6.885 Algebra and Computation November 9, 2005

Lecture 17

Lecturer: Madhu Sudan *Scribe: Paul Valiant*

The problem which we consider today is the following: given m polynomials in n variables over a field K,

$$f_1, ..., f_m \in K[x_1, ..., x_n],$$

does there exist a vector $\bar{a} \stackrel{\text{def}}{=} (a_1, ..., a_n) \in K^n$ such that

$$f_1(\bar{a}) = ... = f_m(\bar{a}) = 0?$$

We work over fields K that are *algebraically closed* – that is, fields where every polynomial has a root.

Questions of this form are said to form the *existential theory over* K.

We point out that this problem is a generalization of problems like 3-SAT. In 3-SAT we want to satisfy m clauses of the form $(x_1 \wedge \neg x_2 \wedge x_3)$. Algebraically, this clause becomes

$$(x_1)(1 - x_2)(x_3),$$

and we wish to simultaneously satisfy n such clauses. Clearly this problem is a special case of the above.

We note that we can make the problem even broader by adding more layers of quantification: given $f_1, ..., f_m$ is it the case that $\exists a_1 \forall a_2 \exists a_3 ... Q a_n$ s.t.

$$f_1(\bar{a}) = ... = f_m(\bar{a}) = 0?$$

These questions form the *quantified theory over* K.

We can also consider the *search* versions of these decision problems: given $f_1, ..., f_m \in K[\bar{x}, \bar{y}]$, for what choices of $\bar{b} \in K^n$ will it be the case that $\exists a_1, ..., a_n$ s.t.

$$f_1(\bar{a}, \bar{b}) = ... = f_m(\bar{a}, \bar{b}) = 0?$$

Returning to the original problem, we note that the question is in the version of NP over K, since we can provide short vectors $\bar{a}$ as witnesses of feasibility. We now ask: can we provide short proofs of infeasibility? We answer this by introducing Hilbert's *Nullstellensatz* theorem.

Theorem 1 (Nullstellensatz) *Given polynomials $f_1, ..., f_n \in K[\bar{x}]$, the following two statements are equivalent:*

$$\exists \bar{a} \quad \textit{such that} \quad f_1(\bar{a}) = ... = f_m(\bar{a}) = 0$$

and

$$\neg \left(\exists q_1, ..., q_m \in K[\bar{x}] \quad \textit{such that} \quad \sum f_i q_i = 1 \right).$$

We might prove this later in the course. For now we note that one direction is clear: if there exists a linear combination of the polynomials f that is uniformly 1, then the polynomials can never be simultaneously 0.

To re-express some of the above in more algebraic language, we make the following definition:

Definition 2 (Variety) *Given polynomials $f_1, ..., f_m$ over some vector $\bar{a} \in k^n$ define the set*

$$S \stackrel{\text{def}}{=} \{\bar{a} | f_1(\bar{a}) = ... = f_m(\bar{a}) = 0\}.$$

We call S a variety.

Motivated by the Nullstellensatz, we look at the ideal

$$I(f_1, ..., f_m) = \{\sum f_i q_i | q_i \in K[\bar{x}]\}.$$

Hilbert's Nullstellensatz may be re-expressed as the following:

$$V(\bar{f}) = \emptyset \Leftrightarrow 1 \notin I(\bar{f}).$$

We note that the definition of variety allows us to go from sets of polynomials to sets of points. To go in the reverse direction, we create an ideal from a set of points: given a set $S \subseteq K^n$, define

$$I_S \stackrel{\text{def}}{=} \{f | \forall \bar{a} \in S, f(\bar{a}) = 0\}.$$

Thus we can go from f to $V(f)$ to $I_{V(f)}$. We may now compare this ideal to the straightforward ideal $I(f)$. By definition, any linear combination of the polynomials in f takes on zero values in $V(f)$, so

$$I(f) \subseteq I_{V(f)}.$$

However, the reverse inclusion does not necessarily hold: suppose $f = h^2$ for some polynomials h. Then $V(f) = V(h)$, which implies $I_{V(f)} = I_{V(h)}$. However, $I(f) \not\subseteq I(h)$.

By the Nullstellensatz, in order to determine if a variety is nonempty, we need only determine if 1 is in a specific ideal. The natural generalization of this is the ideal membership problem:

Given $f, f_1, ..., f_m \in K[\bar{x}]$, is $f \in I(\bar{f})$?

Unfortunately for us, Mayr and Meyer proved that this problem is EXPSPACE complete, even if all the coefficients are 0 or 1 and all the degrees are at most 2.

Nevertheless, we try to proceed with our intuitions to see what will develop.

The natural thing to try is the following: given $f_1, ..., f_m \in K[\bar{x}]$, let $g = gcd(f_1, ..., f_m)$. Thus to determine whether $f \in I(\bar{f})$, we need only determine whether $f \in I(g)$, which can be computed via another gcd. However, in the multivariate case this will not work because some of these notions become ill-defined.

Consider, for example when

$$\begin{aligned} f &= x^2y + xy^2 + y^2 \\ f_1 &= xy - 1 \\ f_2 &= y^2 - 1. \end{aligned}$$

The algorithm we expect to work would try to find the smallest remainder of f after subtracting out linear combinations of f_1 and f_2, and if this remainder is non-zero, we conclude that $f \notin I(f_1, f_2)$. The problem comes when we try to define smallest. Consider the following three candidate representations of f:

$$\begin{aligned} f &= (x + y)f_1 + f_2 + (x + y + 1) \\ f &= xf_1 + (x + 1)f_2 + (2x + 1) \\ f &= (x + 2y)f_1 + (1 - x)f_2 + (2y + 1), \end{aligned}$$

where the three remainders produced are $x+y+1, 2x+1$, and $2y+1$ respectively. Which of these should we call the true remainder? Which of these should we call the smallest?

