DD3501 Current Research in Proof Complexity Dec 5, 2011

Lecture 8

Lecturer: Jakob Nordström *Scribe: Emma Enström, Jakob Nordström*

The purpose of today's lecture is to do as much as we can of the lower bound on PCR proof size from Alekhnovich and Razborov [AR03], which we started on last time. Let us begin by recalling what was said then.

1 Recapitulation of Material from Last Lecture

For a CNF formula F, we can define a bipartite graph $G(F)$ with the clauses of F as the left vertex set and the variables of F as the right vertex set, and edges between variable vertices and clause vertices if the corresponding variable occurs in the corresponding clause. Note that this is not a bijection between formulas and bipartite graphs—we are throwing away all information about the signs of the literals in the clauses, for instance—but nevertheless the graph $G(F)$ tells us something about the structure of the formula F.

We say that the *unique neighbours* of a subset S of vertices on one side are the vertices on the other side which have exactly one incoming edge from that subset (i.e., they are the neighbour of a unique vertex in S). The unique neighbours are also called the *boundary* of the set of neighbours of S. Informally, we say that a graph G is a *unique-neighbour expander*, or a *boundary expander*, if all vertex sets on the left hand side of at most moderate size have many unique neighbours on the right hand side. More formally, a bipartite graph $G = (U \cup V, E)$ is a (d, s, e)-boundary expander if all vertices on the left have at most d incident edges and all left vertex sets $U' \subseteq U$ of size at most s have at least $e \cdot \left|U'\right|$ unique neighbours on the right. It is a well-known fact that a random constant-degree bipartite graph, i.e., a graph where for each vertex on the left we pick a fixed-size subset of neighbours on the right uniformly and independently at random, is a good expander almost surely provided that we choose the right parameters.[1]

We intend today to prove a PCR proof size lower bound for randomly generated formulas. Formally, we say that F is a *random k-CNF formula*, denoted $F \sim \mathcal{F}_k^{n,\Delta}$, if F consists of Δn clauses chosen uniformly and independently and random, and with replacement, from all $\binom{n}{k} \cdot 2^k$ possible k-clauses over n variables. If the *density* Δ is large enough, then F is almost surely unsatisfiable. What density is needed for this to hold depends on k, but once k is fixed we can also fix Δ to be a constant. And this constant need not be particularly large—for instance, for $k = 3$ we can pick $\Delta = 4.6$.

If we generate a random formula F in this way and then look at the graph $G(F)$, it is not hard to see that this is just a complicated way of generating a random bipartite graph. Hence $G(F)$ will almost surely be an excellent expander (again for the right choice of parameters, but we will not go into too much detail here). Alekhnovich and Razborov get their lower bounds for random k-CNF formulas by focusing on the graphs $G(F)$ and establishing that if such a graph has sufficiently good expansion, then the degree of any refutation of F has to be large.[2] And as we proved last time, a sufficiently strong lower bound on degree also implies a strong lower bound on proof size both for PC and PCR.

[1] For a sequence of events X_n parameterized by n in some natural way, for instance the event that a random k-CNF formula over n variables is unsatisfiable, we say that X_n happens *almost surely* (sometimes abbreviated "a.s.") if the probability of X_n approaches 1 as n approaches infinity.

[2] The attentive reader might ask what happens if F is not unsatisfiable in the first place. Note that for any formula F, we can let F^+ be the same formula with all literals flipped to be positive. The latter formula is clearly satisfiable, but $G(F^+) = G(F)$. However, the claim still holds for a vacuous reason—if the formula is satisfiable, then any refutation must have very high degree indeed (since there is no refutation at all).

Therefore, to get the lower bound that we are after it is sufficient to prove the following theorem (the statement of which we recall from last lecture).

Theorem 1.1 (Theorem 3.13 in [AR03]). *If $G(F)$ is a (d, s, e)-boundary expander, then*

$$Deg_{\mathcal{PC}}(F \vdash \bot) \geq \frac{se}{2}$$

over any field $\mathbb{F}$.

We remark that when applying this theorem, ideally we would want s to be linear in the formula size n and e to be a constant. Such a linear lower bound on degree would give an exponential lower bound on size by the theorem in [IPS99] that we proved last time.

2 Some Notation Used in This Lecture

We already defined some notation towards the end of last lecture, but repeat it again here to have it fresh in memory.

For a CNF formula $F = C_1 \wedge C_2 \wedge \cdots \wedge C_m$ over variables $x_1, x_2, \ldots, x_n$, we let $f_1, f_2, \ldots, f_m$ denote the corresponding polynomials in the PC-translation of the formula (i.e., the formulas are over variables x_i only and we do not have any "negative variables" $\overline{x}_i$).

We write T_n to denote the set of all multilinear monomials (or *terms*, as we will perhaps more frequently refer to them today)[3] over the n variables $x_1, x_2, \ldots, x_n$. We let $T_{n,d} = \{t \in T_n \mid Deg(t) \leq d\}$ denote all multilinear monomials of degree at most d.

