LECTURE 15

$\boxed{\text{Last lecture}}$: Size-space trade-offs in proof complexity from time-space trade-offs in pebbling

Roughly: If matching upper and lower bounds for black & black-white pebbling [or trade-offs for "not-too-white" black-white pebbling] then:
- size-space trade-offs with same parameters for resolution
- size-space trade-offs with log factor loss for polynomial calculus

EXAMPLE RESULTS

*) also need $g(n) = O(n^{1/7})$

THEOREM 1 [BN11] (Using pebbling results in [Nordström '12])

Let $g(n) = \omega(1)$ be arbitrarily slowly growing function* and fix any $\varepsilon > 0$. Then $\exists$ explicitly constructible 6-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that

(a) F_n refutable in total space $O(g(n))$ in resolution and PC

(b) F_n refutable in simultaneous size $O(n)$ and total space $O((n/(g(n))^2)^{1/3})$ in resolution and PC

(c) Any resolution refutation / PCR refutation in clause space / monomial space $O((n/((g(n))^2 \log n))^{1/3-\varepsilon})$ must have superpolynomial size in resolution and PCR

THEOREM 2 [BN11] Using pebbling results in [Nordström '12]

There is a family of explicitly constructible 6-CNF formulas of size $\Theta(n)$ such that

(a) F_n refutable in total space $O(n^{1/11})$ in resolution and PC

(b) F_n refutable in simultaneous size $O(n)$ and total space $O(n^{3/11})$ in resolution and PC

(c) Any resolution refutation PCR refutation in clause space monomial space at most $n^{2/11}/(10 \log n)$ must have size at least $(n^{1/11})!$ in resolution and PCR

$\boxed{\text{Technical core}}$: Lift trade-offs between length and variable space to trade-offs between size and clause space for resolution

For polynomial calculus:

- Use that variable space is the same measure as for resolution
- Do substitution with XOR + random restriction argument

Pebbling formulas just happen to have such nice trade-offs

OPEN PROBLEM 1: Are there other such formulas?

OPEN PROBLEM 2: Can we get tight results also for polynomial calculus?

Strength and weakness of results:

- Upper bounds for total space and syntactic proof systems
- Lower bounds for clause space / monomial space and $\boxed{\text{SEMANTIC PROOF SYSTEMS}}$: Anything implied can be derived in single step

But in this model all formulas are refutable in simultaneous linear size and linear space.

Traditionally, time-space trade-offs look something like

$$(\text{Space}) \cdot (\text{Time}) \geq n^2$$

or stronger

$$(\text{Space}) \cdot \log(\text{Time}) \geq n$$

These results say nothing about superlinear space

Recall question from last lecture:

$$\boxed{\text{If } F \text{ is refutable in length/size } L\text{, can } F \text{ be refuted in length poly}(L) \text{ and linear clause/monomial space } O(S(F)) \text{ simultaneously?}}$$

NO! For regular resolution and resolution [Beame, Beck, & Impagliazzo '12, '16]

Tighter results for resolution
+ polynomial calculus
[Beck, Nordström, & Tang '13]

THEOREM 3 [BNT'13]
For $w = w(n)$ with $3 \leq w(n) \leq n^{1/4}$ there are explicitly constructible 8-CNF formulas $\{F_n\}_{n=1}^{\infty}$ of size $\Theta(n)$ such that

(a) F_n refutable in clause space $O(w \log n)$ and length $\exp(O(w \log n))$ in resolution

(b) F_n refutable in length $n^{O(1)} \exp(w)$ and clause space $\exp(w) + n^{O(1)}$ in resolution

(c) For any PCR refutation π_n over a field $\mathbb{F}$ s.t. $\mathrm{char}(\mathbb{F}) \neq 2$, the proof size is bounded by

$$S(\pi_n) = \left(\frac{\exp(\Omega(w))}{MSp(\pi_n)} \right)^{\Omega\left(\frac{\log \log n}{\log \log \log n}\right)}$$

Fix $\boxed{w = K \log n}$ for suitably large $\overbrace{K}^{\text{constant}}$

Then resolution can refute formulas in

- length $\approx n^K$
- clause space $O(\log^2 n)$

But clause space, say, $n^{k/2}$ causes superpolynomial blow-up in proof size (Need to adjust constants for precise statement)

