DD3501 Current Research in Proof Complexity Dec 19, 2011

Lecture 11

Lecturer: Jakob Nordström *Scribe: Karl Palmskog*

The goal of today's lecture is to finish the proof of the lower bound on space in polynomial calculus resolution by Filmus et al. [FLN+12] that we started doing last time (i.e., in the previous regular lecture before the guest lecture on SAT solving).

1 Recap of Terminology and Notation

Recall from last time that for a variable x we introduced the notation $x^0 \equiv x$ and $x^1 \equiv \overline{x}$ for positive and negative literals over x. This means that $x^b = 0$ in PCR (i.e., x^b is true) if and only if $x = b$.

The formulas we consider are so-called *bitwise pigeonhole principle formulas*, which are encodings of the functional pigeonhole principle where the functionality condition that every pigeon should only go into one hole does not require extra axiom clauses but is hard-coded into the encoding. For k a positive integer we let $[0, k)$ denote the set $\{0, 1, \ldots, k-1\}$.

For $n = 2^\ell$, the *bitwise PHP formula* BPHP^m_n has propositional variables $x[p, i]$ for each $p \in [0, m)$ and $i \in [0, \ell)$, with $[0, m)$ the set of pigeons and $[0, n)$ the set of holes (where we assume $m > n$). Each pigeon p is sent to the hole whose binary encoding is given by the string $x[p, \ell-1] \cdots x[p, 1]x[p, 0]$, where we say that the variables $x[p, i]$ are *associated* with the pigeon p.

The intuition behind the formula BPHP^m_n is that it specifies that no two pigeons map to the same hole. Specifically, for every two pigeons $p_1 \neq p_2 \in [0, m)$ and every hole $h \in [0, n)$ we have a *hole axiom*

$$H(p_1, p_2, h) = \bigvee_{i=0}^{\ell-1} x[p_1, i]^{1-h_i} \vee \bigvee_{i=0}^{\ell-1} x[p_2, i]^{1-h_i} \tag{1.1}$$

stating that either p_1 is not mapped to h or p_2 is not mapped to h, where $h_{\ell-1} \cdots h_0$ is the binary encoding of h.

Let α be a total assignment to the variables of BPHP^m_n, and let $S \subseteq [0, m)$ be a set of pigeons. We say that α is *well-behaved on* S if the holes assigned by α to the pigeons in S are all distinct.

A *(disjunctive) commitment* is a 2-clause of the form $x[p_1, i_1]^{b_1} \vee x[p_2, i_2]^{b_2}$, where $p_1 \neq p_2$. A *commitment set* is a set of commitments where all pigeons are distinct. The *domain* of a commitment set $\mathcal{A}$, written $\mathrm{dom}(\mathcal{A})$, is the set of pigeons mentioned in $\mathcal{A}$. The *size* of a commitment set $\mathcal{A}$, denoted $|\mathcal{A}|$, is the number of commitments in $\mathcal{A}$. An assignment α is *well-behaved on and satisfies* a commitment set $\mathcal{A}$ if α is well-behaved on $\mathrm{dom}(\mathcal{A})$ and satisfies $\mathcal{A}$.

A commitment set $\mathcal{A}$ *entails* a PCR-configuration $\mathbb{P}$ over well-behaved assignments if every assignment α which is well behaved on and satisfies $\mathcal{A}$ also satisfies $\mathbb{P}$.

2 Some Key Technical Results

Last time we proved the following lemma.

Lemma 2.1. *Let S be a set of fewer than $n/2$ pigeons; let α be an assignment well-behaved on S; and let $x[p, i]^b$ a literal associated with a pigeon $p \notin S$. Then it is possible to modify α by reassigning p so that the new assignment is well-behaved on $S \cup \{p\}$ and satisfies the literal $x[p, i]^b$.*

The proof was a simple counting argument, and from this lemma we immediately obtained the following corollary.