We write $S_n(\mathbb{F})$ to denote the set of all multilinear polynomials. For simplicity, we will assume today that all Boolean axioms $x_i^2 - x_i$ for all i are applied implicitly so that we always have multilinear polynomials.[4] In particular, from now on "monomial" and "polynomial" will always mean "multilinear monomial" and "multilinear polynomial" unless explicitly stated otherwise. Formally, $S_n(\mathbb{F})$ is an $\mathbb{F}$-*algebra*, i.e., a vector space with multiplication of vectors, but we will not be interested in getting too formal today. We let $S_{n,d}(\mathbb{F})$ denote the set of all multilinear polynomials of degree at most d. This is a vector space but *not* an algebra since it is not closed under multiplication (since the degree can get too large). Just to make sure we are on the same page, an example of a monomial is $x_1x_2x_3x_5$ and an example polynomial is $5x_1x_2x_3x_5 + 2x_2x_6$.

We next state some definitions that play an important role in [AR03].

Definition 2.1 (Admissible ordering). An *admissible ordering* $\preceq$ of all monomials is any (total) ordering such that

1. If $Deg(t_1) < Deg(t_2)$ then $t_1 \preceq t_2$.
2. If $t_1 \preceq t_2$ and $t \in T_n$ does not mention variables in t_1 or t_2, then $tt_1 \preceq tt_2$.

One example of an admissible ordering is to order the monomials first with respect to degree and then lexicographically, so that, for instance, $x_1 \prec x_2 \prec x_3 \prec x_4$, $x_1x_2x_4 \prec x_1x_3x_4$, and $x_1x_3x_4 \prec x_1x_2x_3x_4$.

Definition 2.2 (Leading term). For any $f \in S_n(\mathbb{F})$, let $LT(f)$ denote the *leading term* of f with respect to $\preceq$ (i.e., the largest monomial with respect to this ordering).

Notice that by property 1 in Definition 2.1 it must hold that the leading term of f is a term of highest degree in f, i.e., $Deg(LT(f)) = Deg(f)$.

Recall that an *ideal* is a subring S that is closed under multiplication by any element of the ambient ring (the ring of which S is a subring). Following the notation in [AR03], we will write $Span(f_1, \ldots, f_m) \subseteq S_n(\mathbb{F})$ to denote the ideal generated by $f_1, \ldots, f_m$,[5] i.e., by the polynomials

[3] Sometimes when one wants to be really precise, a *monomial* is a product of variables while a *term* is a monomial multiplied by a coefficient, but we will not need to make this distinction.

[4] This is without loss of generality—if we wanted, we could have defined the multiplication rule in PC to always yield a multilinear result, and in fact some papers use this definition.

[5] This is not completely standard, but we are just trying to be consistent with the notation in [AR03] here.

corresponding to the clauses in F. The reason we care about ideals is that a PC derivation from a formula F is no more and no less than a computation generating elements in the ideal spanned by the clauses of F. To prove lower bounds on PC refutations, we therefore need to understand the structure of this ideal. To this end, we next give a quick recap of, and elaborate on, some standard algebra.

3 Some Standard Algebra

For simplicity, let us start with a concrete example, choosing the integers $\mathbb{Z}$ as our ring. Consider the subring $I_k = \{kn \mid n \in \mathbb{Z}\}$ for some constant k, or the subring $I_5 = \{5n \mid n \in \mathbb{Z}\}$ to be even more concrete. This subring is indeed a ring—the result of any multiplication or addition of elements of the set is easily seen to be contained in the set—but it is not only a subring but even an ideal since it is closed under multiplication by *any* integer of the ambient ring $\mathbb{Z}$. That is (and to be overly formal), for any element $h \in I_5$ and any any $a \in \mathbb{Z}$ it holds that $ah \in I_5$, since h already had a factor 5 in it. We can write any $m \in \mathbb{Z}$ uniquely as $m = q + r$ where

- q is in I_5, and
- r is the minimal irreducible term in $\Psi = \{0, 1, 2, 3, 4\}$.

This is just a more complicated way of describing what we all know as modular arithmetic, and what is usually written as $r \equiv m \bmod 5$. For reasons that will (hopefully) become clear shortly, let us continue to be overly formal and give the next definition.

Definition 3.1 (Reduction operator). For any ideal $I_k = \{kn | n \in \mathbb{Z}\}$, we say that the *reduction operator* $R_{I_k}(m)$ maps m to $R_{I_k}(m) = r$ such that $r \equiv m \bmod k$ and $r \in [0, k-1]$.