To resolve the above difficulties, we create an order on monomials, and then extend this to an order on polynomials. If we do this right, then we should be able to:

a. Find the least possible remainder

b. If this remainder is nonzero, conclude that $f \notin I(\bar{f})$

In the following, we take the notation $\bar{x}^{\bar{a}}$ to mean $x_1^{a_1} \cdot x_2^{a_2} \cdot ....$
We call an ordering "$\leq$" *admissible* if it satisfies the following conditions:

1. $\leq$ is a total order
2. $\bar{x}^{\bar{a}} \leq \bar{x}^{\bar{b}} \Rightarrow \bar{x}^{\bar{a}}\bar{x}^{\bar{c}} \leq \bar{x}^{\bar{b}}\bar{x}^{\bar{c}}$
3. $\forall \bar{a}, \quad \bar{x}^0 \leq \bar{x}^{\bar{a}}$

We thus define the remainder of f with respect to $f_1, ..., f_m$ to be the minimal polynomial r such that $(f - r) \in I(\bar{f})$.

We define the degree of a polynomial $deg(f)$ to be that maximal $\bar{a}$ under $\leq$ such that $\bar{x}^{\bar{a}}$ is in the support of f.

The lexicographic ordering, and its variants are standard examples of *admissible* orders.

Given a polynomial $f(\bar{x}) = c_{\bar{a}}x^{\bar{a}} + ...$ where $\bar{a}$ is the largest monomial, we define the *leading monomial* of f to be $\bar{x}^{\bar{a}}$, the *leading coefficient* of f to be $c_{\bar{a}}$, and the *leading term* to be $c_{\bar{a}}x^{\bar{a}}$.

We are now in a position to begin the analysis of an algorithm for finding the least remainder. We will see that the analysis shares many features with the analysis of the permutation group membership algorithm.

We start out, as with the permutation group membership algorithm, with a definition of a *basis* that at first appears neither feasible nor helpful.

Definition 3 (Groebner basis) *Given $J = I(f_1, ..., f_m)$, the polynomials $g_1, ..., g_t$ form a Groebner basis for J if the following two conditions hold:*

1. $I(g_1, ..., g_t) = J$
2. $I(LT(g_1), ..., LT(g_t)) = I(LT(J))$

where $LT(f)$ is the leading term of f, and $LT(J)$ is the set of all the leading terms of elements of J.

We then define the remainder as follows:

Definition 4 (Weak Remainder) *The* weak remainder *of f with respect to $f_1, ..., f_m$ is an r such that for all monomials $\bar{x}^{\bar{a}} \in r$, and for all i,*

$$LT(f_i) \not| \bar{x}^{\bar{a}}.$$

The crucial fact here is that weak remainders with respect to Groebner bases are *unique*. That is, the weak remainder is independent of the choice of Groebner basis. We prove this here.

We first prove a few preliminaries. The first is about the structure of ideals generated by monomials.

Claim 5 *Given $\bar{x}^{\bar{a}} \in I(\bar{x}^{\bar{a_1}}, ..., \bar{x}^{\bar{a_m}})$, there exists i such that $\bar{x}^{\bar{a_i}} | \bar{x}^{\bar{a}}$.*

Proof If $\bar{x}^{\bar{a}}$ is in the ideal, then $\bar{x}^{\bar{a}}$ is expressible as

$$\bar{x}^{\bar{a}} = \sum_{j=1}^{m} q_j \bar{x}^{\bar{a}_j},$$

for polynomials q_j. Since $\bar{x}^{\bar{a}}$ is a monomial, it must be divisible by some monomial from the right, as desired. ■

We next claim the following:

Claim 6 *If r is a weak remainder of f with respect to a Groebner basis $g_1, ..., g_t$ for an ideal J, then for all monomials $\bar{x}^{\bar{a}} \in r$, $\bar{x}^{\bar{a}} \notin I(LT(J))$.*

Proof By the definition of a Groebner basis

$$I(LT(J)) = I(LT(g_1), ..., LT(g_t)).$$

Further, by the contrapositive of the previous claim, if

$$\forall i, \;\; LT(g_i) \not| \bar{x}^{\bar{a}}$$

then

$$\bar{x}^{\bar{a}} \notin I(LT(g_1), ..., LT(g_t)).$$

Since this last condition is the definition of a weak remainder, the claim is proved. ■

We now prove uniqueness of weak remainders.

Claim 7 *The weak remainder of f with respect to a Groebner basis $g_1, ..., g_t$ is unique.*

Proof Suppose for the sake of contradiction that there were two different weak remainders, r_1, r_2. Thus

$$r_1 = f - A$$

and

$$r_2 = f - B$$

for $A, B \in J = I(g_1, ..., g_t)$. Thus $r_1 - r_2 \in J$, where by assumption $r_1 - r_2 \neq 0$.

Taking the leading term of this expression, we have that

$$LT(r_1 - r_2) \in LT(J).$$

Since all the monomials of $r_1 - r_2$ are monomials of r_1 or r_2, without loss of generality we may assume that $LT(r_1 - r_2)$ is a monomial of r_1. However, by the previous claim, no monomials of r_1 are in $LT(J)$, the desired contradiction. Thus the weak remainder is unique. ■