PROOF COMPLEXITY AS A COMPUTATIONAL LENS: LECTURE 12

Width was identified as a crucial resource of resolution proofs already by [Galil '77]

If F is a k-CNF formula ($k = O(1)$) over n variables, then:

(1) $W(F \vdash \bot) = w \Rightarrow L(F \vdash \bot) = n^{O(w)}$

(2) $L_{tr}(F \vdash \bot) = \exp(\Omega(W(F \vdash \bot)))$

(3) $L(F \vdash \bot) = \exp\left(\Omega\left(\frac{W(F \vdash \bot)^2}{n}\right)\right)$

Next week, we will talk about clause space and will see that

(4) $Sp(F \vdash \bot) = \Omega(W(F \vdash \bot))$

How tight are these results?

(a) Can we do something smarter than the counting argument in (1)?

(b) Could (3) be improved to something like (2)?

(c) Does small width also imply small clause space?

(d) In order to prove (3), we take a short proof and turn it into a narrow one, but the length increases exponentially — is this necessary?

TODAY we will answer questions (a) and (b). We will return to (c) & (d) in future lectures.

Note that we have completely analogous results ① and ③ for size vs degree in polynomial calculus, so same questions (a) - (d) make sense for PC as well.

Our results today apply for both resolution and PC.

Note

- No analogue of ② for tree-like PC (as far as I know)
- Analogue of ④ is believed to be true but is open — will talk about this in later lectures

TECHNICAL WARM-UP

Extended 3-CNF version $\widetilde{F}$ of F:

Convert wide clauses to 3-CNF with extension variables

For $C = \ell_1 \vee \ell_2 \vee \dots \vee \ell_W$

If $W \leq 3$ $\quad \widetilde{C} = C$

Otherwise

$$\widetilde{C} = (\ell_1 \vee \ell_2 \vee a_2) \wedge (\bar{a}_2 \vee \ell_3 \vee a_3) \wedge (\bar{a}_3 \vee \ell_4 \vee a_4) \wedge \vdots \wedge (\bar{a}_{W-2} \vee \ell_{W-1} \vee \ell_W)$$

for fresh variables a_i specific to C.

For $F = \bigwedge_{i=1}^{m} C_i$, define

$\tilde{F} = \bigwedge_{i=1}^{m} \tilde{C}_i$

OBSERVATION 1

For any unsatisfiable CNF formula F it holds that

$$L_R(\tilde{F} \vdash \bot) = O(L_R(F \vdash \bot))$$
$$W_R(\tilde{F} \vdash \bot) \leq W_R(F \vdash \bot)$$
$$S_{PCR}(\tilde{F} \vdash \bot) = O(S_{PCR}(F \vdash \bot))$$
$$Deg_{PCR}(\tilde{F} \vdash \bot) \leq Deg_{PCR}(F \vdash \bot)$$

For two CNF formulas F and G, say that G is a STRENGTHENING of F if for every $C \in F$ exists $D \in G$ such that $D \subseteq C$

OBSERVATION 2

If G is strengthening of unsatisfiable CNF formula F, then

$$L_R(G \vdash \bot) \leq L_R(F \vdash \bot)$$
$$W_R(G \vdash \bot) \leq W(F \vdash \bot)$$
$$S_{PCR}(G \vdash \bot) = O(S_{PCR}(F \vdash \bot))$$
$$Deg_{PCR}(G \vdash \bot) \leq Deg_{PCR}(F \vdash \bot)$$

PLAN FOR TODAY

(a) Show ③ is essentially best possible by exhibiting k-CNF formulas F_n over n variables s.t. $Deg_{PCR}(F_n \vdash \bot) = \Omega(\sqrt{n})$ but $L_R(F_n \vdash \bot) = O(n^{3/2})$

(b) Show ① tight by exhibiting k-CNF formulas F_n over n variables s.t. $W(F_n \vdash \bot) \leq w$ but $S_{PCR}(F_n \vdash \bot) = n^{-\Omega(w)}$

