6.966 Algebra and Computation	October 19, 1998
Lecture 12	
Lecturer: Madhu Sudan	*Scribe: Rocco Servedio*

1 Introduction

This was the first lecture on the second main course topic, the complexity of solving multivariate polynomial equations. In this lecture we outlined the general family of problems we will be interested in and began to discuss the ideal membership problem. We gave an algorithm for "dividing" a fixed polynomial by several polynomials and introduced the notion of a Groebner basis. A useful reference for this material is the book "Ideals, Varieties and Algorithms" by Cox, Little and O'Shea.

2 The general question

In its most general form, the question we will be concerned with for this portion of the course is the following: let $f_1(x_1, \ldots, x_n), \ldots, f_s(x_1, \ldots, x_n)$ be a collection of s polynomials, where for each polynomial the degree in x_i is bounded by d_i. Let $Q_1, \ldots, Q_n$ be some sequence of quantifiers, i.e. each Q_i is either "$\exists$" or "$\forall$." What is the computational complexity of determining the truth value of the following statement:

$$\text{“}Q_1 x_1 Q_2 x_2 \cdots Q_n x_n \text{ such that } f_1(x_1, \ldots, x_n) = \ldots = f_s(x_1, \ldots, x_n) = 0\text{”?}$$

A moment's thought shows that this is a very general question (and that the computational complexity can be quite high). If we allow an arbitrary sequence of quantifiers, then we can express PSPACE-hard problems in the above format. We noted back in Lecture 1 that even if all quantifiers are restricted to be existential and $d_i = 2$, we can still express the NP-hard problem 3SAT. Similarly, if we take all quantifiers to be "$\forall$" we can capture co-NP; and by bounding the number of alternations between "$\exists$" and "$\forall$" we obtain the different levels of the polynomial hierarchy.

Today we will focus on the "NP" version of this problem in which all quantifiers are existential. More specifically, we consider the following question: let K be a field and let $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$ be polynomials each of degree at most d_i in x_i. Does there exist $(a_1, \ldots, a_n) \in K^n$ such that

$$f_1(a_1, \ldots, a_n) = \ldots = f_s(a_1, \ldots, a_n) = 0 \text{ ?}$$

3 Varieties and Ideals

The question we have asked leads naturally to the following definition:

Definition 1 *Let $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$. The* variety *of $(f_1, \ldots, f_s)$ is the set*

$$V(f_1, \ldots, f_s) = \{(a_1, \ldots, a_n) \in K^n \mid \forall i \in \{1, \ldots, s\} f_i(a_1, \ldots, a_n) = 0\}.$$

So the variety of a set of polynomials is the set of their common zeros. The question we are interested in could thus be succinctly written as "is $V(f_1, \ldots, f_s)$ nonempty?" This problem has an NP flavor to it because if the answer is yes, a short proof can be given which can be easily verified – the proof is simply a point $(a_1, \ldots, a_n) \in K^n$, and verification consists of checking that it is a zero for each polynomial f_i.

Now let $f_1, f_2 \in K[x_1, \ldots, x_n]$, and suppose that $(a_1, \ldots, a_n)$ is such that $f_1(a_1, \ldots, a_n) = f_2(a_1, \ldots, a_n) = 0$. Then it is clear that for any $q_1, q_2 \in K[x_1, \ldots, x_n]$, we will have

$$q_1(a_1, \ldots, a_n) f_1(a_1, \ldots, a_n) + q_2(a_1, \ldots, a_n) f_2(a_1, \ldots, a_n) = 0.$$

This observation motivates the following definition:

Definition 2 *Let* $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$. *The* ideal generated by $f_1, \ldots, f_s$ *is the following set of polynomials:*

$$I(f_1, \ldots, f_s) = \{\sum_{i=1}^{s} q_i(x_1, \ldots, x_n) f_i(x_1, \ldots, x_n) \;\mid\; \forall i \in \{1, \ldots, s\}\ q_i \in K[x_1, \ldots, x_n]\}.$$

So the ideal generated by a nontrivial set of polynomials is an infinite set of polynomials which includes the original set. Note that if $f_1, \ldots, f_s$ all vanish at a point $(a_1, \ldots, a_n)$, then so does every polynomial in $I(f_1, \ldots, f_s)$. This leads to the following simple claim:

Claim 3 *If* $1 \in I(f_1, \ldots, f_s)$, *then* $\nexists (a_1, \ldots, a_n) \in K^n$ *such that* $\forall i\ f_i(a_1, \ldots, a_n) = 0$.