Now let us generalize this a bit. Fix a field $\mathbb{F}$ and let $V \subseteq S_n(\mathbb{F})$ be an ideal in the multilinear polynomial ring over $\mathbb{F}$. We say that a term $t \in T_n$ is *reducible modulo* V if there exists a polynomial f in V such that $LT(f) = t$, that is, if t is the leading term of some polynomial in the ideal. Otherwise, t is *irreducible*. We let Ψ denote the set of all irreducible monomials. It is a standard fact in algebra that any polynomial $p \in S_n(\mathbb{F})$ can be written uniquely as $p = q + r$ where

- q is in V,
- r is in the linear subspace over Ψ,

and furthermore V and Ψ are linearly independent. Therefore, just as in Definition 3.1 above we can define a reduction operator R_V that maps the term t to the unique polynomial $R_V(t)$ such that $t - R_V(t)$ is in V.

Actually, to make these notes self-contained let us give a proof of this standard fact. Let us write p as $p = q + r$ and $p = q' + r'$ for $q, q' \in V$ and r, r' in the linear subspace over Ψ. Then $q - q' = r' - r$. Clearly, $q - q' \in V$. If $r \neq r'$, then $LT(r' - r)$ is irreducible by assumption, but on the other hand we have $LT(r' - r) = LT(q - q')$ which shows that it is reducible. Hence, we must have $r = r'$. Linear independence follows in the same way by observing that if there are polynomials q and r such that $q + r = 0$, then their leading terms cancel.

4 Ideals, Pseudoideals, and a Degree Lower Bound Approach

In order to prove lower bounds on refutation degree, and hence on refutation size, we show the equivalent statement that no PC derivation of too low degree can derive contradiction.

In what follows, we will write $V_{n,d}(f_1, \ldots, f_m)$, or more briefly just $V_{n,d}$, to denote the set of all polynomials in $S_{n,d}(\mathbb{F})$ derivable in degree at most d. Note that it is sufficient for us to

prove that $1 \notin V_{n,d}$. However, $V_{n,d}$ is *not* an ideal, since it is not closed under multiplication because of the degree constraints on this set. But it turns out that we can sort of pretend that $V_{n,d}$ is an ideal and use the algebraic concepts in Section 3 anyway. This idea of studying such *pseudoideals* $V_{n,d}$ was first proposed in [CEI96]. So let us describe how to do this.

The addition rule in PC means that $V_{n,d}$ is a vector space. We can also extend the concept of reducibility in the natural way as follows.

Definition 4.1 (Reducible and irreducible terms). The term $t \in T_{n,d}$ is *reducible modulo* $V_{n,d}$ if there exists an $f \in V_{n,d}$ such that $t = LT(f)$, and *irreducible* otherwise.

Let us denote the set of all irreducible terms in the sense of Definition 4.1 by $\Psi_{n,d}$. Then we can write any $p \in S_{n,d}(\mathbb{F})$ uniquely as $p = q + r$ where

- q is in $V_{n,d}$ and
- r is in the linear space over $\Psi_{n,d}$.

(More formally, the reason for this is that $S_{n,d}(\mathbb{F})$ is the direct sum of $\Psi_{n,d}$ and $V_{n,d}$, but we do not want to go too deep into this or we will get lost.) Hence, we can let $R_{n,d}$ be the reduction operator that maps $p = q + r$ to r as described above, and this will be analogous to the other reduction operators that we have already seen.

Given these definitions, we can now describe what we want to prove in three equivalent ways (for some degree bound d):

1. Any PC refutation (of the formula F with graph $G(F)$) must have degree strictly larger than d.
2. The constant term 1 (which is the encoding of contradiction in PC) is not contained in $V_{n,d}$.
3. The reduction operator $R_{n,d}$ is nontrivial, i.e., $R_{n,d} \not\equiv 0$ (because if $1 \in V_{n,d}$, then it holds that $V_{n,d} = S_{n,d}$ and there are no irreducible terms).

If we recall that the *kernel* of an operator R is the set of all elements $\ker(R) = \{v \mid R(v) = 0\}$ sent to zero by R, we can rewrite condition 3 as $1 \notin \ker(R_{n,d})$, or equivalently $\ker(R_{n,d}) \neq S_{n,d}(\mathbb{F})$. Explaining in natural language, the reason for this equivalence is that if 1 is in the ideal, then everything is in the ideal, which means that we can always write $p = q + r$ for $r = 0$ and all terms are reducible.

Thus, to prove a lower bound on PC refutation degree, all we need to do is to study the linear operator $R_{n,d}$ and establish that 1 is not in its kernel. The problem is, though, that we have not really done any real work so far—we have just restated our original problem in more fancy algebraic terms. And it is very unclear how this will help us understand $R_{n,d}$.

Nevertheless, by rephrasing the original problem in this way we can finally make some progress by using the following observation: Suppose that we could somehow define a "stronger" operator R, which had the property that $\ker(R_{n,d}) \subseteq \ker(R)$ and such that R was easier to understand than $R_{n,d}$, while at the same time R was not too strong in the sense that it still held that $1 \notin \ker(R)$. If we found such an R, then clearly we would be done (by condition 3). And this is exactly what we will do.

