DD3501 Current Research in Proof Complexity Oct 31, 2011

Lecture 3

Lecturer: Jakob Nordström *Scribe: Sangxia Huang*

1 Quick Recap of Last Lecture

In the previous lecture we proved that large resolution width implies large resolution length. More specifically, for general resolution, we have the following theorem by Ben-Sasson and Wigderson, which was published in [BW99] and later appeared as a full-length journal version in [BW01].

Theorem 1.1 ([BW01]). *For any unsatisfiable CNF formula F it holds that*

$$W(F \vdash \bot) \leq W(F) + \sqrt{8n \ln L(F \vdash \bot)} \ , \tag{1.1}$$

where n is the number of variables in F.

Today, we will use Theorem 1.1 to prove exponential lower bounds on proof length in resolution. The proof we do is also from [BW01]. This lower bound is one of the early classics in proof complexity, proved already in 1985, but we will reprove it using the tools developed in [BW01]. We remark that one way of viewing this result from an applied perspective is that this shows that SAT solvers based on resolution cannot possibly solve certain formulas efficiently.

Since our goal is to prove lower bounds on proof length, it will be more convenient to use the following corollary, which is easily verified to be an immediate consequence of Theorem 1.1.

Corollary 1.2. *For any unsatisfiable CNF formula F it holds that*

$$L(F \vdash \bot) \geq \exp\left(\frac{(W(F \vdash \bot) - W(F))^2}{8n}\right), \tag{1.2}$$

where n is the number of variables in F.

2 A Lower Bound for Pigeonhole Principle Formulas

The pigeonhole principle says that if $m > n$, then there is no way to fit m pigeons into n holes while having at most one pigeon in each hole. To formulate (the negation of) this statement as a CNF formula, we let variables x_{ij} encode whether pigeon i sits in hole j. The formula PHP^m_n consists of the following clauses:

$$P^i = \bigvee_{j=1}^{n} x_{ij} \qquad \text{for } i \in \{1, 2, \cdots, m\}, \tag{2.1}$$

$$H^{ii'}_j = \overline{x}_{ij} \vee \overline{x}_{i'j} \qquad \text{for } i, i' \in \{1, 2, \cdots, m\}, i \neq i', j \in \{1, 2, \cdots, n\}. \tag{2.2}$$

We call P^i a *pigeon axiom*, which says that "pigeon i sits in some hole," and $H^{ii'}_j$ a *hole axiom*, which says that "hole j does not hold both pigeon i and i'." PHP^m_n is unsatisfiable whenever $m > n$. Intuitively, PHP^m_n is hardest when $m = n + 1$, when the formula is *almost* satisfiable, and therefore PHP^{n+1}_n will be today's focus. Our goal is to prove an exponential lower bound on proof length in resolution for these formulas.

The lower bound for PHP^{n+1}_n stated next was proved by Armin Haken, who happens to be the son of Wolfgang Haken (of Four-Colour Theorem fame).

Theorem 2.1 ([Hak85]). *$L(PHP_n^{n+1} \vdash \bot) = \exp(\Omega(n))$.*

Note that PHP_n^{n+1} is a CNF formula with $\Theta\big(n^2\big)$ variables and $\Theta\big(n^3\big)$ clauses, so if we let N be the size of the formula PHP_n^{n+1}, then the lower bound is $\exp\big(\Omega\big(\sqrt[3]{N}\big)\big)$.

Remark 2.2. If one increases the number of pigeons m, the formulas become slightly easier. Buss and Pitassi [BP97] showed that for $m \approx \exp\big(\sqrt{n \log n}\big)$ pigeons, the formulas PHP_n^m have resolution refutations with length polynomial in $\exp\big(\sqrt{n\ \log n}\big)$ (which is significantly better than $\exp(\Omega(n))$). However, the hardness is still exponential in n, being of order $exp(n^{\epsilon})$ even for $m = \infty$ pigeons. This was proven in an amazing paper by Raz in 2001 (journal version in [Raz04]), later simplified and slightly improved by Razborov [Raz01]. (Note that having more than, say, 2^{n^2} or so pigeons does not help anyway, since it will take more than the trivial length upper bound just to look at axiom clauses for all these pigeons, so we never need to consider more than $m \approx 2^{n^2}$ pigeons).

2.1 Problems with Applying Theorem 1.1 and a Work-Around

There are two problems if we try to use Theorem 1.1:

1. There exists refutation for PHP_n^{n+1} of width $\mathrm{O}(n) = \mathrm{O}\left(\sqrt{|\mathrm{Vars}(PHP_n^{n+1})|}\right)$.
2. The width of the formula is $\Omega(n)$.