Beame, Beck, & Impagliazzo have much sharper results for regular resolution

> OPEN PROBLEM 3: Improve the parameters in the trade-offs in Thm 3. Is it possible to extend the much stronger trade-off results for regular resolution also to general resolution? What about exponential trade-offs? *

*) Not for the formulas we talk about today

Trade-off formulas: $\boxed{\text{TSEITIN CONTRADICTIONS}}$

Graph $G = (V, E)$

Charge function $\chi : V \to \{0,1\}$ such that $\sum_{v \in V} \chi(v) \equiv 1 \pmod 2$

(ODD CHARGE)

$$\boxed{PARITY_{v,\chi} = \text{"} \sum_{e \ni v} x_e \equiv \chi(v) \pmod 2 \text{"}}$$

$$= \left\{ \bigvee_{e \ni v} x_e^{1-b_e} \;\middle|\; \sum_{e \ni v} b_e \not\equiv \chi(v) \pmod 2 \right\}$$

Recall $x^b = \begin{cases} x & \text{if } b = 1 \\ \overline{x} & \text{if } b = 0 \end{cases}$

$$\boxed{Ts(G,\chi) = \bigwedge_{v \in V} \text{PARITY}_{v,\chi}}$$

Suppose G is connected.
Then $Ts(G,\chi)$ unsatisfiable
$\Updownarrow$
χ odd-charge function

Exact charge function does not matter — only whether charge is odd or even
Can use substitution to convert between different odd-charge functions

We know: $\boxed{G \text{ expander} \Rightarrow Ts(G,\chi) \text{ exponentially hard}}$

But we want only moderately hard formulas. Use rectangular grids with w rows and ℓ columns $(w \ll \ell)$

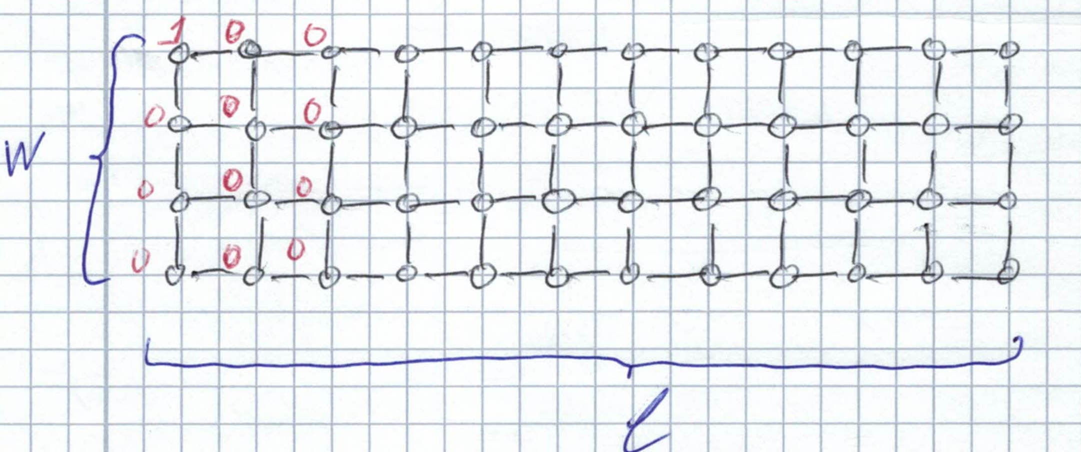

Topmost left vertex charge 1, all others 0

(We will need to tweak this a bit, but this is the idea

PROPOSITION 4

If F has resolution refutation in depth d, then tree-like resolution can refute F in simultaneous

- length $2^{d+1} - 1$
- clause space $d + 2$

Proof sketch

- Make resolution refutation tree-like — does not increase space
- # nodes in proof DAG $\leq 2^{d+1} - 1$
- Black-pebble proof DAG to get the stated bounds ▨

PROPOSITION 5

Let G be $w \times \ell$ grid and let $\chi : V \rightarrow \{0,1\}$ be odd-charge function. Then $Ts(G, \chi)$ can be refuted in depth $O(w \log \ell)$