Let us describe in more detail what kind of R we are looking for, using a lemma from [Raz98] (restated as Lemma 2.10 in [AR03]) that formalizes this approach.

Lemma 4.2 ([Raz98]). *Suppose that $f_1, \ldots, f_m$ are multilinear polynomials over $x_1, \ldots, x_n$ of degree at most $d < n$. If there exists a linear operator R on $S_{n,d}(\mathbb{F})$ such that*

1. *$R \not\equiv 0$,*

2. *$R(f_i) = 0$ for all i,*

3. *for all terms t of degree $Deg(t) < d$ and all variables x it holds that $R(xt) = R(x \cdot R(t))$,*

then there is no PC refutation (or PCR refutation) of $f_1, \ldots, f_m$ in degree d or less.

Proof. In order to prove the lemma, we first show that for any PC derivation in degree at most d the reduction operator R sends any polynomial derived to 0. To see this, we argue by induction:

- The claim is true for axioms (by the assumption about R in property 2).
- For addition steps, the claim holds simply by linearity.
- For multiplication steps, if the polynomial $p = \sum_i t_i$ is multiplied by x we get by linearity and property 3 that

$$R(xp) = R(\textstyle\sum_i xt_i) = \sum_i R(xt_i) = \sum_i R(xR(t_i)) = \\ = R(\textstyle\sum_i xR(t_i)) = R(xR(p)) = R(x \cdot 0) = 0 \quad (4.1)$$

  since $R(p) = 0$ by the inductive hypothesis.

Now this means that if there were a PC derivation of 1 in degree at most d, then we would have $R(1) = 0$. But this in turn would mean that R would be sending everything in $S_{n,d}(\mathbb{F})$ to 0 (by property 3, since $t = t \cdot 1$), which means that R is trivial contradicting the assumption in property 1. The lemma follows. □

5 Constructing a Reduction Operator

Now we know that if we can find a reduction operator R meeting the criteria of Lemma 4.2, this would suffice to prove Theorem 1.1. So let us build such an R. In doing so, we will in fact move back from pseudoideals to ideals again, and to reduction operators over ideals, since they are nicer to work with. It turns out we can do this if we choose our ideals carefully.

Before starting to define this stronger R, it is important to note that by linearity, it is sufficient to define R on isolated *monomials* rather than on a full polynomial in one go. This of course holds in general: in order to define a linear operator on a vector space, it is sufficient to define how it acts on the elements in a basis since it is then uniquely defined on the whole space by linearity. And in this case the set of monomials of degree at most d clearly form a basis of the vector space $S_{n,d}$.

Informally speaking, for t any term (that is, monomial) we will define the result $R(t)$ of reducing t in three steps by:

- finding set of axioms $\{f_{i_1}, \ldots, f_{i_l}\}$ on which t "depends" (we will elaborate later on what this means),
- looking at the (ordinary) ideal $V(t) = Span(f_{i_1}, \ldots, f_{i_l})$ generated by these axioms, and
- letting R send t to the polynomial that is the "remainder" of t modulo the ideal $V(t)$ as described above, i.e., $R(t) = R_{V(t)}(t)$.

Our hope is that we can do this without the ideals $V(t)$ getting too large, so that in particular $R(1) \neq 0$.

Let us now (admittedly somewhat out of the blue) describe how to construct such a reduction operator R. Recall that the clauses of our CNF formula F are encoded as polynomials $f_1, \ldots, f_m$. We build the bipartite graph $G(F)$ as explained at the beginning of this lecture, and let $N(f_i)$ denote the neighbours on the right-hand side in $G(F)$ of the polynomial f_i encoding

the clause C_i. By a little abuse of notation, for any term t we will write $N(t)$ to denote the vertices on the right-hand side in $G(F)$ corresponding to the variables in t. We can think of this as temporarily adding a "bogus vertex" t at the bottom of the left-hand side of $G(F)$ and drawing edges from this term t to the variables on the right appearing in it. We will look at the graph $G(F)$ with t added on the left in this way to determine which subsets of the axioms/polynomials $f_1, \ldots, f_m$ should be chosen to generate the ideal $V(t)$ that we will use to construct $R(t)$.

In order to give the formal definition, we introduce some new notational conventions that will be used in the rest of these notes. We will let I (with or without subindex) denote any subset of the natural numbers $[m] = \{1, \ldots, m\}$. We will associate such subsets of integers I with the subsets of axioms $\{f_i \mid i \in I\}$ indexed by them. In fact, in what follows we will allow ourselves to be a bit sloppy and sometimes *identify* I with $\{f_i \mid i \in I\}$ when no misunderstanding can occur. Hence, in particular, ∂I is a shorthand for the boundary vertices $\partial\{f_i \mid i \in I\}$ on the right-hand side of $G(F)$.