Proof Complexity as a Computational Lens 2025/26 Dec 4, 2025

Lecture 12

Lecturer: Jakob Nordström *Scribe: Lukáš Poláček, Add-scribe-name-here*

1 Quick Recap of Lower Bounds on Length in Terms of Width

Recall that in previous lectures we proved two bounds by Ben-Sasson and Wigderson [BW01] on length in terms of width. For an (unsatisfiable) CNF formula F over n variables, we have the bound

$$L_{\mathcal{R}}(F \vdash \bot) \geq \exp\left(\Omega\left(\frac{(W(F \vdash \bot) - W(F))^2}{n}\right)\right) \tag{1.1}$$

in general resolution, whereas for tree-like resolution we proved the cleaner (and stronger) bound

$$L_{\mathcal{T}}(F \vdash \bot) \geq 2^{W(F \vdash \bot) - W(F)}\,. \tag{1.2}$$

A natural question is: can we improve (1.1) to an expression similar to (1.2)? The short answer is: no, we cannot. The purpose of today's lecture is to show that the bound (1.1) is in some sense essentially optimal for general resolution. As a byproduct, we will also see an example which shows that general resolution is exponentially stronger than tree-like resolution with respect to length.

INSTRUCTOR'S COMMENT 1: *The discussion above needs to be updated with that a similar bound as in (1.1) holds for polynomial calculus size versus degree, as we have talked about in previous lectures, and that a similar question can be asked about the tightness of that relation.*

2 Formulas with Wide and Short Refutations

Suppose F is a k-CNF formula for some bounded $k = O(1)$. Then what (1.1) says is that if $W(F \vdash \bot) = \omega\left(\sqrt{n \log n}\right)$, it must hold that $L_{\mathcal{R}}(F \vdash \bot)$ is superpolynomial. Rephrasing our question above, what we are asking is whether a weaker lower bound on width can also provide guarantees for superpolynomial length lower bounds. We will prove that if one weakens the bound by only a $\sqrt{\log n}$ factor, there are no longer any guarantees. Namely, we will find k-CNF formulas that require width on the order of $\sqrt{n}$ but nevertheless have refutations of polynomial length.

The fact that the relation between size on the one hand and width/degree on the other hand is tight was proven for resolution by Bonet and Galesi [BG01] and for polynomial calculus by Galesi and Lauria [GL10].

INSTRUCTOR'S COMMENT 2: *We should maybe sharpen the statement of the result to say that the size of the formulas is $O(n^{3/2})$ and that the length/size of the refutations is also $O(n^{3/2})$.*

Theorem 2.1 ([BG01]). *There are 3-CNF formulas F_n over n variables with $\mathrm{poly}(n)$ clauses such that $L_{\mathcal{R}}(F_n \vdash \bot) = \mathrm{poly}(n)$ but $W(F_n \vdash \bot) = \Theta\left(\sqrt{n}\right)$.*

As discussed above, this implies that (1.1) is essentially optimal. By (1.2), a refutation in tree-like resolution for F_n has to have length at least $2^{\Omega(\sqrt{n})}$. Thus, we have the following corollary.

Corollary 2.2. *General resolution is exponentially stronger than tree-like resolution with respect to length.*

There are stronger separations of tree-like and general resolution—the best one is in [BIW04] as far as the lecturer is aware—but the nice thing with Corollary 2.2 is that we get it "for free" from Theorem 2.1.

A word of caution: *Above, the parameter n was the number of variables in the formula. This will change in a few seconds, and from then on n will instead be the natural parameter for a family of formulas $\{F_n\}_{n=1}^{\infty}$ with $\mathrm{poly}(n)$ variables. This might be a bit confusing, but these are standard conventions in*

the literature. Thus, rather than trying to shield course participants from the harsh realities of life by changing these conventions, we instead add this caveat to help deal with these realities head on.

Nevertheless, as an extra service to the reader we will also try to denote the number of variables by N below for increased clarity.

