Started this course looking at RESOLUTION
Lines in proofs are clauses
Resolution rule

$$\frac{C \vee x \quad D \vee \bar{x}}{C \vee D}$$

Then detour to CUTTING PLANES
Geometric proof system
Lines in proofs are linear inequalities
= hyperplanes in $\mathbb{R}^n$ (n = # variables)

Now go back to proof systems where lines are Boolean formulas

Can allow more general formulas as proof lines
Can generalize resolution rule to
CUT RULE

$$\frac{F \vee H \quad G \vee \neg H}{F \vee G}$$

(Plus other rules for syntactic manipulation of formulas)

Most general form with essentially no restrictions on formulas
FREGE PROOF SYSTEMS
Very strong. Do not know how to prove strong lower bounds (or almost any lower bounds) for these systems
Best lower bounds are quadratic.

BOUNDED-DEPTH FREGE

Formulas must have bounded number of alternations between $\wedge$ and $\vee$

Only clauses $\bigvee_i a_i$ Depth-1 Frege = resolution

DNF formulas $\bigvee_i \bigwedge_j a_{ij}$ Depth-2 Frege

Disjunctions of CNF formulas $\bigvee_i \bigwedge_j \bigvee_\ell a_{ij\ell}$ Depth-3 Frege

Et cetera

Bounded-depth Frege (vanilla definition)

Start with CNF formula, as before
All lines in proof are formulas of depth $\leq d$ for some fixed $d \in \mathbb{N}^+$
using cut rule or syntactic massaging rules
Prove lower bound by showing for some CNF formula family $\{F_n\}_{n=1}^{\infty}$ that minimum-size refutations scale superpolynomially

Cutting planes and bounded-depth Frege both right at the research frontier
Major progress has been made in the last few years. But still many open problems.

Already talked a bit about cutting planes
Now move on to bounded-depth Frege

But a bit cautiously...
Next few lectures focus on proof systems strictly between resolution and depth-2 Frege

"Resolution with bounded conjunctions" or "d-DNF resolution" or "Res(d)" or "$\mathcal{R}(d)$"

Lines in proofs are d-DNF formulas (disjunctions not of literals, but of conjunctions or <u>terms</u> of size d)

Resolution rule generalized to

$$\frac{F \vee (a_1 \wedge a_2 \wedge \cdots \wedge a_d) \qquad G \vee \bar{a}_1 \vee \bar{a}_2 \vee \cdots \vee \bar{a}_d}{F \vee G}$$

Proof system defined by Krajíček in 2001

We will follow papers [Segerlind, Buss, Impagliazzo 2004]
[Alekhnovich 2011]
(journal versions of slightly earlier conference papers) to show

(1) PHP formulas hard for d-DNF resolution
(2) Random k-CNF formulas — " —

We will not show, but it is true, that:

(3) (d+1)-DNF resolution is exponentially stronger than d-DNF resolution

Side note 1

Proof system often called k-DNF resolution. But since we want to talk about refutations of k-CNF formulas, and that k is often different, we will try (at least in this lecture) to say "d-DNF resolution". We will sometimes use the shorthand "$\mathcal{R}(d)$" for brevity. Also, will try to reserve "F" for CNF formulas from now on, using G and H for d-DNF formulas.
Let us now give some formal definitions

PRELIMINARIES

Term — conjunction $a_1 \wedge a_2 \wedge \ldots \wedge a_d$
a_i literals, size of term $= d$

d-term — Term of size $\leq d$

DNF formula — Disjunction of terms

d-DNF formula — Disjunction of terms of size $\leq d$

(A clause is a 1-DNF formula.)

Two terms t, t' are CONSISTENT if the term $t \wedge t'$ is satisfiable, i.e., if there is no literal a with $a \in t$, $\bar{a} \in t'$ (We assume that a term is consistent with itself)

Will also consider terms as sets of literals when convenient, so, e.g., $|t| =$ size of term

d-DNF resolution inference rules

G, H d-DNF formulas
t, t' d-terms
a_i literals (as usual)

d-CUT $$\frac{(a_1 \wedge \cdots \wedge a_{d'}) \vee G \qquad \bar{a}_1 \vee \cdots \vee \bar{a}_{d'} \vee H}{G \vee H} \quad d' \leq d$$

$\wedge$-INTRODUCTION $$\frac{G \vee t \qquad G \vee t'}{G \vee (t \wedge t')} \quad |t \cup t'| \leq d$$