Definition 5.1 (Support [AR03]). For t a term, we say it is possible to *t-infer* I' from I if

1. $|I'| \leq s/2$, and
2. $\partial I' \subseteq N(I) \cup N(t)$, i.e., $\partial\{f_i \mid i \in I'\} \subseteq N(\{f_i \mid i \in I\}) \cup N(t)$.

We define the *support* of the term t, denoted $Sup(t)$, to be the largest subset of $[m]$ that can be t-inferred by repeated applications of the derivation rule above starting from $I_0 = \emptyset$.

That is, the support of t is derived by some sequence

$$\begin{aligned} \emptyset &\vdash_t I_1 \\ I_1 &\vdash_t I_2 \\ (I_1 \cup I_2) &\vdash_t I_3 \\ (I_1 \cup I_2 \cup I_3) &\vdash_t I_4 \end{aligned}$$

et cetera, until it is not possible to add any more axioms/integers to I_j. Observe that part 1 of the derivation rule enforces that the set I_j added in any one step is not too large, while part 2 makes sure that any new (compared to $\bigcup_{i<j} I_i$) boundary variables for I_j are represented in the term t (in this sense the polynomials and t "depend" on each other).

Remark 5.2. Perhaps we should also point out explicitly that the support of a term is well-defined in the sense that there is one unique largest subset that we can t-infer. To see this, suppose that we t-infer two subsets S_1 and S_2. Then we can also t-infer $S_1 \cup S_2$ by first inferring S_1 and then inferring S_2 starting with S_1 instead of $\emptyset$. It is straightforward to verify that this is in accordance with the rules. Also, we remark that the total size of the inferred set is allowed to be larger than $s/2$; it is only the set of new integers inferred in each step that has this restriction on its size.

We will use Definition 5.1 to define the polynomials generating the ideal $V(t)$ that we will use to reduce t. Intuitively, for any t we want $V(t)$ to be not too large (so that R is nontrivial). This is where the expansion of the graph comes into play. Because of expansion, condition 2 should be difficult to satisfy for (medium-)large sets. Also by expansion, $Sup(c)$ for a constant $c \in \mathbb{F}$ is the empty set, and we cannot c-infer anything from it.

Formally, we let $V(t) = Span(\{f_i \mid i \in Sup(t)\})$ and define $R(t) = R_{V(t)}(t)$. We want to prove that this R is a linear operator as in Lemma 4.2. Let us start by proving that the generating set $Sup(t)$ of the ideal $V(t)$ cannot be too large if t has not too large degree.

Lemma 5.3 (Lem 3.16 in [AR03]). *If $Deg(t) \leq se/2$, then $|Sup(t)| \leq s/2$.*

Proof. By contradiction. Suppose that there exists a sequence $I_0 = \emptyset, I_1, I_2, \ldots, I_\tau$ such that $\bigcup_{j=1}^{r-1} I_j \vdash_t I_r$ for all $r = 1, 2, \ldots, \tau$ and $\left|\bigcup_{j=1}^{\tau} I_j\right| > s/2$. Fix $\tau^* \leq \tau$ to be minimal such that $\left|\bigcup_{j=1}^{\tau^*} I_j\right| > s/2$. Since we could t-infer all the I_j's, by condition 1 in Definition 5.1 the last set to be t-inferred, I_{τ^*}, must have been of size at most $s/2$. By the minimality of τ^* we also have $\left|\bigcup_{j=1}^{\tau^*-1} I_j\right| \leq s/2$, and putting these two facts together we get that $|\bigcup_{j=1}^{\tau^*} I_j| \leq s$. This means that this set of vertices on the left-hand side of $G(F)$ is small enough to meet the expansion guarantees, so it follows that $\left|\partial(\bigcup_{j=1}^{\tau^*} I_j)\right| > se/2$ (overloading notation as explained above).

By condition 2 in Definition 5.1, we must have that every new boundary element at each step belongs to $N(t)$. A conservative estimate of all of these boundary elements is $\partial(\bigcup_{j=1}^{\tau^*} I_j)$. This is so since an element can be in $\partial(\bigcup_{j=1}^{r-1} I_j)$ at time $r-1$ and then disappear from the boundary in $\partial(\bigcup_{j=1}^{r} I_j)$ if it is contained in I_r, but if an element is in the boundary at time τ^* then it certainly was so also at the time when it was added. Thus the size of the boundary at time τ^* is a lower bound on the number of variables in t and hence on the degree, i.e., $\left|\partial(\bigcup_{j=1}^{\tau^*} I_j)\right| \leq Deg(t)$.