INSTRUCTOR'S COMMENT 3: *Not clear if N is a great choice, since I think we reserved this for the size of the input formula in the first lecture... Maybe think about whether this new usage of N here could be eliminated without losing much.*

In the rest of the lecture, we will prove Theorem 2.1. We will use formulas which encode *ordering principles*. Suppose we have a finite set $S_n = \{e_1, \ldots, e_n\}$ of partially (or totally) ordered elements, then S_n must have a minimal element. Note that this statement is not true for infinite sets, e.g. $\{1/n \mid n \in \mathbb{N}^+\}$.

Since we want unsatisfiable formulas, we will use CNF formulas saying that S_n is ordered but that despite of this there is no minimal element in the set. Below, we should interpret the variable x_{ij} to mean that $e_i < e_j$. We will use the following 4 types of clauses, where indices i, j, k range from 1 to n.

$$A(i,j,k) = \overline{x}_{ij} \vee \overline{x}_{jk} \vee x_{ik} \quad \text{for all } i \neq j \neq k \neq i \quad \text{(transitivity)} \tag{2.1a}$$

$$B(i,j) = \overline{x}_{ij} \vee \overline{x}_{ji} \quad \text{for all } i \neq j \quad \text{(anti-symmetry)} \tag{2.1b}$$

$$C_n(j) = \bigvee_{1 \leq i \leq n, i \neq j} x_{ij} \quad \text{for all } j \quad \text{(non-minimality)} \tag{2.1c}$$

$$D(i,j) = x_{ij} \vee x_{ji} \quad \text{for all } i \neq j \quad \text{(totality)} \tag{2.1d}$$

Clauses of type (2.1a), (2.1b) and (2.1c) form the *partial ordering principle formulas* and we denote the formula for S_n by POP_n. By adding clauses of type (2.1d), we get *linear ordering principle formulas* and we denote these formulas by LOP_n. We remark that these formulas go under a number of different names in the literature, but we will stick to POP_n and LOP_n in this lecture It is easy to check that both types of formulas have $\Theta(n^2)$ variables and $\Theta(n^3)$ clauses.

3 An Upper Bound on Refutation Length of Ordering Principles

The ordering principle formulas were conjectured to be hard to refute in resolution by Krishnamurthy [Kri85] (note that this is close in time to the first exponential lower bounds for resolution by Haken [Hak85]), but were instead proven to be easy by Stålmarck [Stå96] a decade later. We will follow an adaptation of Stålmarck's resolution refutation by Bonet and Galesi [BG01].

Note that $POP_n \subseteq LOP_n$, so a refutation for POP_n is also a refutation for LOP_n and hence it holds that $L(LOP_n \vdash \bot) \leq L(POP_n \vdash \bot)$. Our goal in this lecture is to prove an upper bound on resolution refutation length for POP_n (which will automatically hold for LOP_n as well).

Already a couple of lectures back, we proved a degree lower bound for graph ordering principle formulas. Upper bounds for POP_n imply upper bound for these formulas, and lower bounds for graph ordering principle formulas imply lower bounds for POP_n.

Theorem 3.1 ([Stå96]). *There exist resolution refutations of POP_n in length $O(n^3)$.*

Proof. For $n = 1$, we have $POP_1 = (\overline{x}_{12} \vee \overline{x}_{21}) \wedge x_{12} \wedge x_{21}$, and this formula can clearly be refuted in a (small) constant number of steps by a resolution refutation π_2. For bigger n our strategy is to derive POP_{n-1} from POP_n in polynomial length. If we have such resolution derivations $\pi_n : POP_n \vdash POP_{n-1}$, we can then string all these derivations $\pi_n, \pi_{n-1}, \ldots, \pi_3, \pi_2$ together to get a refutation of POP_n.

Note that clauses of type (2.1a) and (2.1b) from POP_{n-1} are all present in POP_n, so we only need to show a way to derive clauses of type (2.1c). Namely, we will derive clauses $C_{n-1}(1), \ldots, C_{n-1}(n-1)$ from $A(i,j,k), B(i,j)$ and $C_n(j)$ (for all needed i, j, k). The intuition behind this derivation is that we can extract a smaller set S_{n-1} from S_n which is also ordered and does not contain a minimal element by showing that since e_n is not minimal, some element in $\{e_1, \ldots, e_{n-1}\}$ must be.