If we plug this into (1.2), the denominator in the exponent becomes $\Theta\big(n^2\big)$ (the number of variables), which will kill anything in the numerator, but this numerator looks like it will only be at most a constant anyway...

We want to get around this problem by reducing the number of variables and make the formula "sparser." To this end, consider a bipartite graph $G = (U \cup V, E)$, where $|U| = m$, $|V| = n$, and let $N(u)$ be the set of *neighbours* of vertex u. Similar to PHP_n^{n+1}, we define the "graph-version" pigeonhole principle $PHP(G)$ to be the conjunction of the following clauses:

$$P^u = \bigvee_{v \in N(u)} x_{uv} \qquad \text{for } u \in U \tag{2.3}$$

$$H_v^{u,u'} = \overline{x}_{uv} \vee \overline{x}_{u'v} \qquad \text{for } v \in V, u \neq u', u, u' \in N(v) \tag{2.4}$$

We have the following observation.

Observation 2.3. *If $G' = (U \cup V, E')$ has $E' \supseteq E$, then $L(PHP(G) \vdash \bot) \leq L(PHP(G') \vdash \bot)$.*

Proof. Consider assignment ρ setting $x_{uv} = 0$ for all $(u, v) \in E' \backslash E$. We claim that $PHP(G')\restriction_{\rho} = PHP(G)$. This means that from any resolution proof π for $PHP(G')$ we can get a proof for $PHP(G)$ by applying the restriction ρ, and thus follows the upper bound. To see that the claim is true, note that the pigeon axioms are the same because for all $(u, v) \in E' \setminus E$, x_{uv} is set to 0 and thus does not appear after restriction, and the hole axioms involving x_{uv} get satisfied and can therefore be ignored after restricting with ρ. □

If we take $K_{m,n}$ to be the complete bipartite graph with m vertices to the left and n vertices to the right, a moment of thought reveals that $PHP(K_{m,n}) = PHP_n^m$. Therefore we can derive lower bound for $L(PHP_n^{n+1} \vdash \bot)$ by giving lower bound for $L(PHP(G) \vdash \bot)$, where G is a carefully chosen bipartite graph with $n + 1$ vertices to the left and n vertices to the right.

Suppose that the graph G has constant left-degree d. Then $PHP(G)$ is a d-CNF formula with $d(n + 1) = \Theta(n)$ variables, which means that we are potentially "back in business." If we can find such a G with $W(PHP(G) \vdash \perp) = \Omega(n)$, then we would have

$$\begin{aligned} L(PHP_n^{n+1} \vdash \perp) &\geq L(PHP(G) \vdash \perp) \\ &\geq \exp\left(\Omega\left(\frac{(n-d)^2}{d(n+1)}\right)\right) \\ &\geq \exp(\Omega(n)) \end{aligned}$$

since $d = \mathrm{O}(1)$.

2.2 Expander Graphs

Let us now step back a bit and consider why PHP_n^{n+1} is hard. Intuitively, one reason is that it is *almost* satisfiable. Every set of $s \leq n$ pigeons fit perfectly into the holes. This means that no "local argument" can derive a contradiction. Therefore the graph we are looking for should have similar properties, namely it should be

1. a sparse graph (with constant left-degree),
2. but with good connectivity properties.

One of the possible candidates is the following class of graphs.

Definition 2.4 (Bipartite vertex expander graph). The bipartite graph $G = (U \cup V, E)$ is a *bipartite vertex* (d, s, e)*-expander graph*, or just a (d, s, e)-expander for short, if

1. G has constant left-degree d;
2. for each $U' \subseteq U$, $|U'| \leq s$, it holds that $|N(U')| \geq e \cdot |U'|$.

Condition 1 implies that G is sparse, and condition 2 means that G is well-connected in that to disconnect a large part of the graph, one has to cut a lot of edges. In other words, it will be hard to make local argument because there will always be many external variables involved.

One could hope that if G is such a graph, then $PHP(G)$ might be hard. It turns out that to carry out this argument, we need something slightly stronger than Definition 2.4, as described next.

Definition 2.5 (Unique-neighbour expander). $G = (U \cup V, E)$ is a (d, s, e)*-unique-neighbour expander* (or sometimes *boundary expander*) if

1. G has constant left-degree d;
2. for each $U' \subseteq U$, $|U'| \leq s$, it holds that $|\partial U'| \geq e \cdot |U'|$, where $v \in \partial U'$ if $|N(v) \cap U'| = 1$.