But combining the preceding two paragraphs we get that $Deg(t) \geq \left|\partial(\bigcup_{j=1}^{\tau^*} I_j)\right| > se/2$. This is a contradiction. The lemma follows. □

Now we can intuitively see that we are in good shape. If $G(F)$ is an expander with expansion rate $e \geq 1$ then any subset of at most s clauses is satisfiable. Hence, any low-degree monomial t as in Lemma 5.3 will get an ideal $V(t)$ generated from a satisfiable subset of clauses, so in this case we will have $1 \notin V(t)$. This is not the way the proof in [AR03] goes, but for random k-CNF formulas it might be that we could actually do (parts of) the argument that remains in this way.

But at this point in the lecture we ran out of time... So we did not have time to prove that our reduction operator R as defined above satisfies the conditions in Lemma 4.2. For completeness, we present a proof of this below.

A The Part of the Proof That We Did Not Do in Class

We want to prove that the reduction operator R defined in Section 5 meets the requirements of Lemma 4.2. A key step in the proof of this is to show that the reduction $R(t)$ of t does not change if we compute it modulo an ideal generated by a slightly larger set than $Sup(t)$, as long as it is not too large. We continue to overload notation so that $Span(I)$ is a shorthand for $Span(\{f_i \mid i \in I\})$.

Lemma A.1 (Lemma 3.17 in [AR03]). *Assume that t is a term and that $I \subseteq [m]$ is such that $I \supseteq Sup(t)$ and $|I| \leq s/2$. Then $R_{Span(I)}(t) = R_{Span(Sup(t))}(t)$.*

In order to establish Lemma A.1 we will need two (even more) technical lemmas. In what follows below, let us adopt vector notation signifying that $\vec{x}$ is a set of variables $x_1, \ldots, x_n$ and $\vec{P}$ is a set of polynomials $P_1, \ldots, P_\ell$ (for ranges of the indices to be determined from context). The first lemma that we need might look rather abstract (which it is), but it will come in handy very soon.

Lemma A.2 (Lemma 3.14 in [AR03]). *Let $\vec{y}, \vec{v}, \vec{z}$ be a partition of $\vec{x} = \{x_1, \ldots, x_n\}$, let $\vec{P} = \vec{P}(\vec{y}, \vec{v})$ be a set of polynomials over $\vec{y} \cup \vec{v}$, and let $Q = Q(\vec{v}, \vec{z})$ be a polynomial over $\vec{v} \cup \vec{z}$ such that $(z - b)$ divides Q for some z in $\vec{z}$ and some b in $\{0, 1\}$. Suppose that the term $t(\vec{y}, \vec{v})$ does not contain any of the z-variables and is reducible modulo $Span(\vec{P}, Q)$ with respect to $\prec$. Then t is reducible modulo $Span(\vec{P})$.*

Proof. By assumption, there exists some polynomial $f \in Span(\vec{P}, Q)$ such that $t = LT(f)$. By the implicational completeness of polynomial calculus,[6] it holds that $f \in Span(f_1, \ldots, f_t)$ if and

[6] Which is discussed in more detail in the scribe notes for lecture 6.

only if $f_1, \ldots, f_t \vDash f$, i.e., if for every α such that $f_i(\alpha) = 0$ for all i, it must also hold that $f(\alpha) = 0$. Hence, it follows from the assumptions in the lemma that $\vec{P}, Q \vDash f$.

Let us now apply the restriction setting $z = b$. Since implications are preserved under restrictions, we have $\vec{P}\restriction_{z=b}, Q\restriction_{z=b} \vDash f\restriction_{z=b}$. Let us look closer at this implication. $\vec{P}$ does not contain z by assumption, so $\vec{P}\restriction_{z=b} = \vec{P}$. Since also by assumption $(z-b)$ is a factor of Q we have $Q\restriction_{z=b} = 0$. Consequently, the implication can be written as $\vec{P} \vDash f\restriction_{z=b}$, and using completeness again we conclude that $f\restriction_{z=b} \in \mathit{Span}(\vec{P})$. Now we observe that since the term t does not contain the variable z it is not affected by the restriction setting $z = b$, so it must remain the leading term also in $f\restriction_{z=b}$, i.e., $t = \mathit{LT}(f\restriction_{z=b}) = \mathit{LT}(f)$. Restating this in words, there is a polynomial in the ideal $\mathit{Span}(\vec{P})$ such that t is the leading term of this polynomial (namely $f\restriction_{z=b}$), which is exactly the definition of what it means that t should be reducible modulo $\mathit{Span}(\vec{P})$. This proves the lemma. □

The second lemma formalizes the (hopefully fairly obvious) claim that if we reduce a term t modulo a set of polynomials $\vec{P}$, then the remaining irreducible terms will not contain any variables other than the ones in t and $\vec{P}$.