More formally, we claim that any clause $C_{n-1}(i)$ can be derived in polynomial length, and state this formally as Lemma 3.2 below. Assuming this lemma, which we will prove shortly, we can apply it for all

$i = 1, 2, \ldots, n-1$ to obtain all clauses of POP_{n-1} from those of POP_n, and Theorem 3.1 follows by induction. □

Lemma 3.2. *For any $j \leq n-1$, the clause $C_{n-1}(j)$ is derivable in polynomial length from $C_n(1), \ldots, C_n(n)$, $A(1,n,j), \ldots, A(n-1,n,j)$, and $B(j,n)$.*

Proof. For any $i \neq j$, we can do the inference step

$$\frac{C_n(j) \qquad A(i,n,j)}{C_{n-1}(j) \vee \overline{x}_{in}} \tag{3.1}$$

by resolving over x_{nj}. We can interpret this step as follows. Suppose e_j is not minimal in S_n and none of the elements $\{e_1, \ldots, e_{j-1}, e_{j+1}, \ldots, e_{n-1}\}$ is smaller then e_j (this corresponds to $C_{n-1}(j)$ being false), then e_n has to be smaller than e_j. This implies that $e_i \not< e_n$, since if on the contrary $e_i < e_n$, then we would also have $e_i < e_j$ by transitivity contrary to assumption. Thus, x_{in} must be false in this case.

To derive $C_{n-1}(j) \vee \overline{x}_{jn}$, we instead resolve

$$\frac{C_n(j) \qquad B(j,n)}{C_{n-1}(j) \vee \overline{x}_{jn}} \tag{3.2}$$

over x_{nj}. In this way, we can derive $C_{n-j} \vee \overline{x}_{in}$ for all i.

Using these clauses one by one and resolving over the variables $x_{1n}, x_{2n}, \ldots, x_{n-1,n}$ in that order, we get the resolution derivation

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{C_n(n) \qquad C_{n-1}(j) \vee \overline{x}_{1,n}}{C_{n-1}(j) \vee \bigvee_{\ell=2}^{n-1} x_{\ell n}} \qquad C_{n-1}(j) \vee \overline{x}_{2n}}{C_{n-1}(j) \vee \bigvee_{\ell=3}^{n-1} x_{\ell n}} \qquad C_{n-1}(j) \vee \overline{x}_{3n}}{C_{n-1}(j) \vee \bigvee_{\ell=4}^{n-1} x_{\ell n}}}{\vdots}}{\cfrac{C_{n-1}(j) \vee x_{n-1,n} \qquad C_{n-1}(j) \vee \overline{x}_{n-1,n}}{C_{n-1}(j)}} \tag{3.3}$$

which also clearly has polynomial length, and which derives $C_{n-1}(j)$. This concludes the proof of Lemma 3.2. □

Note that all the clauses appearing in the resolution refutation have width $O(n)$, which is $O(\sqrt{N})$ if we let N denote the number of variables in POP_n.

4 3-CNF Versions of Wide CNF Formulas

So now we are done with the first half of our program for today, and if we can also prove a lower bound on width for LOP_n we will be done. Except...

Except that if we look closer at the definitions of POP_n and LOP_n, these formulas have clauses of width n, but Theorem 2.1 is stated for 3-CNF formulas. Intuitively, if the width of the clauses themselves is already n, then we would have to work very hard indeed to prove a lower bound on $W(LOP_n \vdash \bot)$ that could give anything interesting in (1.1).

So what can we do? One approach would be to proceed as for the pigeonhole principle formulas, where we defined "sparser versions" with constant-width clauses and showed that these formulas were also hard. That is in fact something one can do for ordering principles as well,[1] although one has to be a bit more careful, but that is not what we are going to do today.

