6.885 Algebra and Computation November 14, 2005

Lecture 18

Lecturer: Madhu Sudan *Scribe: Alexey Spiridonov*

1 Today

- We cover Gröbner basis recognition, generation, and the resulting algorithm for ideal membership.
- We won't produce any complexity estimates this lecture, but only a finite time decision procedure for testing membership in ideals.

2 Notation and Definitions

Some of these overlap with the previous lecture, but I repeat them here for the sake of completeness.

Ambient polynomial ring We'll take k to be an algebraically closed field, and take all our polynomials in $k[x_1, \ldots, x_n]$.

Monomial ordering A total ordering $\geq$ on monomials $x^\alpha = x_1^{\alpha_1} \ldots x_n^{\alpha_n}$ satisfying

1. $x^\alpha \geq 1$ for al
2. $x^\alpha \geq x^\beta \Rightarrow x^{\alpha+\gamma} \geq x^{\beta+\gamma}$

For our purposes the lexicographic ordering will suffice (but there are many other useful ones). We compare x^α and x^β thus: if the first k indices agree: $\alpha_i = \beta_i, i \leq k$ and the kth differ, we decide based on that index $\alpha_i \leq \beta_i \Rightarrow \alpha \leq \beta$, and the reverse.

Leading monomial Denoted $LM(f)$, this is the greatest monomial of $f \in k[x_1, \ldots, x_n]$ according to our chosen ordering.

Leading coefficient The coefficient in front of $LM(f)$, denoted $LC(f)$.

Leading term $LT(f) = LC(f)LM(f)$.

3 Ideal Membership Problem

Given an ideal $J = (f_1, f_2, \ldots, f_m)$ and a polynomial f_0 in the ring, we would like to decide whether $f \in J$. The idea is simple: f_0 is in J if and only if it can be written $\sum p_i f_i$. If we "divide" the latter by representation by f_i and take the remainder, we eliminate the term containing f_i; dividing out by all f_i we ought to get 0, if and only if $f_0 \in J$. However, generically, division is poorly defined: remainders depend on the order of the division, choice of basis. Ideally, we'd like to divide by the entire (infinite) ideal J; that's impractical, but last time we saw that dividing by a Gröbner basis is well-behaved.

So, here's an outline of our algorithm:

1. Fix a monomial ordering (say, lexicographic).
2. Find a Gröbner basis $g_1, \ldots, g_t$ for J. (*a priori*, it's not clear that a finite one exists given our starting basis $f_1, \ldots f_m$). Dividing by this basis in order is a well-behaved operation.
3. Divide f_0 by $g_1, \ldots, g_t$.
4. If the remainder is 0 then $f_0 \in J$, else $f_0 \in J$.

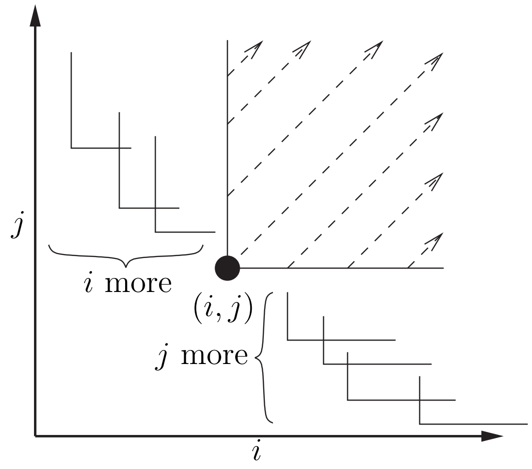

Figure 1: The two-variable case of Dickson's Lemma.

4 Gröbner Bases

Definition. A *Gröbner basis (GB)* is a set of polynomials $g_1, \ldots, g_t \in k[x_1, \ldots, x_n]$ satisfying the following properties: (we give two version of property 1 – 1a from the previous lecture, and 1b which we will show is equivalent)

1. $g_1, \ldots, g_t \in J$ (last time: generate J; we'll see these are equivalent)
   (a) Old variant: $g_1, \ldots, g_t$ generate J
   (b) New variant: $g_1, \ldots, g_t \in J$
2. $I(LT(g_1), \ldots, LT(g_t)) = I(LT(J))$, where $LT(J)$ is the set of leading terms of J. (Note that it's also an ideal.)

So, we need a set of monomials generating the monomial ideal $I(LT(J))$. The set of generators $LT(J)$ is infinite, so we'd better make sure that we can actually produce a finite GB for this ideal. That's the subject of the next section.

In the process, we will also see that 1b implies that $g_1, \ldots, g_t$ generate J.

5 Hilbert Basis Theorem

Theorem (Hilbert Basis Theorem)**.** *Every ideal in* $k[x_1, \ldots, x_n]$ *has a finite basis.*

Remark. Though we assumed at the start of the lecture that k is a field, this theorem holds for k any Nötherian ring (ring in which every ideal is finitely generated).

The proof follows easily from the following lemma.