Lemma A.3. *Let t be a term and $\vec{P}$ be a set of polynomials, and let $t = q + r$ for $q \in \mathit{Span}(\vec{P})$ and $r = R_{\mathit{Span}(\vec{P})}(t)$. Then $\mathit{Vars}(r) \subseteq \mathit{Vars}(t) \cup \mathit{Vars}(\vec{P})$.*

Proof. Write $t = q + r$ for $q \in \mathit{Span}(\vec{P})$ and $r = R_{\mathit{Span}(\vec{P})}(t)$, and suppose there is a variable $z \in \mathit{Vars}(r) \setminus \mathit{Vars}(t) \cup \mathit{Vars}(\vec{P})$. As in the proof of Lemma A.2 we can apply the restriction $z = b$ on the implication $\vec{P} \vDash q = t - r$. Since z does not occur in $\vec{P}$ or t we get that $\vec{P} \vDash t - r\restriction_{z=b}$, and we can write $t = q' + r'$ for $q' \in \mathit{Span}(\vec{P})$ and $r' = r\restriction_{z=b}$. Note that r' is also a linear combination of irreducible terms since restricting an irreducible term can never make it reducible. But if so we get $q = q'$ and $r = r' = r\restriction_{z=b}$ by uniqueness, which shows that z does not appear in r after all. □

Now we can prove Lemma A.1.

Proof of Lemma A.1. By assumption $\mathit{Sup}(t) \not\vdash_t I \setminus \mathit{Sup}(t)$. Hence by condition 2 in Definition 5.1 there exists a $j \in I \setminus \mathit{Sup}(t)$ such that f_j contains a variable $z \in \partial I$ that does *not* occur in t. Apply Lemma A.2 with $\vec{P} = \{f_i \mid i \in I \setminus \{j\}\}$ and $Q = f_j$. Since f_j encodes a disjunctive clause C_j, we know that the fact that z appears implies that f_j is divisible by $z - b$ for $b \in \{0, 1\}$ chosen depending on the sign of z in C_j. Hence, Lemma A.2 says that if t is reducible modulo $\mathit{Span}(\vec{P}, Q)$, t is also reducible modulo $\mathit{Span}(\vec{P})$. (Restating this in our overloaded notation we have that if t is reducible modulo $\mathit{Span}(I)$, then t is also reducible modulo $\mathit{Span}(I \setminus \{j\})$.

Furthermore, we claim that $R_{\mathit{Span}(\vec{P})}(t) = R_{\mathit{Span}(\vec{P},Q)}(t)$, i.e., that all irreducible terms in $R_{\mathit{Span}(\vec{P})}(t)$ are irreducible also modulo $\mathit{Span}(\vec{P}, Q)$. This follows from repeated application of Lemma A.2 if we can prove that none of these terms contains the variable z. To see that this latter statement must be true, write $t = q + r$ for $q \in \mathit{Span}(\vec{P})$ and $r = R_{\mathit{Span}(\vec{P})}(t)$ and appeal to Lemma A.3.

Summing up, we have proven that $R_{\mathit{Span}(I)}(t) = R_{\mathit{Span}(I\setminus\{j\})}(t)$ for $j \in I \setminus \mathit{Sup}(t)$. By induction, we can eliminate all $j \in I \setminus \mathit{Sup}(t)$ one by one, which proves the lemma. □

With the help of Lemma A.1 we are now able to verify that R satisifies the conditions in Lemma 4.2.

1. R **is linear.** This is by construction. We define the operator on all monomials in $S_{n,d}(\mathbb{F})$ and extend it by linearity.

2. **For any axiom f_i it holds that** $R(f_i) = 0$. Consider $t_i = \prod_{x_j \in \mathit{Vars}(f_i)} x_j$. For all terms $t \in f_i$ we have that $\mathit{Vars}(t) \subseteq \mathit{Vars}(t_i)$. Lemma 5.3 yields the upper bound

$|Sup(t_i)| \leq s/2$. Clearly, for each $t \in f_i$ it holds that $Sup(t) \subseteq Sup(t_i)$. Since $i \in Sup(t_i)$, it holds that $R_{Span(Sup(t_i))}(f_i) = 0$ (this is just a complicated way of stating the obvious fact that $f_i \in Span(f_i)$, so "the remainder" modulo the ideal in this case is 0). But then it follows that $R(f_i) = 0$, since for each $t \in f_i$ Lemma A.1 says that $R_{Span(Sup(t))}$ agrees with $R_{Span(Sup(t_i))}(t)$—i.e., it does not change the reduction operator when we replace $Span(Sup(t))$ by the somewhat larger ideal $Span(Sup(t_i))$—and by linearity it then holds that

$$\begin{aligned} R(f_i) &= \textstyle\sum_{t\in f_i} R(t) = \sum_{t\in f_i} R_{Span(Sup(t))}(t) = \\ &= \textstyle\sum_{t\in f_i} R_{Span(Sup(t_i))}(t) = R_{Span(Sup(t_i))}\big(\sum_{t\in f_i} t\big) = R_{Span(Sup(t_i))}(f_i) = 0 \ . \end{aligned} \tag{A.1}$$

3. **For any term t such that $Deg(t) < se/2$ and any variable x it holds that $R(xt) = R(xR(t))$.** This is clearly the crucial property where we should expect to have to work the hardest, so let us argue carefully step by step.