We refer to $\partial U'$ as the *boundary* of U', which consists of all vertices $v \in V$ that have only one neighbour in U'. Intuitively, if a set U' expands very well, the sets of neighbours of its vertices could not intersect too much, and thus the boundary $\partial U'$ should not be small. More formally, the following proposition shows that a good vertex expander must also be a good unique-neighbour expander.

Proposition 2.6. *Any* (d, s, κ)*-expander is a* $(d, s, 2\kappa - d)$*-unique-neighbour expander.*

Proof. Left as an exercise. □

2.3 Two Key Lemmas

We will establish Theorem 2.1 by proving the two lemmas below and then combining them.

Lemma 2.7. *For a (d, s, e)-unique-neighbour expander with $e \geq 1$, $W(PHP(G) \vdash \bot) \geq s \cdot e/2$.*

Lemma 2.8. *There is a constant $c > 1$ such that for all n large enough, there are graphs $G = (U \cup V, E)$, $|U| = n+1$, $|V| = n$, that are $(5, n/c, 1)$-unique-neighbour expanders.*

Proof sketch for Lemma 2.8. There are constructive proofs, but we will not go there. Instead, one can use the *probabilistic method* to show that $(5, n/c, 3)$-expanders exist. For all $u \in U$, pick 5 neighbours in V uniformly and independently at random among all $\binom{n}{5}$ subsets. For the right c, such a graph is an expander with overwhelming probability. This is proved by showing that all sets $U' \subseteq U$ with $|U'| \leq s$ are expanding almost surely. And the probability could not be large (indeed, could not be greater than 0) unless such expanders exist. So they do exist.

Some helpful facts for these calculations are:

1. Union bound: $\Pr[A \cup B] \leq \Pr[A] + \Pr[B]$.

2. $\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k$.

3. If $m > n$, then
$$\frac{\binom{n}{k}}{\binom{m}{k}} < \left(\frac{n}{m}\right)^k.$$

The detailed calculations are left as an exercise. □

Remark 2.9. There are also explicit constructions of expanders, for instance in [CRVW02], but note that we do not need explicitness—this is just a construction inside the lower bound proof, which is the proof that we care about.

Even more importantly, in proof complexity lower bounds for non-explicit formulas are perfectly fine anyway. To see this, recall what we discussed during the first lecture about doing proof complexity as a way to separate NP from $\mathsf{co\text{-}NP}$. Now, to do this, it would work just as fine to prove non-constructively that for any propositional proof system, there is some non-explicit formula family F_n of polynomial size such that the proofs for these formulas grow superpolynomially.

2.4 Proof Strategy for Lemma 2.7

We now present the strategy used to prove Lemma 2.7. We want to define a "measure of progress" $\mu : \{\text{clauses}\} \to \mathbb{N}$ such that given any resolution refutation $\pi = \{D_1, D_2, \cdots, D_L\}$, we can use μ to determine, for any clause in the refutation, how much progress we have made towards deriving contradiction. We want the measure μ to have the following properties:

1. $\mu(\text{axioms}) \leq 1$;

2. $\mu(\text{the final clause } \bot)$ is large;

3. μ can only increase gradually, so there must exist some $D_i \in \pi$ with medium-sized measure;

4. Such a "medium-progress" clause D_i must contain many literals, i.e., $W(D_i)$ is large, which shows that the resolution refutation π is at least as wide as this clause.

Since this holds for any resolution refutation π, we get the width lower bound that we are after.

To construct such a measure μ, let $\mathcal{H}$ denote the set of all hole axioms and let $\mathcal{P}$ denote the set of all pigeon axioms, where as before we write P^u for the the pigeon axiom associated with pigeon u. Then we define

$$\mu(D) = \min \left\{ |U'| : \mathcal{H} \wedge \bigwedge_{u \in U'} P^u \vDash D \right\} \quad . \tag{2.5}$$

Intuitively, μ measures how many pigeon axioms have we used so far to derive D. We now verify that μ satisfies the requirements above.

First of all, we have $\mu(D) \leq n$ for all $D \in \pi$, because with all the pigeon axioms as well as all the hole axioms we can derive the empty clause $\bot$, which implies everything.