INSTRUCTOR'S COMMENT 4: *Well, in this edition of the course this is* exactly *what we will do, or already did, in fact.*

[1] See Lemma 8.17 in [SBI04] for the details.

$W(F \vdash \bot) = w \Rightarrow L(F \vdash \bot) = n^{O(w)}$

How tight is this argument?

For all "classic" formulas in the proof cplx literature, we have that as soon as $W(F \vdash \bot) = o(\sqrt{n})$ it holds that $L(F \vdash \bot) = n^{O(1)}$ with exponent independent of width (except possible for padding arguments, but no lower bounds $n^{\Omega(w)}$).

[CEI '96] showed $W(F \vdash \bot) \leq w \Rightarrow S_{PC}(F \vdash \bot) = n^{O(w)}$ but asked whether tighter simulation of resolution by PC possible (also asked by [Beame & Pitassi '98])

[AFT '11] showed (kind of) that if $W(F \vdash \bot)$, then a standard CDCL SAT solver, except that decisions are made randomly, will decide F in time $n^{O(w)}$ asymptotically almost surely.

Is this tight?

YES, these results are all tight up to (not too bad) constants in the exponent

Also beyond resolution and polynomial calculus for Sherali-Adams and sum-of-squares (but we will focus on resolution and PC)

THEOREM [Atserias, Lauria, Nordström '16]

Let $w = w(n) = O(n^c)$ for $c \in \mathbb{R}^+$, $c < 1/2$

Then there exist 3-CNF formulas $F_{n,w}$ with $O(wn)$ clauses over $O(n)$ variables such that

(1) $F_{n,w}$ has a resolution refutation π with

$$\boxed{\begin{aligned} L(\pi) &= n^{O(w)} \\ W(\pi) &= O(w) \\ Sp(\pi) &= O(w) \end{aligned}}$$

← TODAY

(2) Any refutation of $F_{n,w}$ in resolution, PCR, or Shenali-Adams must have size $n^{\Omega(w)}$

Why is such a result hard to prove?

(i) Ben-Sasson & Wigderson:

Width $\geq w \Rightarrow$ length $\exp(w^2/n)$

Gives nothing for $w \leq \sqrt{n}$

And can never be larger than 2^w —

we want n^w

(ii) Random restriction method

Formula family $\{F_n\}_{n \in \mathbb{N}^+}$

Random restrictions ρ_n s.t.

$F_n \upharpoonright \rho_n = F_m$ (after renaming of variables)

Set random literal to true $\Rightarrow$

satisfies and removes $\frac{w}{2n}$ fraction of wide clauses in $\pi_n : F_n \vdash \bot$

Suppose proof in size/length $S(\pi) = S$

After r rounds of random restrictions expect $\leq \left(1 - \frac{w}{2n}\right)^r S$ wide clauses

Choose $r = \frac{2n}{w} \log S \Rightarrow$ expected number $\approx 0 \Rightarrow$ exists good ρ

But argue separately that $F_n \wedge \rho = F_m$ still requires large width

(Our PHP lower bound can be cast in this framework)

So we must have killed formula pretty much completely $\Rightarrow r \gtrsim n$

$\Rightarrow \log S \gtrsim w \Rightarrow S \gtrsim 2^w$

But this is the limit — we cannot have $r > n$, since then there is no formula left after the restriction.

So standard techniques seem stuck at $2^{\Omega(n)}$ lower bounds.

REFINED RANDOM RESTRICTION METHOD

Choose $\{F_n\}_{n \in \mathbb{N}^+}$ & ρ_n s.t. $F_n \wedge \rho_n = F'_m$

Analyse clauses in $\pi_n : F_n \vdash \bot$

(a) If C very wide, $C \wedge \rho = 1$ almost surely

NEW! (b) If C only somewhat wide, $W(C \wedge \rho)$ small

So get $\pi_n \wedge \rho_n : F'_m \vdash 1$ in small width if π_n short

Argue separately $W(F'_m \vdash 1)$ large

So π_n must have been long

RELATIVIZED PIGEONHOLE PRINCIPLE FORMULAS

"There is a way to choose k out of n pigeons so that these pigeons can fly to $k-1$ holes in one-to-one fashion"

More formally: There exist partial functions

$p : [k] \to [n]$

$q : [n] \to [k-1]$

p is one-to-one & defined on $[k]$

q is one-to-one on & defined on range(p)

Encode as formula $RPHP^{k,n}_{k-1}$

Variables

$p_{u,v}$	"is $p(u) = v$?"
$q_{v,w}$	"is $q(v) = w$?"
r_v	"is v in the range of p?"