Lemma (Dickson's Lemma)**.** *Every monomial ideal* J *in* $K[x_1, \ldots, x_n]$ *has a finite basis.*

Proof. We will work by induction on n, the number of variables.

The case $n = 1$ is trivial. If our ideal is generated by $\{x^{i_1}, x^{i_2}, \ldots, x^{i_n}, \ldots\}$, it is also generated by x^{i_0}, with $i_0 = \min\{i_1, \ldots, i_n\}$.

For two variables, $J = \{x^{i_1}y^{j_1}, x^{i_2}y^{j_2}, \dots\}$, we can draw a picture; see Figure 1. It depicts all monomials on an integer grid (e.g. x^5y^2 is at $(5,2)$). Pick a monomial x^iy^j from the generating set such that $i+j$ is minimal. Then, it generates all monomials in the dashed region in the figure. That leaves a vertical strip i wide and a horizontal strip j wide. There's room only for $i+j$ more generating monomials in those strips, so the ideal has at most $i+j+1$ generators.

In more variables, the bounds stop having a nice form and we can't draw pictures any more. So, we'll settle for arguing, in the same vein, that the generating set is finite.

Suppose that all $n-1$-variable ideals are finitely generated. Consider an n-variable ideal J. Now, take $x^\alpha = x_1^{\alpha_1}\dots x_n^{\alpha_n} \in LT(J)$. For an index i and a degree β, define

$$J_{(i,\beta)} = \{x_1^{\gamma_1}x_2^{\gamma_2}\dots x_{i-1}^{\gamma_{i-1}}x_{i+1}^{\gamma_{i+1}}\dots x_n^{\gamma_n} \text{ such that } x_1^{\gamma_1}x_2^{\gamma_2}\dots x_{i-1}^{\gamma_{i-1}}x_i^{\beta}x_{i+1}^{\gamma_{i+1}}\dots x_n^{\gamma_n}\}.$$

Now, $J_{(i,\beta)}$ is a set of monomials in $n-1$ variables, so the corresponding ideal is finitely generated. Let C be the union of all sets of generators of $J_{(i,\beta)}$ for $i = 1\dots n, \beta = 0\dots\alpha_i - 1$. Then, we claim that $C \cup \{\alpha\}$ generates J. Indeed, take a monomial x^δ in J; if all $\delta_i \ge \alpha_i$, then x^α is a generator. Otherwise, there's some i such that $\delta_i < \alpha_i$, and in that case $x^\delta \in J_{i,\delta_i}$. That finishes the proof of the lemma. □

That doesn't give any nice complexity bound on the size of the generating set. We will not show this, but the complexity of the resulting ideal membership algorithm is very bad (EXPSPACE).

Now, we use Dickson's lemma to prove Hilbert's Basis Theorem. Actually, we will prove more:

Theorem 5.1. *Every polynomial ideal has a finite Gröbner basis.*

Proof. The idea of the proof is: we will pick out some polynomials from the ideal J, such that $\{LT(g_i)\}$ generate $LT(J)$. This requires this lemma we promised we'd prove:

Lemma. *In the definition of the Gröbner basis,* $g_1, \dots, g_t \in J \Rightarrow (g_1, \dots, g_t) = J$ *(assuming part* 2 *of the definition).*

Proof. We will prove this by contradiction. Take $g_1, \dots, g_t$ as in the statement of the theorem.

Last lecture we proved the following helpful fact: if we have polynomials $g_1, \dots, g_t$ satisfying 1a and 2, it follows we can canonically divide by these polynomials. As discussed before, this is an ideal membership test: if a polynomial f is in $I(g_1, \dots, g_t)$, the division returns 0, otherwise a nonzero remainder.

Suppose $\exists f \in J$ such that $f \notin I(g_1, \dots, g_t)$. Let's compute the remainder after dividing f by $g_1, \dots, g_t$; we get $f = r + \sum g_i q_i$. Moreover (again from last lecture), no monomial of r is divisible by $LT(g_i)$, for any i. Now, consider $LT(r)$; since $r \in J$, we get $LT(r) \in LT(J)$, so $LT(r) = I(LT(g_1), \dots, LT(g_t)) = LT(J)$, so there exists i such that $LT(g_i) | LT(r)$. Contradiction. □ So, we

saw that 1b implies 1a.

Now, consider $LT(J)$; by the lemma, this is generated by a finite set $g_1', g_2', \dots, g_t'$. By definition of $LT(J)$, every g_i' was obtained from J by taking the leading term of some $g_i \in J$. Look at the set $g_1, \dots, g_t$; by the lemma, this Gröbner basis generates J, and so we are done. □ Next, we

will see that a Gröbner basis is *essentially* unique. There are two obvious problems that get in the way of uniqueness. First, a GB plus arbitrary element is still a GB, so we need the following condition.