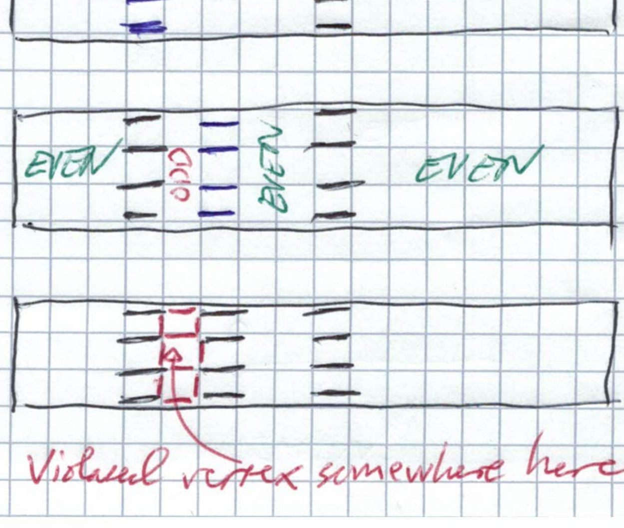

Proof sketch

Use that tree-like resolution
$\Updownarrow$
decision tree

Do binary search

Query middle column
– disconnects graph

Recurse on odd-charge component

$O(\log \ell)$ recursive steps

w queries per step ▨

PROPOSITION 6

Let G $w \times l$ grid and χ odd-charge. Then $Ts(G, \chi)$ can be refuted in simultaneous length $lw\, 2^{O(w)}$ and clause space $2^{O(w)}$.

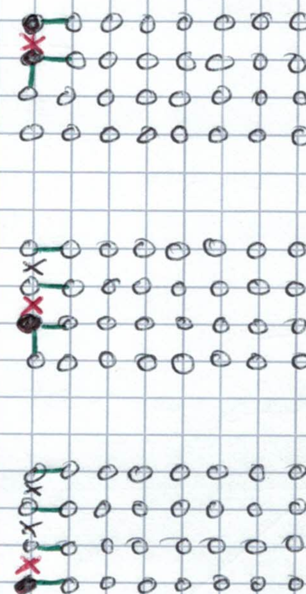

Proof Order edges from right to left and from top to bottom in each column.

Resolve all clauses containing top left vertical edge.

Keep resolvent in memory.

Download all axioms for third vertex in first column.

FACT Resolving over all variables in fixed order yields resolution refutation.

DAVIS-PUTNAM RESOLUTION or VARIABLE ELIMINATION

Prove, e.g., by induction over # variables

In this case: ~~Invariant~~ is that sum of charges of cut edges is odd

Space $\leq w+1$ edges = variables $\Rightarrow \leq 2^{w+1}$ clauses

Length wl vertices $\times\ 2^{O(w)}$ steps per vertex $\square$

w-parameter is tree-width of graph.
This is special case of more general result.

How to prove trade-off?

HIGH-LEVEL IDEA

1. Formalize notion of PROGRESS of proof

2. Divide proof into large number of equal-sized EPOCHS

3. Prove the following claims:
   (a) If epochs are small, then no single epoch makes very much progress
   (b) If space is small, then not much progress can be carried over from one epoch to the next
   (c) To refute formula, proof needs to make substantial progress summed over all epochs

4. Hence, a proof that is too short and uses too little space cannot refute the formula

Fix grid graph with w rows and ℓ columns

Vertices indexed by (i,j) $i \in [w]$, $j \in [\ell]$

Edges from (i,j) to $(i, j\pm 1)$, $(i \pm 1, j)$

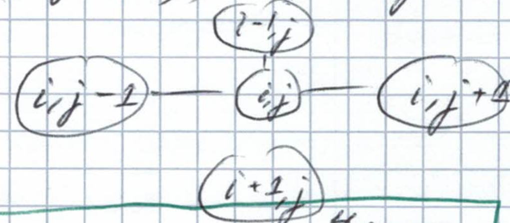

Choose w so that $\boxed{\log \ell \leq w \leq \sqrt[4]{\ell}}$

Do binary XOR substitution in $Ts(G, \chi)$ to get $Ts(G,\chi)[\oplus_2]$

Same thing as letting $G' = G$ with two copies of every edge and taking $Ts(G', \chi)$