$\wedge$-ELIMINATION $$\frac{G \vee t}{G \vee t'} \quad t' \subseteq t$$

WEAKENING $$\frac{G}{G \vee H} \quad H \text{ any } d\text{-DNF formula}$$

There are slightly different definitions in the ~~literature~~ literature, but the variations don't really matter. We will go with this definition, at least for now

d-DNF refutation of CNF formula F

$\pi = (G_1, G_2, \ldots, G_L)$

Each G_i either (a) axiom clause $\in F$ or (b) derived from previous lines using the inference rules above

G_L is empty d-DNF formula (denoted $\perp$)

Side note 2

Empty clause = falsity; no literal can satisfy it
Empty d-DNF formula = falsity; no term can satisfy it
Empty term = truth; no literal can falsify it

Somewhat analogous to

$$\sum_{i=1}^{0} a_i = 0 \qquad \prod_{i=1}^{0} a_i = 1$$

LENGTH of refutation $\pi = (G_1, \ldots, G_L)$ is # lines L

SIZE = total # literals counted with repetitions

Measures polynomially related
Every line can contain at most

$$\sum_{i=1}^{d} 2^i \binom{n}{i} = O(n^d) \text{ terms}$$

of size d (assuming $d = O(1)$)
so we usually don't care too much to distinguish between the two.

We also don't care too much about exactly how d-DNF resolution is defined, because we will only use that the derivation rules are STRONGLY SOUND as defined next

DEFINITION 1 (STRONG SOUNDNESS)
Let G_i, $i \in [s]$, and H be d-DNF formulas. We say that $G_1, \ldots, G_s$ STRONGLY IMPLIES H if for any set of mutually consistent terms $t_i \in G_i$, $i \in [s]$ it holds that there is a term $t \in H$ such that

$$\bigwedge_{i=1}^{s} t_i \vDash t$$

We say that a proof system (operating with d-DNFs) is STRONGLY SOUND if whenever H is derived from $G_1, \ldots, G_s$ it holds that $G_1, \ldots, G_s$ strongly implies H

Recall that G implies H, denoted $G \vDash H$, if for any (total) truth value assignment α s.t. $\alpha(G) = 1$ it holds that $\alpha(H) = 1$

OBSERVATION 2
If $G_1, \ldots, G_s$ strongly imply H, then $G_1, \ldots, G_s$ imply H in the standard sense of $G_1 \wedge \cdots \wedge G_s \vDash H$. The opposite direction does not hold.

<u>Proof</u> Exercise (likely to end up on problem set 2).

So strong implication is indeed stronger than just implication.

OBSERVATION 3
d-DNF resolution is strongly sound.

Proof Inspect the rules and do a (simple) case analysis.

[Detour: Hardness of PHP and random k-CNFs]

Important part of proof complexity
Study different families of formulas
Understand their hardness in different proof systems.

Pigeonhole principle PHP^m_n extremely well-studied (but still many interesting open problems).
What happens when we vary not only $n \to \infty$ but also $m = m(n)$?

$m \leq n$ — satisfiable; not so interesting

$m = n+1$ — exponentially hard for resolution [Haken '85]
Also for bounded-depth Frege [Krajíček, Pudlák, Woods '95] [Pitassi, Beame, Impagliazzo '93]
But easy for Frege [Buss '87]

$m = cn$, $c > 1$ constant — Haken's proof still works for resolution

$c = 2$ Proof crashes for stronger proof systems

[Segerlind, Buss, Impagliazzo '04]
$m = 2n$ pigeons hard for $\mathcal{R}(d)$

But for $R(\text{polylog}(n))$

DNFs with slowly growing terms

$\text{polylog}(n) = \log^k(n)$ for some $k \in \mathbb{N}^+$

PHP_n^{2n} is not too hard

Exist refutations in quasipolynomial size [Maciel $\boxed{\text{Woods}}$ $\boxed{\text{Pitassi}}$ '02]

$\exp(\text{polylog}(n)) = n^{\log^k(n)}$ for some $k \in \mathbb{N}^+$

Superpolynomial, but not by too much.

$\underline{m = n^2}$ "Weak pigeonhole principle"

Haken's proof breaks down

But actually even for arbitrarily many pigeons m it holds that resolution refutations require length

$\exp(\Omega(n^\delta))$ for some $\delta > 0$

[Raz '01]

In fact true for all flavours of resolution, even with functionality and/or onto axioms

[Razborov '03, '04] (These are deep papers.)

How hard are random k-CNF formulas?

Hard for resolution [Chvátal-Szemerédi '88]

d^2-CNFs hard for $R(d)$ [SBI '04]

3-CNFs hard for $R(d)$ [Alekhnovich '11]

Bounded-depth Frege? Big open problem.