Clauses

(R1)	$p_{u,1} \vee p_{u,2} \vee \cdots \vee p_{u,n}$	$u \in [k]$
(R2)	$\overline{p_{u,v}} \vee \overline{p_{u',v}}$	$u \neq u' \in [k], v \in [n]$
(R3)	$\overline{p_{u,v}} \vee r_v$	$u \in [k], v \in [n]$
(R4)	$\overline{r_v} \vee q_{v,1} \vee q_{v,2} \vee \cdots \vee q_{v,k-1}$	$v \in [n]$
(R5)	$\overline{r_v} \vee \overline{r_{v'}} \vee \overline{q_{v,w}} \vee \overline{q_{v',w}}$	$v \neq v' \in [n]$, $w \in [k-1]$

(R1)-(R2): $p : [k] \to [n]$ injective / one-to-one

(R3): range of p

(R4)-(R5): $q : [n] \to [k-1]$ defined and injective on range of p

But we need to make this into 3-CNF

So convert (R2) to

$$p_{u,1} \vee p_{u,2} \vee y_{u,2}$$
$$\overline{y}_{u,2} \vee p_{u,3} \vee y_{u,3}$$
$$\vdots$$

(R1') $\quad \overline{y}_{u,r} \vee p_{u,r+1} \vee y_{u,r+1}$

$$\vdots$$
$$\overline{y}_{u,n-2} \vee p_{u,n-1} \vee p_{u,n}$$

and (R4) to

$$\overline{r}_{v} \vee q_{v,1} \vee z_{v,1}$$

(R4') $\quad \overline{z}_{v,w} \vee q_{v,w+1} \vee z_{v,w+1}$

$$z_{v,k-3} \vee q_{v,k-2} \vee q_{v,k-1}$$

and (R5) to

$$\overline{r}_{v} \vee \overline{r}_{v'} \vee r_{v,w'}$$

(R5') $\quad \overline{r}_{v,v'} \vee \overline{q}_{v,w} \vee \overline{q}_{v,w'}$

Denote this formula $\widetilde{RPHP}_{k-1}^{k,n}$

$O(kn^2)$ clauses over $O(n^2)$ variables

(Would get slightly better bound for 4-CNF formulas if we didn't convert (R5))

Game plan

(i) Define restriction $\widetilde{RPHP} \rightsquigarrow PHP$ (but with holes)

(ii) Prove width lower bound for $\widetilde{PHP}$

(iii) Prove that if $\pi : \widetilde{RPHP} \vdash \bot$ is short, then $\pi|_\rho : \widetilde{PHP} \vdash \bot$ has small width

(iv) Conclude $d(\pi) = n^{-\Omega(k)}$

(v) Show $W(\widetilde{RPHP}^{k,n}_{k-1} \vdash \bot) \leq O(k)$

RANDOM RESTRICTION DISTRIBUTION $\mathcal{D}$

Pick $S \subseteq [n]$, $|S| = k$ uniformly at random

Fix arbitrary bijection $\psi : [k] \to S$

Define ρ on variables as follows

$$\rho(r_v) = \begin{cases} \top & \text{if } v \in S \quad [v \text{ is PICKED}] \\ \bot & \text{otherwise} \end{cases}$$

$$\rho(r_{v,v'}) = \rho(r_v) \wedge \rho(r_{v'}) \qquad v \neq v'$$

$$\rho(p_{u,v}) = \begin{cases} \top & \text{if } v = \psi(u) \\ \bot & \text{otherwise} \end{cases}$$