Corollary 2.2. *Let S and T be disjoint sets of pigeons such that $|S \cup T| \leq n/2$. Let X be a set containing exactly one literal associated with the pigeon p for each $p \in T$. Then, any assignment α that is well-behaved on S can be modified, by reassigning pigeons in T, into an assignment β that is well-behaved on $S \cup T$ and also satisfies all literals in X.*

We also claimed, but did not prove, the next lemma.

Lemma 2.3 (Locality lemma). *Let $\mathcal{A}$ be a commitment set and $\mathbb{P}$ a PCR-configuration such that $\mathcal{A}$ entails $\mathbb{P}$ over well-behaved assignments and $|\mathcal{A}| \leq n/4$. Then there exists a commitment set $\mathcal{B}$ of size $|\mathcal{B}| \leq 2 \cdot Sp(\mathbb{P})$ such that $\mathcal{B}$ entails $\mathbb{P}$ over well-behaved assignments.*

Let us first see how the lower bound on PCR space that we discussed last time now follows, and then prove the key technical result in Lemma 2.3 needed for the proof of the space lower bound.

3 Proof of the PCR Space Lower Bound

In this section, we establish the following theorem, which is the lower bound we set out to prove last time.

Theorem 3.1. $Sp_{\mathcal{PCR}}(BPHP^m_n \vdash \perp) > n/8$.

Proof. Let $\pi = \{\mathbb{P}_0 = \emptyset, \ldots, \mathbb{P}_\tau\}$ be a PCR-derivation from $BPHP^m_n$ in monomial space at most $n/8$. We will prove that for all $\mathbb{P}_t \in \pi$, there exists a commitment set $\mathcal{A}_t$ such that $|\mathcal{A}_t| \leq 2 \cdot Sp(\mathbb{P}_t) \leq n/4$ and $\mathcal{A}_t$ entails $\mathbb{P}_t$.

If we can do this, then the theorem follows. To see this, note first that by Corollary 2.2 (with $S = \emptyset$), $\mathcal{A}_t$ is satisfiable by some well-behaved assignment α (in fact, we can satisfy *all literals* in $\mathcal{A}_t$). This in turn implies that α must satisfy $\mathbb{P}_t$ as well, which in particular means that $\mathbb{P}_t$ does not contain the constant polynomial 1 encoding contradiction. It follows that no PCR-derivation in space at most $n/8$ can refute $BPHP^m_n$.

We construct the commitment sets $\mathcal{A}_t$ by structural induction on the derivation π. For the base case, we have $\mathbb{P}_0 = \emptyset$ and can set $\mathcal{A}_0 = \emptyset$. Suppose inductively that we have defined $\mathcal{A}_t$ and want to construct $\mathcal{A}_{t+1}$. There are three cases depending on which derivation step is performed at time $t+1$: axiom download, inference, or erasure. We analyze the different cases in this order.

Axiom download We distinguish two download cases: (a) complementarity axioms of the form $x + \overline{x} - 1$ or boolean axioms of the form $x^2 - x$ and (b) hole axioms $H(p_1, p_2, h)$. In the former case, we can simply set $\mathcal{A}_{t+1} = \mathcal{A}_t$ since any truth value assignment satisfies such an axiom by definition. In the latter case, we distinguish the following subcases (by symmetry).

1. $\{p_1, p_2\} \subseteq \mathrm{dom}(\mathcal{A}_t)$: Note that any well-behaved assignment sends p_1 and p_2 to different holes and thus satisfies $H(p_1, p_2, h)$. Hence, $\mathcal{A}_t$ already entails $\mathbb{P}_{t+1}$ over well-behaved assignments, and we can set $\mathcal{A}_{t+1} = \mathcal{A}_t$.

2. $\{p_1, p_2\} \cap \mathrm{dom}(\mathcal{A}_t) = \emptyset$: We add a commitment $C = x[p_1, 0]^{1-h_0} \vee x[p_2, 0]^{1-h_0}$, say (any literal for p_1 and any literal for p_2 is fine), to $\mathcal{A}_t$, i.e. we set $\mathcal{A}_{t+1} = \mathcal{A}_t \cup \{C\}$. Then since p_1 and p_2 are now both in the domain of $\mathcal{A}_{t+1}$, any well-behaved assignments will send these two pigeons to distinct holes and in particular satisfy $H(p_1, p_2, h)$, and since $\mathcal{A}_t \subseteq \mathcal{A}_{t+1}$ everything else in $\mathbb{P}_{t+1}$ is also entailed by $\mathcal{A}_{t+1}$.