A key tool for the lower bounds that we will prove in future lectures are <u>random restrictions</u>. Let us recall what a restriction is.

<u>RESTRICTIONS</u>

Partial assignment ρ to variables in a formula

$$\rho : Vars(F) \to \{0, 1, *\}$$

$\rho(x) = *$ means that ρ does not assign x

When representing ρ, can omit $*$-assignments

Can hit formula with restriction and simplify "in obvious way".

For terms: $\quad 0 \wedge t = 0 \qquad 1 \wedge t = t$

For clauses: $\quad 0 \vee C = C \qquad 1 \vee C = 1$

For formulas of higher depth, simplify layer by layer.

<u>Example</u> $\quad G = (x \wedge y \wedge z) \vee (\overline{x} \wedge \overline{z} \wedge w)$

$\rho_1 = \{x \mapsto 1\} \qquad G\restriction_{\rho_1} = y \wedge z$

$\rho_2 = \{x \mapsto 1, z \mapsto 0\} \qquad G\restriction_{\rho_2} = \bot$ (contradiction/falsity)

$\rho_3 = \{x \mapsto 1, y \mapsto 1, z \mapsto 1\} \qquad G\restriction_{\rho_3} = 1$ (truth/trivial)

Sometimes also think of ρ as $\{$ literals set to true $\}$, so

$\rho_2 = \{x, \overline{z}\}$, $\rho_3 = \{x, y, z\}$ Can be convenient for notation

General fact (true, with small twists, for all proof systems we will study).

For $\pi = (G_1, G_2, \ldots, G_L)$ define
$\pi\restriction_\rho = (G_1\restriction_\rho, G_2\restriction_\rho, \ldots, G_L\restriction_\rho)$

Then if π is a refutation of F, it holds that $\pi\restriction_\rho$ is a refutation of $F\restriction_\rho$

(Might also end up as an exercise for problem set 2 for cutting planes and/or d-DNF resolution).

Random restrictions ρ sampled from suitable distributions $\mathcal{D}$ can be a powerful tool to prove lower bounds. Let us see an example, for simplicity using resolution.

Take formulas PHP_n^{n+1} again
Prove, firstly, that for any resolution refutation $\pi : PHP_n^{n+1} \vdash \bot$ there has to be some clause which mentions a lot about $\cancel{n/3}$ $n/2$ pigeons $I \subseteq [n]$ in the following way
For every $i \in I$, $|I| = \cancel{n/3}$ $n/2$, one of two cases hold

① $\exists$ literal $\boxed{\text{negated}}$ $\bar{x}_{i,j}$ in C

② $\exists$ $n/2$ positive literals $x_{i,j_1}, \ldots, x_{i,j_{n/2}}$
for different $j_1, \ldots, j_{n/2}$

Secondly take a random matching M of $n/2$ pigeons to $n/2$ and define the restriction ρ_M by

$$\rho(x_{ij}) = \begin{cases} 1 & \text{if } (i,j) \in M, \\ 0 & \text{if } (i,j') \in M \text{ for } j' \neq j, \\ * & \text{otherwise.} \end{cases}$$

Prove that for any clause C that for $|I'| \geq n/4$ pigeons $i \in I'$ contains

① $\overline{x}_{ij}$, or

② $x_{i,j_1} \vee x_{i,j_2} \vee \cdots \vee x_{j,n/4}$

Call such a clause _fat_.

it holds that $\Pr[C\restriction_\rho \neq 1] \leq p$

(think of $p = 2^{-\delta n}$, $\delta > 0$).

Now set $L = 1/p$ and suppose $\pi : PHP_n^{n+1} \vdash \bot$ has length $\leq L$

We have that

$\pi\restriction_{\rho_M}$ for any ρ_M

is in fact a resolution refutation of $PHP_{n/2}^{n/2+1}$ (after renaming of variables)

What is the probability (over ρ_M) that $\pi\restriction_{\rho_M}$ contains some fat clause?

