  - extend $\prec$ to $\prec_0$ with $j \prec_0 i$, set $\sigma_{v0} = \sigma_{S',\prec_0}$
  - extend $\prec$ to $\prec_1$ with $i \prec_1 j$, set $\sigma_{v1} = \sigma_{S',\prec_1}$

Continue extending the tree until $S(\sigma_v) \geq n/2$ in every leaf v.

Resolution Trees with Lemmas

A Resolution tree with lemmas (RTL) for formula F is an ordered binary tree labelled with clauses s.t.

- $C_{\text{root}} = \square$
- if v has 2 children u and u', then

  C_v is obtained by resolution from C_u and $C_{u'}$
- if v has 1 child u, then

  $C_v \supseteq C_u$
- if v is a leaf, then

  $C_v \in F \quad$ or $\quad C_v = C_u$ for some $u \prec v$ (lemma)

$\prec$ is the post-order on trees.

Clause learning and *RTL*

Theorem

If unsatisfiable formula F is refuted by CDCL without restarts in s steps, then F has an RTL-refutation R of size $s \cdot n^{O(1)}$.

Moreover, the lemmas used in R are among the clauses learned by the algorithm.

In fact, there is a subsystem $WRTI < RTL$
for which a sort of converse also holds.

Fact: regular Resolution $\leq_{\neq}$ regular RTL $\leq$ Resolution

Lower bounds for $RTL(k)$

A refutation R in RTL is in $RTL(k)$, if every lemma C used in R is of width $w(C) \leq k$.

Theorem

For $k \leq n/2$, every $RTL(k)$-refutation of PHP_n^{n+1} is of size $2^{\Omega(n \log n)}$.

Shows: Learning short clauses does not help to refute PHP.

Lower bound for the Pigeonhole Principle

Lemma

For R an $RTL(k)$-refutation of F, there is R' that contains no lemma $D \supset C$ for $C \in F$, and $|R'| \leq 2|R|$.

Lower bound is shown for $FPHP_n^{n+1}$ with functional clauses:

- $F_{i;j,k}$ $\quad \bar{x}_{i,j} \vee \bar{x}_{i,k}$ $\quad$ for $j < k$

Main Lemma

Let C be a clause of width $w(C) \leq k \leq n/2$, such that

- C not subsumed by hole clause $\bar{x}_{i,j} \vee \bar{x}_{i',j}$
- C not subsumed by functional clause $\bar{x}_{i,j} \vee \bar{x}_{i,j'}$

Then there is a matching restriction ρ with $C\lceil_\rho = 0$ and $|\rho| \leq k$.

Lower bound for the Pigeonhole Principle

Proof of the lower bound:

- Let R be a $RTL(k)$-refutation of $FPHP_n^{n+1}$.
- W.l.o.g. no lemma is subsumed by hole or functional clause.
- Let C be the first clause in R used as a lemma, so $w(C) \leq k$.
- Subtree R_C below C is tree-like resolution derivation of C.
- By the Main Lemma, there is matching restriction ρ with $C\restriction_\rho = 0$ and $|\rho| \leq k$.
- Thus $R_C\restriction_\rho$ is a tree-like refutation of $FPHP_n^{n+1}\restriction_\rho = FPHP_{n-k}^{n-k+1}$.
- Therefore $|R| \geq |R_C| \geq |R_C\restriction_\rho| \geq (n/8)^{n/8} \geq 2^{\Omega(n \log n)}$.

Lower bound for the Ordering Principle

Theorem

For $k < n/4$, every RTL(k)-refutation of Ord_n is of size $2^{\Omega(n)}$.

Proof strategy: imitate proof for *PHP*.

Problem 1: proof shows it takes long to derive C sufficiently short
not true here!

Problem 2: need notion of restriction preserving Ord_n formulas
ordering restrictions don't work!

New and improved ordering restrictions

Ordering restriction: defined by $S \subseteq [n]$
and an ordering $\prec$ on S.

$$\sigma(x_{i,j}) = \begin{cases} 1 & \text{if } i,j \in S \text{ and } i \prec j \\ 0 & \text{if } i,j \in S \text{ and } j \prec i \\ x_{s,j} & \text{if } i \in S \text{ and } j \notin S \\ x_{i,s} & \text{if } i \notin S \text{ and } j \in S \\ x_{i,j} & \text{otherwise,} \end{cases}$$

where $s \in S$ is fixed.

Property: $Ord_n\lceil_\sigma \equiv Ord_{n-|S|+1}$.

Cyclic clauses

For clause C, the graph $G(C)$ has edges

(i,j) for $\bar{x}_{i,j} \in C$ and

(j,i) for $x_{i,j} \in C$

Definition: C is *cyclic*, if $G(C)$ contains a cycle.

Lemma: A cyclic clause C has a tree-like resolution derivation from Ord_n of size $O(w(C))$.

The main lemmas

Lemma

If there is an $RTL(k)$-refutation of Ord_n of size s, then there is another one using no cyclic lemmas of size $O(sk)$.

Proof: Replace each cyclic lemma by its derivation of size $O(k)$.

Lemma

If C is acyclic with $w(C) \leq k$, then there is an ordering restriction σ with $|\sigma| \leq 2k$ such that $C\lceil_\sigma = 0$.

Proof: For C acyclic $G(C)$ is a dag

$\rightsquigarrow$ obtain σ as a topological ordering of $G(C)$.

Lower bound for the Ordering Principle

Proof of the lower bound:

- Let R be a $RTL(k)$-refutation of Ord_n.
- W.l.o.g. no lemma is cyclic.
- Let C be the first clause in R used as a lemma, so C is acyclic and $w(C) \leq k$.
- Subtree R_C below C is tree-like resolution derivation of C.
- By the Main Lemma, there is an ordering restriction σ with $C\lceil_\sigma = 0$ and $|\sigma| \leq 2k$.
- Thus $R_C\lceil_\sigma$ is a tree-like refutation of $Ord_n\lceil_\sigma = Ord_{n-2k+1}$.
- Therefore $|R| \geq |R_C| \geq |R_C\lceil_\rho| \geq 2^{n/8} \geq 2^{\Omega(n)}$.