   Firstly, as we already noted several times by now we have

$$|Sup(xt)| \leq s/2 \tag{A.2}$$

   by Lemma 5.3. By combining the degree bound in Equation (A.2) with what we just proved in Lemma A.1, we deduce that

$$R_{Span(Sup(t))}(t) = R_{Span(Sup(xt))}(t) \ . \tag{A.3}$$

   Secondly, for any term $t' \in R_{Span(Sup(t))}(t)$ it holds that

$$Sup(xt') \subseteq Sup(xt) \ . \tag{A.4}$$

   This can be derived from the fact that

$$N(t') \subseteq N(t) \cup N(Sup(t)) \ , \tag{A.5}$$

   i.e., that all variables in t' must come from somewhere on the right-hand side of (A.5), which is what we showed in Lemma A.3. To prove that (A.4) follows from (A.5), it is sufficient to show that if I_j xt'-infers I_{j+1} for $I_j \subseteq Sup(xt)$, then $Sup(xt)$ xt-infers I_{j+1}. Suppose $I_j \vdash_{xt'} I_{j+1}$. By Definition 5.1 and Equation (A.5) we obtain that

$$\begin{aligned} \partial I_{j+1} &\subseteq N(xt') \cup N(I_j) \\ &= N(x) \cup N(t') \cup N(I_j) \\ &\subseteq N(x) \cup N(t) \cup N(Sup(t)) \cup N(I_j) \\ &\subseteq N(xt) \cup N(Sup(xt)) \ , \end{aligned} \tag{A.6}$$

   which again by Definition 5.1 means that $Sup(xt) \vdash_{xt} I_{j+1}$. Combining (A.4) and (A.2) and applying Lemma A.1 again we deduce that

$$R_{Span(Sup(xt'))}(xt') = R_{Span(Sup(xt))}(xt') \ . \tag{A.7}$$

   Finally, it is a straightforward to check that

$$R_{Span(Sup(xt))}\big(xR_{Span(Sup(xt))}(t)\big) = R_{Span(Sup(xt))}(xt) \tag{A.8}$$

(this is the analogue of the statement for integers a, b, c that $ab \equiv a(b \bmod c) \pmod c$). Putting all the pieces together, we get

$$\begin{aligned}
R(xR(t)) &= \textstyle\sum_{t' \in R(t)} R(xt') && \text{[by linearity]} \\
&= \textstyle\sum_{t' \in R(t)} R_{Span(Sup(xt'))}(xt') && \text{[by the definition of } R\text{]} \\
&= \textstyle\sum_{t' \in R(t)} R_{Span(Sup(xt))}(xt') && \text{[by (A.7)]} \\
&= R_{Span(Sup(xt))}(xR(t)) && \text{[by linearity]} \\
&= R_{Span(Sup(xt))}(xR_{Span(Sup(t))}(t)) && \text{[by the definition of } R\text{]} \\
&= R_{Span(Sup(xt))}(xR_{Span(Sup(xt))}(t)) && \text{[by (A.3)]} \\
&= R_{Span(Sup(xt))}(xt) && \text{[by (A.8)]} \\
&= R(xt) && \text{[by the definition of } R\text{]}
\end{aligned}$$

which is exactly what we needed to show.

4. $R \not\equiv 0$**, i.e., in particular** $R(1) \neq 0$**.** Consider Definition 5.1 with $t = 1$ and $I = \emptyset$. Then $N(t) = N(1) = \emptyset$, but by the expansion properties of $G(F)$ it holds for any I_1 such that $|I_1| \leq s/2$ that $\partial I_1 \neq \emptyset$. Hence, we cannot t-infer anything from $\emptyset$ for $t = 1$, which means that $Sup(1) = \emptyset$. Thus, $R(1) = R_{Span(Sup(1))}(1) = R_{Span(\emptyset)}(1) = 1$.

We have now verified all the conditions in Lemma 4.2, so Theorem 1.1 follows. Hence, for random k-CNF formulas $F \sim \mathcal{F}_k^{n,\Delta}$ with $\Delta = \mathrm{O}(1)$ large enough it holds almost surely that F is unsatisfiable and requires linear refutation degree (and thus exponential refutation size in both PC and PCR).

Acknowledgements

Many people have helped improve these notes since the lecture was first given. Among all these people, the lecturer would like to especially acknowledge the assistance of Yuval Filmus, Mladen Mikša, and Bangsheng Tang.

References

[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.

[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.

[Raz98] Alexander A. Razborov. Lower bounds for the polynomial calculus. *Computational Complexity*, 7(4):291–324, December 1998.