Two cases

① C not fat $\Rightarrow$ clearly $C\restriction_{\rho_M}$ is fat with probability 0 (for any restriction, actually)

② C is fat $\Rightarrow$ then we have proven $\Pr[C\restriction_{\rho_M} \neq 1] < \rho = 1/L$ (and satisfied clauses are just removed from the restricted refutation).

$$\Pr[\pi\restriction_{\rho_M} \text{ contains fat clause}] \leq$$

$$\leq \Pr[\exists C \in \pi \text{ s.t. } C \text{ fat and } C\restriction_{\rho_M} \neq 1] \leq$$

$$\leq \sum_{C\in\pi} \Pr[\overset{C \text{ fat and}}{C\restriction_{\rho_M}} \neq 1] <$$

$$< \sum_{C\in\pi} 1/L \leq L \cdot 1/L = 1$$

So $\Pr[\pi\restriction_{\rho_M} \text{ contains fat clause}] < 1$

This means $\exists \rho^* \in \mathcal{D}$ s.t. $\pi\restriction_{\rho^*}$ has no fat clauses (otherwise we would have probability $= 1$)

But $\pi\restriction_{\rho^*}$ is a refutation of $PHP^{n/2+1}_{n/2}$, and we started by arguing that any such ~~clause~~ refutation has fat clause. Contradiction! Hence no refutation in length $\leq L = 1/\rho = 2^{\delta n}$ can exist!

Sounds familiar?

This is actually the Prosecutor-Defendant lower bound we did in lecture 2, only described in a slightly different way (and with slightly different numbers).

~~~

How to prove $R(d)$ lower bounds for PHP?

Terms are only of size d. Random restriction likely to either

(a) satisfy a term (and that whole line, removing it)

(b) falsify the term (removing it from the formula)

(c) leave just one literal $\Rightarrow$ which will give use a clause.

So if we are lucky:

Take $R(d)$ refutation π of PHP_n^{n+1}

Hit with random restriction

Argue that we only have clauses left, so this is a resolution refutation.

Use resolution lower bound!

<u>Bad news</u>: Won't work.

But something in a somewhat similar spirit will.

Take short $R(d)$ refutation & hit with random restriction.

Turns refutation into sequence of small decision trees.

Can be transformed to resolution refutation with clauses of small width. But such refutations cannot exist.
~~~

DECISION TREE T for function $f: \{0,1\}^n \to \{0,1\}$

Binary tree.

Internal nodes labelled by variables x, y, z

Outgoing edges labelled 0 & 1

Leaves labelled 0 / 1

Every total assignment α defines path to leaf v_α

T <u>represents</u> f if $f(\alpha) = \text{label}(v_\alpha)$.

T represents d-DNF $\mathcal{H}$ if it represents the function computed by $\mathcal{H}$

Path from root to leaf v defines partial assignment ρ_v

For $b \in \{0,1\}$, let $br_b(T) = \{$ paths/partial assignments to leaves labelled $b\}$

$T\restriction_\rho$ is the decision tree obtained by

(a) deleting edges conflicting with ρ (and subtrees below these edges)

(b) contracting edges agreeing with ρ

T <u>strongly represents</u> d-DNF formula $\mathcal{H}$

if for every $\rho \in br_0(T)$ and for all $t \in \mathcal{H}$ it holds that $t\restriction_\rho = 0$; and for every $\rho \in br_1(T)$ ~~and every~~ there exists a $t \in \mathcal{H}$ such that $t\restriction_\rho = 1$

The <u>representation height</u> $h(\mathcal{H})$ of a DNF formula $\mathcal{H}$ is the minimum height of a decision tree strongly representing $\mathcal{H}$