3. $p_1 \in \mathrm{dom}(\mathcal{A}_t), p_2 \notin \mathrm{dom}(\mathcal{A}_t)$: Fix $p^* \notin \mathrm{dom}(\mathcal{A}_t) \cup \{p_2\}$ and let $C = x[p_2, 0]^{1-h_0} \vee x[p^*, 0]^0$, say. (This is just to get p_2 added to the domain of $\mathcal{A}_{t+1} = \mathcal{A}_t \cup \{C\}$ in a way that is consistent with the formal requirements on how commitment sets may look). Again, any well-behaved assignment will now have to send p_1 and p_2 to distinct holes and so satisfy $H(p_1, p_2, h)$.

In an axiom download step, the space increases by 1 and we never add more than $1 < 2$ commitments in any of the cases above, so we clearly maintain the invariant $|\mathcal{A}_{t+1}| \leq 2 \cdot Sp(\mathbb{P}_{t+1})$.

Inference Let $\mathbb{P}_{t+1} = \mathbb{P}_t \cup \{Q\}$, where the polynomial Q is derived from $\mathbb{P}$. Then since PCR is sound, Q certainly follows semantically from $\mathbb{P}_t$. That is, if it holds for all $P \in \mathbb{P}_t$ that $P(\alpha) = 0$, then it also holds that $Q(\alpha) = 0$. Set $\mathcal{A}_{t+1} = \mathcal{A}_t$. Then all well-behaved α satisfying $\mathcal{A}_{t+1} = \mathcal{A}_t$ must satisfy $\mathbb{P}_t$ by the induction hypothesis and hence also Q, so all of $\mathbb{P}_{t+1}$ is satisfied.

Here, space is increasing, but the space of the commitments is unchanged, so we are still on track.

Erasure Set $\mathbb{P}_{t+1} = \mathbb{P}_t \setminus \{Q\}$ for some $Q \in \mathbb{P}_t$. As we discussed last time, this is the hard case. We know that $\mathcal{A}_t$ entails $\mathbb{P}_{t+1} \subseteq \mathbb{P}_t$, but the size of $\mathcal{A}_t$ may now be far too large if Q contained a lot of monomials. Thus we need to find a smaller commitment set that still entails $\mathbb{P}_{t+1}$, but it is not at all clear how to do this.

However, here the "magic" Lemma 2.3 comes to the rescue. We just appeal to this lemma with $\mathcal{A} = \mathcal{A}_t$ and $\mathbb{P} = \mathbb{P}_{t+1}$ and let $\mathcal{A}_{t+1}$ be the $\mathcal{B}$ produced by the lemma. By the statement of the lemma, $\mathcal{A}_{t+1}$ has all the properties needed for the induction step to go through.

This takes care of all inductive cases, and the theorem follows. □

4 Proof of the Locality Lemma

Looking at the proof of Theorem 3.1 above, it seems that we never really did any hard work. The only case that looked tricky was the erasure steps, and there we just appealed to Lemma 2.3 and were immediately done. Thus, if there is any work at all needed to be done in proving Theorem 3.1, we should expect this to be in the proof of Lemma 2.3. We now finally prove this "magic lemma."

Proof of Lemma 2.3. Consider a bipartite graph with on the left-hand side the set M of all distinct monomials in $\mathbb{P}$ (which is certainly a lower bound on the monomial space), and on the right-hand side the set of all disjunctive commitments in $\mathcal{A}$. We let there be an edge between a monomial $m \in M$ on the left and a commitment $C \in \mathcal{A}$ on the right if there is a pigeon p mentioned in both (i.e., there is some variable $x[p,i]^b \in \mathit{Vars}(m)$ and some literal $x[p,i']^{b'} \in \mathit{Lit}(C)$, where i and i' might be the same or distinct, and similarly for b and b'). To follow the rest of the argument, it might be helpful to consider the example graph drawn in Figure 1(a).