$\rho(y_{u,v})$ selected so as to satisfy (R1')

$$\rho(q_{v,w}) = \begin{cases} * & \text{if } v \in S \\ \text{random } b_v \in \{\top, \bot\} & \text{otherwise} \end{cases}$$

$$\rho(z_{v,w}) = \begin{cases} * & \text{if } v \in S \\ \text{same random } b_v & \text{otherwise} \end{cases}$$

ρ converts $\widetilde{RPHP}^{k,n}_{k-1}$ into 3-CNF version of pigeonhole principle

Rename picked pigeons in S to $\{1, \ldots, k\}$
Then clauses in $\widetilde{RPHP}^{k,n}_{k-1} \restriction_S$ are

$\widetilde{PHP}^{k}_{k-1}$

$\overline{q_{v,1}} \vee z_{v,1}$ $\quad v \in [k]$

$\overline{z_{v,w}} \vee \overline{q_{v,w+1}} \vee z_{v,w+1}$ $\quad v \in [k]$, $w \in [k-2]$

$\overline{z_{v,k-2}} \vee \overline{q_{v,k-2}} \vee \overline{q_{v,k-1}}$ $\quad v \in [k]$

$\overline{q_{v,w}} \vee \overline{q_{v',w}}$ $\quad v \neq v' \in [k]$, $w \in [k-1]$

Define the PIGEON-WIDTH of a clause D $|\{ v \mid v \in [k], \text{ exist } q_{v,w} \text{ or } z_{v,w} \in \text{Vars}(D)\}|$ as the number of pigeons mentioned in D

LEMMA 1 Every resolution refutation $\pi : \widetilde{PHP}^{k}_{k-1} \vdash \bot$ has pigeon-width $\geq k-1$

Proof Exercise. Use Prosecutor-Defendant game, but adapt it to pigeon-width. Show that if Prosecutor remembers too few pigeons, then Defendant has winning strategy

To get analogous result for polynomial calculus, need degree lower bound

LEMMA 2 Every polynomial calculus refutation of $\widetilde{PHP}^{k}_{k-1}$ has pigeon-degree $\lceil \frac{k-1}{2} \rceil$

Proof Requires a bit more work
Use lower bound in [Razborov '98] plus some rewriting, since encodings don't match.
Or use [ALN '24] but then (much) worse constants.

LEMMA 3

Let $k, \ell, n \in \mathbb{N}^+$ such that $n \geq 16$ and $\ell \leq k \leq n/4\log n$. Let A be clause (or monomial) over $Vars(\widetilde{RPHP}^{k,4}_{k-1})$ and let $\mathcal{D}$ random restriction distribution $\mathcal{D}$. Then pigeon-width of $A\restriction_\rho$ is less than ℓ width probability at least $1 - (4k\log n)^k / n^\ell$

Defer proof

Proof of lower bound (for resolution)

Let π: $\widetilde{RPHP}^{k,4}_{k-1} \vdash \bot$ have length/size S

Hit π with random restriction $\rho \sim \mathcal{D}$

Since restrictions preserve resolution refutations, $\pi\restriction_\rho$ is a refutation of

$$\widetilde{RPHP}^{k,4}_{k-1} = \widetilde{PHP}^{k}_{k-1}.$$

This refutation $\pi\restriction_\rho$ must have pigeon-width at least $k-1$ with probability 1.

But by Lemma 3, each individual clause in $\pi\restriction_\rho$ is unlikely to have this large width

Set $\ell = k-1$ and take union bound of at most S clauses in $\pi\restriction_\rho$ to compute that the probability that $\pi\restriction_\rho$ has pigeon-width $\geq k-1$ is at most

$$S \cdot (4k\log n)^k / n^{k-1} \qquad (*)$$

Since $(*)$ must be ≥ 1, we get

$$S(\pi) = S \geq n^{k-1} / (4k\log n)^k$$

□

LEMMA 4

$\widetilde{RPHP}^{k+n}_{k+1}$ has tree-like resolution refutations on length/size $O(k^k n^k)$ and (standard) width $2k+1$

Proof *Exercise.* Use (standard) Prosecutor-Defendant game plus equivalence to resolution.