Ex $H = (x \wedge z) \vee (y \wedge \bar{z})$

T_1

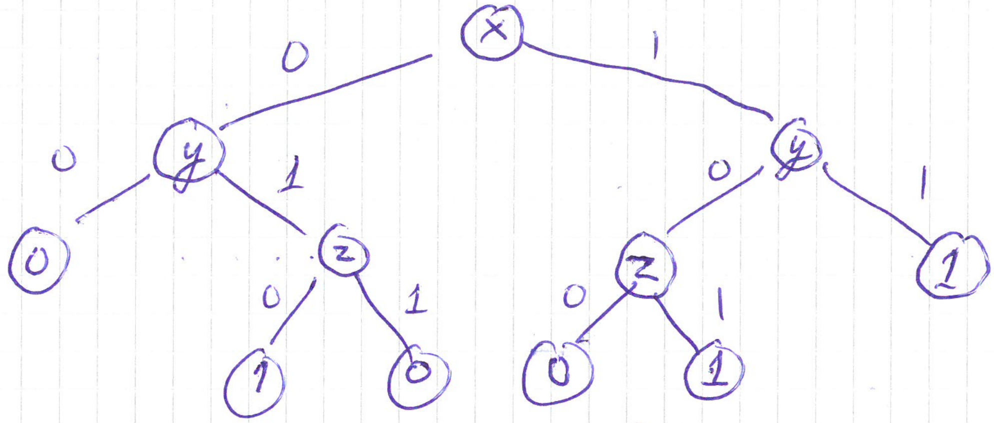

T_1 represents H, but does <u>not</u> represent H strongly, because

$$H\restriction_{\{x \mapsto 1,\, y \mapsto 1\}} = z \vee \bar{z}$$

T_2

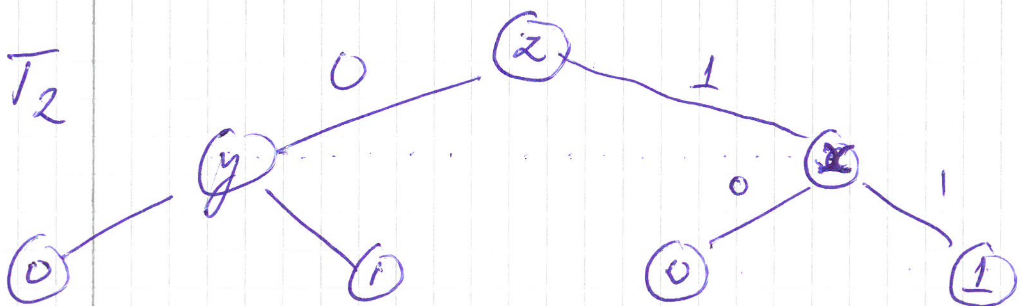

T_2 represents H strongly.

OBSERVATION

If $f : \{0,1\}^n \to \{0,1\}$ can be computed by a decision tree of height h, then f can be computed both by a h-CNF formula and a h-DNF formula.

Proof sketch The CNF formula F can be chosen as $\bigwedge_{\rho \in br_0(T)} \bigvee_{a \in \rho} \bar{a}$ and the DNF formula H as $\bigvee_{\rho \in br_1(T)} \bigwedge_{a \in \rho} a$

(if we think of $\rho = \{$ set of literals set to true by $\rho \}$) $\boxtimes$

How likely a d-DNF formula $\mathcal{H}$ is to be satisfied by a random restriction can be estimated by how much variables in different terms in $\mathcal{H}$ intersect. This is captured by the concept of COVERING NUMBER

DEFINITION Let H be a DNF formula and S be a set of variables. We say that S is a cover of $\mathcal{H}$ if every term $t \in \mathcal{H}$ contains a variable from S. The covering number of H, denoted $cov(\mathcal{H})$, is the minimum cardinality of a cover of H

Example For a clause $C = a_1 \vee \dots \vee a_w$, the covering number is $=$ width of C.

$H = (x \wedge y \wedge z) \vee (\bar{x}) \vee (y \wedge w)$ has $cov(H) = 2$

THEOREM (Switching lemma for d-DNF formulas)

Let $d, s, t \in \mathbb{N}^+$; let $\gamma, \delta \in (0,1]$, and let $\mathcal{D}$ be a distribution on partial assignments such that for every d-DNF formula G it holds that

$$\Pr_{\rho \sim \mathcal{D}}\left[G\restriction_\rho \neq 1 \right] \leq t \cdot 2^{-\delta(\mathrm{cov}(G))^\gamma}.$$

Then for every d-DNF formula H it holds that

$$\Pr_{\rho \sim \mathcal{D}}\left[h(H\restriction_\rho) > 2s \right] \leq dt \cdot 2^{-\delta' s^{\gamma'}}$$

where $\delta' = 2(\delta/4)^d$ and $\gamma' = \gamma^d$.

Compare For PHP and resolution, our random restriction satisfies ("kills") a clause of width w except with probability sth like $2^{-\sqrt{w}}$!
Since restriction is partial matching, never falsifies formula.

For d-DNF resolution, want to find random restrictions that (a) don't falsify the formula F being refuted (b) kill d-DNF formula H except with probability roughly

$$2^{-(\mathrm{cov}(H))^\gamma} \quad \text{for some } \gamma > 0.$$

THEOREM Let F be a h-CNF formula. If F has a d-DNF resolution refutation π such that for every line $H \in \pi$ it holds that $h(H) \leq h$, then there is a resolution refutation π' of F in width $W(\pi') \leq 2h$.

Given these two theorems,
~~to me~~ we can prove lower bounds for d-DNF resolution in the following way.

1. Find a good distribution $\mathcal{D}$ of random restrictions that satisfy the assumption of the switching lemma.

2. Argue that for short (d-DNF resolution) refutations π of F, we get shallow decision trees for $\pi\restriction_\rho$ for some ρ (by a union bound argument)

3. Prove that $F\restriction_\rho$ requires too large width in resolution for the decision-tree-to-resolution-refutation conversion to work.

Hence, there could be no short d-DNF resolution refutation of F.

Next lecture we will prove these two theorems. (Switching lemma and translation from decision trees to resolution refutations)