Let $\Gamma \subseteq M$ be a set of maximal size such that $|N(\Gamma)| \leq 2 \cdot |\Gamma|$. Note that Γ is not necessarily unique, but such a maximal set always exists. If nothing else, $\Gamma = \emptyset$ satisfies the requirement that $|N(\Gamma)| \leq 2 \cdot |\Gamma|$. If we would have $\Gamma = M$, then it is not too hard to prove directly that we can conclude with $\mathcal{B} = N(\Gamma)$ (although this is not quite immediate but requires a short argument), so in what follows it is helpful to think of the case when $\Gamma \neq M$ (although the proof works for any Γ such that $\emptyset \subseteq \Gamma \subseteq M$).

If $\Gamma \neq M$, it must hold for all $S \subseteq M \setminus \Gamma$ that $|N(S) \setminus N(\Gamma)| > 2 \cdot |S|$, since otherwise we could add S to Γ to get a larger set. But this means that there is a matching of every $m \in M \setminus \Gamma$ to two distinct commitments $C', C'' \in \mathcal{A} \setminus N(\Gamma)$ such that no two m, m' share any commitment. To see this, just make two copies of each monomial/vertex in $m \in M \setminus \Gamma$ with the same edges from both copies to the vertices on the right, apply Hall's theorem, and then identify the two copies of the monomial again (this step is illustrated in Figure 1(b), where Γ and $N(\Gamma)$ are in the upper half of the graph).

Fix such a monomial $m \in M \setminus \Gamma$ and suppose it has been matched to the two disjunctive commitments $C' = x[p',i']^{b'} \vee x[q',j']^{c'}$ and $C'' = x[p'',i'']^{b''} \vee x[q'',j'']^{c''}$ as shown in Figure 1(c). By construction, m must mention at least one pigeon each from C' and C'', so suppose without

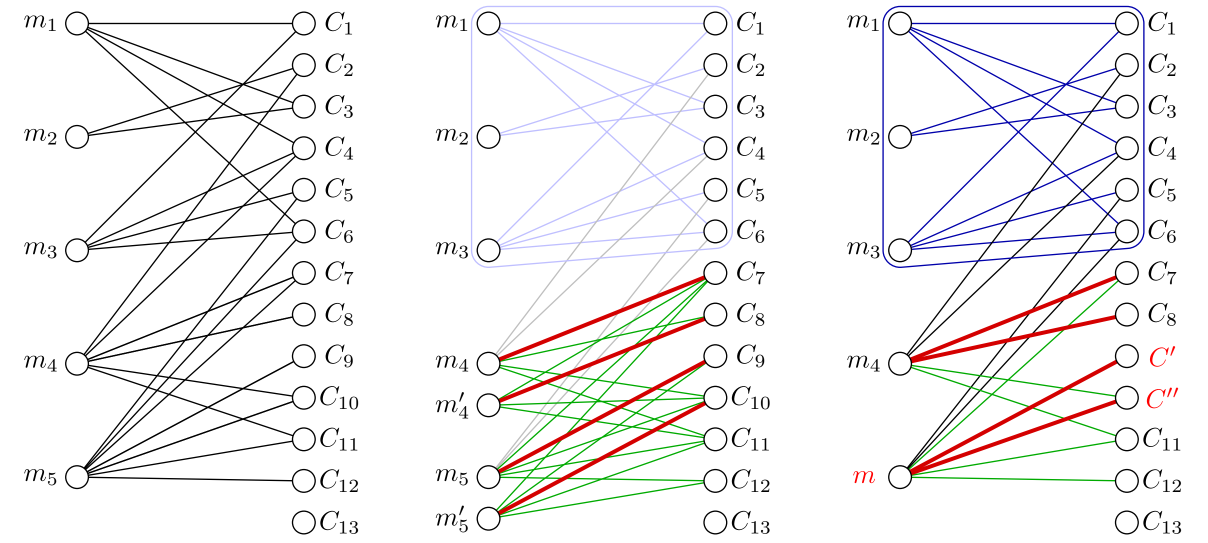

(a) Bipartite graph with monomials $m_i \in \mathbb{P}$ and commitments $C_j \in \mathcal{A}$.