For all $C \in \mathcal{H}$, we have $\mu(C) = 0$ because if $\mathcal{H}$ is satisfied, then all clauses $C \in \mathcal{H}$ have to be satisfied and no pigeon axioms at all are needed. Similarly, for all $C \in \mathcal{P}$, it holds that $\mu(C) = 1$ since picking $U = \{u\}$ for $P^u = C$ is sufficient (as an aside, observe that $\mu(C) > 0$ because when all variables x_{ij} are assigned 0, $\mathcal{H}$ is satisfied but not P^u).

For the final empty clause $\bot$ we have $\mu(\bot) > s$, since any set U' of size no more than s fits into $|N(U')| \geq |U'|$ distinct holes. This follows from the next theorem.

Theorem 2.10 (Hall's Marriage Theorem). *For $G = (U \cup V, E)$, there is a matching of U into V if and only if for all $U' \subseteq U$, we have $|N(U')| \geq |U'|$.*

This means that for all vertex sets U' of size no more than s, $\mathcal{H} \wedge \bigwedge_{u \in U'} P^u$ is satisfiable, so $\mathcal{H} \wedge \bigwedge_{u \in U'} P^u \nvDash \bot$.

Remark 2.11. The condition in Theorem 2.10 is clearly necessary. The interesting thing is that it is also sufficient. We will not prove the theorem, however, since it is a standard fact in combinatorics.

Returning to our measure μ, if $\frac{D \vee x \quad D' \vee \overline{x}}{D \vee D'}$ is an application of the resolution rule, then $\mu(D \vee D') \leq \mu(D \vee x) + \mu(D' \vee \overline{x})$. This is so because if for sets of pigeons U_1 and U_2 we have $\mathcal{H} \wedge \bigwedge_{u \in U_1} P^u \vDash D \vee x$ and $\mathcal{H} \wedge \bigwedge_{u \in U_2} P^u \vDash D' \vee \overline{x}$, then by the soundness of the resolution rule it holds for $U_1 \cup U_2$ that $\mathcal{H} \wedge \bigwedge_{u \in U_1 \cup U_2} P^u \vDash D \vee D'$. This means that in each resolution step, μ can at most double.

Hence there must exist some $D \in \pi$ such that $s/2 \leq \mu(D) < s$. Fix such a D, and then fix also a set of pigeons U' of minimal size such that the implication

$$\mathcal{H} \wedge \bigwedge_{u \in U'} P^u \vDash D \tag{2.6}$$

holds. Observe that by construction, we have $s/2 \leq |U'| = \mu(D) < s$.

Now we want to argue that all unique neighbours of U' must be represented by variables in the clause D. This gives us $W(D) \geq |\partial U'|$. Note that this is sufficient to establish Lemma 2.7, since the expansion properties of G give us that $|\partial U'| \geq e \cdot |U'| \geq es/2$, and hence $W(\pi) \geq W(D) \geq es/2$. Thus, it will be sufficient for us to prove the following claim.

Claim 2.12. For all $v \in \partial U'$ some variable $x_{u_v, v}$ occurs in D.

Proof. Fix $v^* \in \partial U'$ and a unique neighbour $u^* \in N(v^*)$. Assume that no variable x_{u', v^*} appears in D. We want to derive a contradiction.

Observe first that if we remove u^* from U', then by the definition of μ and since U' was chosen minimal it holds that

$$\mathcal{H} \wedge \bigwedge_{u \in U' \setminus \{u^*\}} P^u \nvDash D \quad .$$

But if so, then there exists an assignment α such that $\alpha \left(\mathcal{H} \wedge \bigwedge_{U' \setminus \{u^*\}} P^u \right) = 1$ but $\alpha(D) = 0$ (that is what it means that the left-hand side does not imply the right-hand side). Furthermore, without loss of generality we can assume that $\alpha(x_{u'v^*}) = 0$ for all $u' \in N(v^*)$. This is so because:

- $\mathcal{H}$ cannot be falsified by flipping variable assignments to false since all variables in $\mathcal{H}$ appear negated.
- P^u cannot not be falsified for $u \in U' \setminus \{u^*\}$, since u^* is the unique neighbor of v^* in U'.
- D cannot get satisfied, since there are no variables $x_{u'v^*}$ in D by assumption. (And from this very assumption we will now very soon derive our contradiction.)