Prosecutor strategy

- first figure out $\rho : [k] \to [n]$
  requires memory $\approx k$
- Remember range(ρ)
- Finally force contradiction for
  $g : [k] \to [k-1]$ on range(ρ)
  requires additional memory $\approx k$ ▨

Completes proof of ATR theorem for resolution. For Sherali-Adams and PC nothing much changes, except

- work a bit harder to show degree lower bound for $\widetilde{PHP}^k_{k-1}$
- For Sherali-Adams, argue that it can simulate resolution efficiently.

Proof of Lemma 3

AZN X

Assume A clause – the proof for monomials is completely analogous

Let $v_1, v_2, \ldots, v_r$ be pigeons mentioned in A sorted in order

Let $a_1, a_2, \ldots, a_r$ be literals in A such that a_i mentions pigeon v_i

We do case analysis:

(1) LARGE PIGEON-WIDTH $r > 2k \log n$

(2) SMALL PIGEON-WIDTH $r \leq 2k \log n$

Want to argue

Case 1: $A|_\rho = 1$ very likely

Case 2: Not sure if $A|_\rho = 1$, but if not, at least ρ shrinks A a lot

$r > 2k \log n$

$$\Pr[A|_\rho \neq 1] \leq \Pr[\rho(a_i) \neq * \text{ for } i \in \lceil 2k \log n \rceil]$$

$$= \prod_{i=1}^{\lceil 2k \log n \rceil} \Pr[\rho(a_i) \neq * \mid \rho(a_j) \neq * \text{ for } j < i]$$

Happens if (i) v_i picked or (ii) random value to a_i wrong

$$\leq \prod_{i=1}^{\lceil 2k \log n \rceil} \Pr[\rho(a_i) \neq * \mid v_j \notin S \text{ for } j < i]$$

(i) most likely if no v_j, $j < i$ picked

(ii) only if (i) doesn't happen, and if so, independent coin flip

$$\leq \prod_{i=1}^{\lceil 2k \log n \rceil} \left(\frac{1}{2} + \frac{k}{n-i+1} \right) < \left(\frac{5}{8} \right)^{2k \log n} < \frac{1}{n^k}.$$

$r \leq 2k \log n$

Very few pigeons are picked, so it is unlikely that sufficiently many of the few pigeons in A survive

choices S with exactly i pigeons mentioned in A is

$$\binom{r}{i}\binom{n-r}{k-i}$$

Consider all possible intersections of size ℓ between S and pigeons mentioned in A. Probability of ℓ surviving pigeons is

$$\sum_{i=\ell}^{k} \binom{r}{i}\binom{n-r}{k-i} \Big/ \binom{n}{k} \leq$$

$$\leq k \binom{\lfloor 2k \log n \rfloor}{k}\binom{n}{k-\ell} \Big/ \binom{n}{k}$$

$$\leq \frac{k \cdot (2k \log n)^k}{k!} \cdot \frac{n!}{(k-\ell)!(n-k+\ell)!} \cdot \frac{(n-k)!\,k!}{n!}$$

$$\leq k (2k \log n)^k \frac{1}{(k-\ell)!} \frac{1}{(n-k)^\ell} <$$

$$< \frac{k (2k \log n)^k}{(n-k)^\ell}$$

Use that $n \geq 16$ and $k \leq n/4\log n$ to get $k \leq n/16$

Show that $k(16/15)^\ell \leq 2^k$ for all $\ell \in [k]$ (not particularly right). Plug this in to get

$$\frac{k(2k\log n)^k}{(n-k)^\ell} \leq \frac{k(2k\log n)^k}{(15n/16)^\ell}$$

$$= k \cdot (16/15)^\ell \cdot \frac{(2k\log n)^k}{n^\ell}$$

$$\leq \frac{(4k\log n)^k}{n^\ell}$$

This concludes the proof