We will prove (or at least sketch proof) for resolution that $Ts(G', \chi)$ does not have resolution refutations in short length and small space simultaneously

Note that upper bounds in Props 5 & 6 still hold just replacing w by $2w$

Will not talk about proof for PCR — this is much more complicated

Let ρ be random restriction that

- picks one copy of edge uniformly and independently at random
- fix this edge to $\top$ or $\bot$ uniformly and independently at random

Then $Ts(G', \chi)\restriction_\rho = \left(Ts(G,\chi)[\oplus]\right)\restriction_\rho = Ts(G, \chi)$

except for renaming variables & flipping polarities

Define/recall complexity measure for clauses derived from Tseitin formulas $Ts(G,\chi)$

$$\boxed{\mu(C) = \min \left\{ |S| : \bigwedge_{v \in S} PARITY_{v,\chi} \vDash C \right\}}$$

Properties

- $\mu(A) = 1$ for $A \in Ts(G,\chi)$
- $\mu(\bot) = |V(G)|$ (if G connected)
- SUBADDITIVITY $\mu(C \vee D) \leq \mu(C \vee x) + \mu(D \vee \bar{x})$

If $S \subseteq V = V(G)$ is such that

$$\bigwedge_{v \in S} PARITY_{v,\chi} \vDash C \quad \text{and} \quad |S| = \mu(C)$$

call S a CRITICAL SET for C

Recall definition of BOUNDARY

$$\boxed{\partial S = \{ (u,v) \in E \mid u \in S,\ v \in V \setminus S \}}$$

LEMMA 7

Let C clause over variables of $Ts(G,\chi)$ and suppose S critical set for C. Then

(a) S is a connected set

(b) $\{ x_e \mid e \in \partial S \} \subseteq Vars(C)$

Proof Suppose $S = S_1 \cup S_2$ with no edges between S_1 & S_2

$$\bigwedge_{v \in S_i} PARITY_{v,\chi} \nvDash C$$

Fix α_i s.t.

$$\alpha_i\left(\bigwedge_{v \in S_i} PARITY_{v,x}\right) = T$$

$$\alpha_i\ (C) \neq T$$

Note that α_1 & α_2 assigns disjoint sets of variables, so $\alpha_1 \cup \alpha_2$ is valid assignment.

$$(\alpha_1 \cup \alpha_2)\left(\bigwedge_{v \in S} PARITY_{v,x}\right) = T$$

$$(\alpha_1 \cup \alpha_2)\ (C) \neq T$$

Extend to α s.t. $\alpha(C) = \bot$ $\lightning$

(b) We did this flipping argument in lecture on Tseitin formula lower bounds

Suppose exists $e \in \partial S$ such that $x_e \notin Vars(C)$

$e = (u,w)$, $u \in S$

$$\bigwedge_{v \in S \setminus \{u\}} PARITY_{v,x} \nvDash C$$

Fix α s.t. $\alpha\left(\bigwedge_{v \in S \setminus \{u\}} PARITY_{v,x}\right) = T$

$\alpha\ (C) = \bot$

$\alpha\ (PARITY_{u,x}) = \bot$ by necessity

Flip α on x_e to get α'

Now $\alpha'(PARITY_{u,x}) = T$

so

$$\alpha'\left(\bigwedge_{v \in S} PARITY_{v,x}\right) = T$$

But since $x_e \notin Vars(C)$, we have $\alpha'(C) = \bot$ $\lightning$ $\square$

Fix $t_0 = w^4$

Say that C has MEDIUM COMPLEXITY if $t_0 \leq \mu(C) \leq |V|/4 = \frac{w\ell}{4}$

(note that $w \geq 4$ since $w \geq \log \ell$ and $\ell \to \infty$)

Say that C has COMPLEXITY LEVEL i for $i \in \mathbb{N}$ if

$$t_0 2^i < \mu(C) \leq t_0 \cdot 2^{i+1}$$

We have $\approx \log \ell$ complexity levels

OBSERVATION 8

By subadditivity, in any resolution refutation $\pi : Ts(G, \chi)$ there are clauses of all complexity levels.