Now let α^* be the same assignment as α except that we set $\alpha^*(x_{u^*v^*}) = 1$. We still have $\alpha^*(\mathcal{H}) = 1$ because each clause in $\mathcal{H}$ is a disjunction of exactly two variables, and we only flipped a single variable, namely $x_{u^*v^*}$. Thus all clauses in $\mathcal{H}$ will still be satisfied under α^*. Furthermore, since $\alpha^*\big(P^{u^*}\big) = 1$ we now have for the full set U' that $\alpha^*\left(\bigwedge_{u \in U'} P^u\right) = 1$. However, it still holds that $\alpha^*(D) = 0$ for the same reason as above.

But this means that $\alpha^*\left(\mathcal{H} \wedge \bigwedge_{u \in U'} P^u\right) = 1$ and $\alpha^*(D) = 0$, which contradicts the the implication in (2.6). Hence, the assumption that no variable x_{u',v^*} occurs in D must have been wrong. The claim follows. □

Putting all the pieces together as described above, Theorem 2.1 now follows.

3 Tseitin Contradictions

At the end of the lecture, we started talking about another family of CNF formulas defined in terms of graphs, the so-called *Tseitin contradictions*. They are defined as follows.

Let G be a connected undirected graph of size n. We say that a function $f : V(G) \to \{0,1\}$ has *odd weight* if $\sum f(v) \equiv 1 \pmod 2$. We introduce a variable x_e for each edge $e \in E(G)$. For each vertex $v \in V(G)$, we define a set of clauses $PARITY_v$ encoding

$$PARITY_v = \bigoplus_{e \ni v} x_e \equiv f(v) \pmod 2 \ ,$$

that is, that if we sum the truth values of all the edges incident to v then this sum is odd if $f(v) = 1$ and even otherwise. Then for a graph G and an odd-weight function f, the Tseitin contradiction $Ts(G, f)$ is defined to be the CNF formula

$$Ts(G, f) = \bigwedge_{v \in V(G)} PARITY_v .$$

How do we encode $PARITY_v$? The trick is to add clauses ruling out all incorrect assignments as explained in the next example.

Example 3.1. Say that we want to write down clauses encoding $x \oplus y \oplus z = 1$. Then there are four assignments that we need to rule out, namely all assignments with an even number of true variables. To exclude an assignment (x_0, y_0, z_0) with $x_0 \oplus y_0 \oplus z_0 = 0$, add a clause with opposite signs for all variables. Since this clause has to be satisfied, this means that some variable must take a value that disagrees with the assignment (x_0, y_0, z_0). To encode $x \oplus y \oplus z = 1$, we thus add the following clauses:

$$\begin{array}{lll} x \vee y \vee z & \text{falsified by} & x = y = z = 0, \\ \overline{x} \vee \overline{y} \vee z & \text{falsified by} & x = y = 1, z = 0, \\ \overline{x} \vee y \vee \overline{z} & \text{falsified by} & x = z = 1, y = 0, \\ x \vee \overline{y} \vee \overline{z} & \text{falsified by} & y = z = 0, x = 1. \end{array}$$

It is straightforward to verify that all valid assignments of x, y, z satisfy all the above clauses, so the above clauses encode $x \oplus y \oplus z = 1$.

Generalizing the above example, one can prove the following lemma.

Lemma 3.2. *If the maximal degree of G is d, then $Ts(G, f)$ is a d-CNF formula with at most $nd/2$ variables and at most $n \cdot 2^{d-1}$ clauses.*

We will return to the Tseitin contradictions next time and prove a lower bound for them in resolution that is in some sense even stronger than the lower bound for the pigeonhole principle formulas that we proved in this lecture.

References

[BP97] Samuel R. Buss and Toniann Pitassi. Resolution and the weak pigeonhole principle. In *Proceedings of the 11th International Workshop on Computer Science Logic (CSL '97)*, volume 1414 of *Lecture Notes in Computer Science*, pages 149–156. Springer, 1997.

[BW99] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC '99)*, pages 517–526, May 1999.

[BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version appeared in *STOC '99*.

[CRVW02] Michael Capalbo, Omer Reingold, Salil Vadhan, and Avi Wigderson. Randomness conductors and constant-degree lossless expanders. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 659–668, May 2002.

[Hak85] Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985.

[Raz01] Alexander A. Razborov. Improved resolution lower bounds for the weak pigeonhole principle. Technical Report TR01-055, Electronic Colloquium on Computational Complexity (ECCC), July 2001.

[Raz04] Ran Raz. Resolution lower bounds for the weak pigeonhole principle. *Journal of the ACM*, 51(2):115–138, 2004. Preliminary version appeared in *STOC '02*.