(b) Hall's theorem applied on $M \setminus \Gamma$ (with monomial copies) and $\mathcal{A} \setminus N(\Gamma)$.

(c) Graphs with Γ and matchings constructed at end of argument.

Figure 1: Illustration of argument in proof of the Locality lemma.

loss of generality that we can pick p' and p'' to be these pigeons. (It can be the case that m also mentions q' or q'' or both, but by construction we are guaranteed that m mentions at least one pigeon in each commitment and this is all we will need here.) Thus, there are some literals $x[p', i_1]^{b_1}$ and $x[p'', i_2]^{b_2}$ such that $m = x[p', i_1]^{b_1} \cdot x[p'', i_2]^{b_2} \cdot m'$. We choose a new commitment C_m for m to be $C_m = x[p', i_1]^{b_1} \vee x[p'', i_2]^{b_2}$. We construct commitments in this way for every $m \in M \setminus N(\Gamma)$, and let our new commitment set be $\mathcal{B} = N(\Gamma) \cup \{C_m \mid m \in M \setminus N(\Gamma)\}$, i.e., the union of all these new assignments and the old commitments from $\mathcal{A}$ in $N(\Gamma)$. We claim that this is the commitment set we are looking for.

Firstly, we can verify that $\mathcal{B}$ is indeed a commitment set. This is so since all pigeons mentioned in $\mathcal{A}$ are different, and the pigeons in $\mathcal{B}$ are just a subset of the pigeons in $\mathcal{A}$. Secondly, with regard to size it clearly holds that $|\mathcal{B}| \leq 2 \cdot |N(\Gamma)| + |M \setminus N(\Gamma)| \leq 2 \cdot |M| \leq 2 \cdot Sp(\mathbb{P})$ (taking a look at Figure 1(c) might be helpful in verifying this). However, we also need to show that $\mathcal{B}$ entails $\mathbb{P}$ over well-behaved assignments. That is, we must prove that every β that is well-behaved on and satisfies $\mathcal{B}$ also satisfies $\mathbb{P}$. Note that this is a priori not clear. We know that this holds for $\mathcal{A}$ by assumption, but $\mathrm{dom}(\mathcal{A})$ is potentially much larger than $\mathrm{dom}(\mathcal{B})$ and so $\mathcal{A}$ only has to deal with much more well-behaved assignments. Also, and more seriously, the commitments in $\mathcal{B}$ are not a subset of those in $\mathcal{A}$, and on the contrary might be in conflict with $\mathcal{A}$ in the sense that satisfying literals in $\mathcal{B}$ falsifies literals in $\mathcal{A}$.

We prove that $\mathcal{B}$ entails $\mathbb{P}$ over well-behaved assignments in a slightly roundabout way by finding, given any assignment β well-behaved on and satisfying $\mathcal{B}$, another assignment α such that

1. $\mathbb{P}(\alpha) = \mathbb{P}(\beta)$, and

2. α is well-behaved on and satisfies $\mathcal{A}$.

By item 2 it follows from the inductive hypothesis that α satisfies $\mathbb{P}$. But if so, then β also satisfies $\mathbb{P}$ by item 1, which is what we want to prove.

To this end, let S be the set of pigeons in $\mathrm{dom}(\mathcal{B})$, and let T be the set of pigeons in $\mathrm{dom}(\mathcal{A}) \setminus \mathrm{dom}(\mathcal{B})$ (notice that $\mathrm{dom}(\mathcal{B}) \subseteq \mathrm{dom}(\mathcal{A})$). Let X be the set of literals that for each $p \in T$ includes the (unique) literal $x[p, i]^b$ associated with p and appearing in $\mathcal{A}$. Note that each commitment in $\mathcal{A} \setminus N(\Gamma)$ will have at least one literal in X (some commitments will potentially have both literals in X). Since $|\mathcal{A}| \leq n/4$, we have $|S \cup T| \leq n/2$. Apply Corollary 2.2 to S, T, and β to get a truth value assignment that is well-behaved on $S \cup T$, agrees with β on pigeons outside T, and satisfies X. We claim that this is the assignment α that we need.