We want to argue that if

$\pi' : Ts(G', \chi) \vdash \bot$

has small length L and clause space s then can find ρ in support of our random restriction distribution such that

$\pi' \restriction_\rho$ is refutation of $Ts(G', \chi)\restriction_\rho = Ts(G, \chi)$

where not all complexity levels appear

For this to work, need to prove for C_1 & C_2 of complexity levels $i_1 \neq i_2$

Dependent events!

(i) $\Pr_\rho[C_i \restriction_\rho \neq T] \leq \exp(-w)$

(ii) $\Pr[\text{For } i=1,2\ C_i \restriction_\rho \neq T] \lesssim \Pr[C_1 \restriction_\rho \neq T] \cdot \Pr[C_2 \restriction_\rho \neq T]$

LEMMA 9

$$\Pr_{\rho}\left[\, C\restriction_{\rho} \text{ has } \geq w \text{ variables} \,\right] \leq (3/4)^{w}$$

Proof If $W(C) < w$ there is nothing to prove. Suppose e & e' twin edges

If $x_e \in Vars(C)$, $x_{e'} \notin Vars(C)$, then

$$\Pr\left[\rho(x_e) = T\right] = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$$

If $x_e, x_{e'}$ both in $Vars(C)$

$$\Pr\left[\text{literals over } x_e \text{ or } x_{e'} \text{ don't satisfy } C \text{ after } \rho\right] = \frac{1}{2} \leq \left(\frac{3}{4}\right)^2$$

So $\Pr\left[\, C\restriction_{\rho} \neq T \,\right] \leq \left(\frac{3}{4}\right)^{W(C)} \leq \left(\frac{3}{4}\right)^{w}$ ∎

COROLLARY 10

For any clauses $C_1, C_2, \ldots, C_k$ it holds that

$$\Pr_{\rho}\left[\begin{array}{l}\text{For all } i \;\; C_i\restriction_{\rho} \neq T \\ \text{and} \\ \{ C_i\restriction_{\rho} \mid i \in [k] \} \text{ contains } \geq w \text{ variables}\end{array}\right] \leq \left(\frac{3}{4}\right)^{w}$$

Proof Just view $C_1, C_2, \ldots, C_k$ as a big clause with all literals concatenated.

No literal must be assigned to true by ρ. Use proof of Lemma 9 ∎

Fix G to be $w \times \ell$ grid,
$\log \ell \leq w \leq \sqrt[4]{\ell}$, ℓ large enough

Consider set $S \subseteq V = V(G)$
Say that column j in grid G is
- $\boxed{\text{full}}$ in S if all vertices in column $\subseteq S$
- $\boxed{\text{empty}}$ in S if no — " —
- $\boxed{\text{partial}}$ otherwise

Any clause C of medium complexity has width $W(C) \geq w$ because the boundary of any critical set S_C for C has $|\partial S| \geq w$ edges

LEMMA 11
For any $S \subseteq V$ s.t. $w^4 \leq |S| \leq |V|/4$ it holds that $|\partial S| \geq w$

Proof If S has w partial columns, then ∂S has w vertical edges, so suppose S has less than w partial columns.
Then $|S| - w^2 \geq w^4 - w^2 > 0$ vertices are in full columns, so S has a full column

Since $|S| \leq |V|/4$, at most $\frac{1}{4}$ fraction of columns are full
By assumption, at most $\frac{w}{\ell} \leq \frac{1}{w^3} \leq \frac{1}{4}$ fraction of columns are partial.

So S has empty columns, and also full columns
Hence in every row there is at least one edge in boundary, and $|\partial S| \geq w$ $\square$

This means that a single medium-complex clause C is likely to get satisfied by ρ.

To get a kind-of-independence result, we prove that clauses $C_1, C_2, \ldots, C_k$ of distinct complexity levels contain a total of $\Omega(kw)$ distinct variables

LEMMA 12

Let $C_1, C_2, \ldots, C_k$ ($k \leq w$) be clauses of distinct and increasing complexity levels as witnessed by critical sets $S_1, S_2, \ldots, S_k$.

Then $\left| \bigcup_{i=1}^{k} \partial(S_i) \right| = \Omega(kw)$.