Definition. $(g_1, \dots, g_t)$ is a *minimal* Gröbner basis for $J = I(g_1, \dots, g_t)$ if for all i,

$$LT(g_i) \notin I(LT(g_1), \dots, LT(g_{i-1}), LT(g_{i+1}), \dots, LT(g_t)).$$

So, in our quest for a unique GB, we will drop elements g_i such that

$$LT(g_i) \in I(LT(g_1), \dots, LT(g_{i-1}), LT(g_{i+1}), \dots, LT(g_t)),$$

one-by-one, until our GB becomes minimal.

Is the resulting basis unique? No. For instance, $\{x, y\}$ and $\{x, x + y\}$ are minimal bases for the same ideal. The second one obviously looks "worse" than the first. The following definition makes the meaning clear.

Definition. A minimal GB $(g_1, \dots, g_t)$ is *reduced* if

$$g_i = Rem(g_i; g_1, \dots, g_{i-1}, g_{i+1}, \dots, g_t)$$

for all i.

We leave it as an exercise to show that if we have two minimal GB for J: $(g_1, \dots, g_t)$ and $(g'_1, \dots, g'_{t'})$, then

$$\{LT(g_1), \dots, LT(g_t)\} = \{LT(g'_1), \dots, LT(g'_t)\}$$

This implies that $t' = t$. To make a reduced basis, we take a minimal basis, and for every g_i replace it by

$$Rem(g_i; g_1, \dots, g_{i-1}, g_{i+1}, \dots, g_t).$$

It's not difficult to check that this procedure does not alter the leading terms, and so the result is a Gröbner basis. It also follows that a reduced GB is unique; take two bases g_i and g'_i, and arrange the indices so that leading terms are pairwise equal. Then, for some j, $g_j \neq g'_j \Rightarrow g_j - g'_j \in J$. But, the bases are reduced means that $LT(g_i) \not| LT(g_j - g'_j) \forall i$, and hence $LT(g_j - g'_j) \notin I(LT(J))$, which contradicts g_i being a Gröbner basis.

Remark. The notion of a reduced GB is somewhat parallel to the notion of a strong generating set from earlier in the course.

6 Recognizing Gröbner Bases

The above proof is almost, but not quite constructive: we didn't specify how to construct a finite monomial basis (although that's not difficult), and it isn't immediately obvious how to get an inverse image in J of a given monomial.

However, we will get an actual algorithm for constructing a GB by answering the question: how do we recognize whether a given basis is a GB?

We'll start by making division by a Gröbner basis even more canonical: we already know that the remainder is unique, but we'd like to also regularize the quotients. To do this, we will consider $Rem(f_0; g_1, \dots, g_t)$ with $g_1, \dots, g_t$ an ordered sequence. The procedure is: take the smallest i such that $LT(g_i)$ divides the largest (in our monomial ordering) monomial of f. That yields $f = f' + m_1 g_1$ with m_1 a monomial. Then, we iterate, each time taking the smallest i so that $LT(g_i)$ divides the largest monomial of $f^{(j)}$. Once there is no such i, we are left with the usual remainder r, and m_i monomials such that:

$$f = r + \sum m_i g_i.$$

However, the pieces of the quotient are special in that they are "reduced" – $m_i g_i \neq q_j g_j + \dots$ with $j < i$ and "..." denoting a remainder with smaller leading monomials. This regularized division will be used in a proof shortly.

Next, we need a special polynomial called the *syzygy* of f and g. It's a special polynomial in $I(f, g)$ of the form:

$$S(f, g) = f \cdot X - g \cdot Y$$

with X and Y chosen so that the leading terms of the two pieces are equal. We can write it explicitly:

$$S(f, g) = \frac{LC(g)LCM(LM(f), LM(g))}{LM(f)} \cdot f - \frac{LC(f)LCM(LM(f), LM(g))}{LM(g)} \cdot g.$$

The key claim now is:

Proposition. *The polynomials $g_1, \ldots, g_t$ form a Gröbner basis iff $\forall i, j$,*

$$Rem(S(g_i, g_j); g_1, \ldots, g_t) = 0.$$

Once we prove the proposition, constructing a Gröbner basis is straightforward. Start with some basis $g_1, \ldots, g_t$; if there are i, j such that $Rem(S(g_i, g_j); g_1, \ldots, g_t) \neq 0$, add this remainder to the basis. Any such remainder r has to be such that $LM(r)$ isn't generated by the $LM(g_i)$, so the monomial ideal $I(\{LM(g_i)\})$ gets bigger with every step. But, we have seen that any monomial ideal is finitely generated, so this algorithm must terminate. Thus, proving the proposition will give us a finite-time decision procedure for finding a Gröbner basis, and by proxy, a finite-time ideal membership procedure.

The proof of the proposition follows from

Claim. If $\forall i, j$, $Rem(S(g_i, g_j); g_1, \ldots, g_t) = 0$, then $\{LM(g_i)\}$ generate $LM(J)$.

The proof of this claim was left to the next lecture.