To see this, note first that no monomial in Γ mentions pigeons in T (by construction), so α and β agree on on monomials in Γ. For $m \in M \setminus \Gamma$, all β satisfying $\mathcal{B}$ must set the monomial m to zero, since this is how the new commitments were constructed. Reassigning pigeons in T can change variables in m, but there is still at least one variable that is set to zero, zeroing the whole monomial. So for all $m \in M$, α gives the same value to m as does β, namely 0. Hence α and β agree on all monomials in M and $\mathbb{P}(\alpha) = \mathbb{P}(\beta)$. This takes care of item 1 above.

By Corollary 2.2, α is well-behaved on $S \cup T = \mathrm{dom}(\mathcal{A})$. Also, since α satisfies X as well as $N(\Gamma)$, α satisfies $\mathcal{A}$. This takes care of item 2, and as already discussed it now follows that $\mathbb{P}(\alpha) = 0$.

Thus, every β that is well-behaved on and satisfies $\mathcal{B}$ must also satisfy $\mathbb{P}$. The lemma follows. □

5 Summary of the Lectures on Polynomial Calculus and PCR

Let us conclude by making a quick summary of what we have covered for polynomial calculus and polynomial calculus resolution, just as we did for resolution in lecture 6.

For PC and PCR, we have studied the proof complexity measures size, space, and degree, where we observed that in some sense degree seemed to play a role analogous to that of width in resolution, and we have proven lower bounds for size [AR03] and space [FLN$^+$12] in PCR (and hence also PC). Let us summarize what we know about the relations between these measures. As we will see, in contrast to what was the case for resolution, here we do not know much.

5.1 Size Versus Degree

The relationship between size and degree is analogous to the relationship between length and width in resolution, that is, small degree implies small size [CEI96] and small size implies (reasonably) small degree [IPS99]. Currently, nothing is known about tradeoffs between size and degree. Thus, here we have the same situation as for length versus width in resolution.

5.2 Space Versus Degree

Since we said that degree is an analogoue of width, and since we know a lot about the relationship between space and width in resolution,[1] it is natural to ask whether there are nontrivial relations between space and degree in PC or PCR. For instance, is there some kind of connection along the lines of [AD08]? The answer is that we do not know. As far as the lecturer is aware, nothing is currently known about these questions.

5.3 Size Versus Space

Does small space imply small size as in resolution? Does small size imply anything about space? Are there any tradeoffs? Again, almost nothing is known. However, there is a very recent result by Huynh and Nordström [HN12] that can be interpreted as giving some indications as to what kind of relations we can expect to hold here.[2]

[1] Although we did not have time to cover all the results so far, but will return to the remaining ones after the Christmas break.

[2] Regarding this result, it is absolutely clear that we will *not* have the time to cover it in any detail—unfortunately, that is beyond the scope of this course—but if there is interest we could perhaps give an overview of roughly how the result goes.

References

[AD08] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, May 2008. Preliminary version appeared in *CCC '03*.

[AR03] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003. Available at http://people.cs.uchicago.edu/~razborov/files/misha.pdf. Preliminary version appeared in *FOCS '01*.

[CEI96] Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 174–183, May 1996.

[FLN$^+$12] Yuval Filmus, Massimo Lauria, Jakob Nordström, Neil Thapen, and Noga Zewi. Space complexity in polynomial calculus. In *Proceedings of the 27th Annual IEEE Conference on Computational Complexity (CCC '12)*, June 2012. To appear.

[HN12] Trinh Huynh and Jakob Nordström. On the virtue of succinct proofs: Amplifying communication complexity hardness to time-space trade-offs in proof complexity. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC '12)*, May 2012. To appear.

[IPS99] Russell Impagliazzo, Pavel Pudlák, and Jiri Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Computational Complexity*, 8(2):127–144, 1999.