Proof Take every third set in $S_1, \ldots, S_k$ (if necessary) to get $S_1', S_2', \ldots, S_{k'}'$ such that

$$t_0 \leq S_1'$$
$$4|S_i'| \leq |S_{i+1}'|$$
$$|S_{k'}'| \leq |V|/4$$
$$k' \geq \lceil k/3 \rceil$$

S_i' and S_{i+1}' are at least 2 complexity levels apart

If some S_i' has $\geq w^2$ partial columns, then $\partial(S_i') \geq w^2 \geq kw$ and we are done, so suppose every S_i' has $\leq w^2$ partial columns.

We want to show that every row in grid has at least $k'-1$ horizontal edges.

Fix a row j

Let l_i = column of leftmost vertex of S_i' in row j

Let r_i = column of rightmost vertex of S_i' in row j

If $l_i \neq 1$, there is boundary edge $(l_i - 1, l_i)$

if $r_i \neq l$, there is boundary edge $(r_i, r_i + 1)$

Let SIGNATURE of horizontal edge be column of left endpoint — uniquely determines edge in row

We take sequences $(l_i - 1)_{i=1}^{k'}$ and $(r_i)_{i=1}^{k'}$ and apply following proposition

PROPOSITION 13

Let $(a_i)_{i=1}^{k}$ and $(b_i)_{i=1}^{k}$ be integer sequences such that for all i it holds that

$$|b_i - a_i| \geq 1$$

$$|b_{i+1} - a_{i+1}| \geq 2 |b_i - a_i|$$

Then $$\left| \bigcup_{i=1}^{k} \{a_i, b_i\} \right| \geq k+1$$

<u>Proof of proposition</u> Exercise. Intuitively, the intervals cannot over-lap too much because of the exponentially increasing sizes. $\boxed{}$

If $(l_i - 1)_{i=1}^{k'}$ and $(r_i)_{i=1}^{k'}$ satisfy conditions of Proposition 13, then we get $k'+1$ distinct numbers. Remove 1 and l.

Still guaranteed $k' - 1 = \lceil \frac{k}{3} \rceil - 1$ horizontal edges

It remains to prove that $(l_i - 1)_{i=1}^{k'}$ and $(r_i)_{i=1}^{k'}$ satisfy conditions.

Let f_i = # full columns in S'_i.
We have

$$f_i \leq r_i - l_i \leq f_i + w^2 \qquad (1)$$

since every column between l_i and r_i is non-empty (since S'_i is connected) and there are $\leq w^2$ partial columns.
There are at most $|S'_i|$ and at least $|S'_i| - w^3$ elements in full columns, so

$$\frac{|S'_i| - w^3}{w} \leq f_i \leq \frac{|S'_i|}{w} \qquad (2)$$

From (1) and (2) we get

$$\frac{|S'_i|}{w} - w^2 \leq r_i - l_i \leq \frac{|S'_i|}{w} + w^2 \qquad (3)$$

from which it follows that

$$\frac{r_{i+1} - l_{i+1}}{r_i - l_i} \geq \frac{|S'_{i+1}| - w^3}{|S'_i| + w^3}$$

$$\geq \frac{4|S'_i| - w^3}{|S'_i| + w^3}$$

(assuming $w \geq 2$)

$$\geq \frac{4 - 1/w}{1 + 1/w} \geq 2$$

$\blacksquare$

LEMMA 14

If $\mathcal{M}$ is a set of clauses, then

$$\Pr_{\rho}\left[\mathcal{M}\restriction_{\rho} \text{ has clauses of } k \text{ distinct complexity levels}\right] \leq \left(|\mathcal{M}|c^{w}\right)^{k}$$

for some $c \in [\frac{3}{4}, 1)$

Proof Fix a subset of k clauses $C_1, \dots, C_k$ of $\mathcal{M}$

If $C_1\restriction_\rho, \dots, C_k\restriction_\rho$ have k distinct complexity levels, then by Lemma 12 they contain $\Omega(kw)$ distinct variables

By Corollary 10, this probability is bounded by $\left(\frac{3}{4}\right)^{\Omega(kw)} \leq c^{kw}$ for some $c < 1$

(With some more care, one can get $c = \frac{3}{4}$)

There are $\binom{|M|}{k} \leq |M|^k$ subsets of k clauses in $\mathcal{M}$, so by a union bound we get probability

$$\leq |M|^k \cdot c^{kw} = \left(|M| \cdot c^{w}\right)^k$$

$\boxtimes$