Proof Suppose to the contrary that $(a_1, \ldots, a_n) \in K^n$ is a zero for each f_i. Since $1 \in I(f_1, \ldots, f_s)$, there exist $q_i \in K[x_1, \ldots, x_n]$ such that $1 = \sum_{i=1}^{s} q_i f_i$; evaluating at the point $(a_1, \ldots, a_n)$, we obtain

$$1 = \sum_{i=1}^{s} q_i(a_1, \ldots, a_n) f_i(a_1, \ldots, a_n) = \sum_{i=1}^{s} 0 = 0$$

which is a contradiction. ■

An equivalent form of Claim 3 is the following: if $1 \in I(f_1, \ldots, f_s)$, then for any $(a_1, \ldots, a_n) \in K^n$, there is some i such that $f_i(a_1, \ldots, a_n) \neq 0$. Later in the course we will prove the following partial converse, which is a weak version of Hilbert's Nullstellensatz ("null space theorem"):

Fact 4 *If* K *is an algebraically closed field and* $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$, *then*

$$1 \in I(f_1, \ldots, f_s) \iff \forall (a_1, \ldots, a_n) \in K^n\ \exists j \in \{1, \ldots, s\} \text{ such that } f_j(a_1, \ldots, a_n) \neq 0.$$

This means that over algebraically closed fields,

$$\exists q_1, \ldots, q_s \in K[x_1, \ldots, x_n]\ \sum_{i=1}^{s} q_i f_i = 1 \iff \forall (a_1, \ldots, a_n) \in K^n\ \exists j \in \{1, \ldots, s\}\ f_j(a_1, \ldots, a_n) \neq 0.$$

At first glance this equivalence seems to relate NP- and co-NP-type problems. However, as we will see in later lectures, the degrees of the polynomials q_i can be doubly exponential in n; hence even if the answer is "yes" there may not be a short proof which can be verified quickly.

4 The Ideal Membership Problem

An instance of the *ideal membership problem* consists of a distinguished polynomial $f \in K[x_1, \ldots, x_n]$ and a finite set of polynomials $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$. The problem is to determine whether $f \in I(f_1, \ldots, f_s)$. Since ideals are infinite objects, this can be viewed as an effective method of computing the ideal generated by $f_1, \ldots, f_s$. The discussion at the end of the last section shows that this is an interesting question even when $f = 1$.

Recall that in the simple case when $s = 1$ and f, f_1 are univariate, we can solve the ideal membership problem by simply diving f by f_1 to obtain a quotient q and a remainder r; we have that $f(x) = q(x) f_1(x) + r(x)$, and $f \in I(f_1)$ if and only if $r(x) = 0$. In the more complicated scenario where f, f_i are multivariate polynomials, we would like to take the same approach and divide f by $(f_1, \ldots, f_s)$ by expressing f as $f = \sum_{i=1}^{s} q_i f_i + r$ where r is a remainder term which we try to make as small as possible. If we can achieve $r = 0$ then f belongs to $I(f_1, \ldots, f_s)$.

However, in order to make this precise we need a criterion for deciding which of two possible remainders is "smaller." For instance, if $f = x^2$ and $f_1 = x^2 - y^2$, we could express f as $0 \cdot f_1 + (x^2)$, in which case the remainder would be x^2, or we could express f as $1 \cdot f_1 + (y^2)$, in which case the remainder would be y^2.

In the next section we define an ordering on monomials in $K[x_1, \ldots, x_n]$; this will enable us to meaningfully discuss the "size" of a remainder. Then we will give an algorithm for "dividing" f by several polynomials $f_1, \ldots, f_s$ and computing the remainder. As we will see, though, this division algorithm will not have all of the properties we want. Ultimately we will develop the machinery of Groebner bases to come up with an effective (but not efficient) algorithm to solve the ideal membership problem.

5 Orderings on Monomials

For $\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathbf{Z}^n_{\geq 0}$ we write x^α to denote the monomial $x_1^{\alpha_1} \ldots x_n^{\alpha_n}$.

Definition 5 *A relation $\succeq$ on the set of monomials in $K[x_1, \ldots, x_n]$ is an* admissible monomial ordering *if it satisfies the following properties:*

1. *$\succeq$ is a total ordering, i.e. for any $\alpha, \beta \in \mathbf{Z}^n_{\geq 0}$, either $\alpha \succ \beta, \alpha = \beta$, or $\alpha \prec \beta$.*
2. *for all $\alpha, \beta, \gamma \in \mathbf{Z}^n_{\geq 0}$, if $x^\alpha \succeq x^\beta$ then $x^{\alpha+\gamma} \succeq x^{\beta+\gamma}$.*
3. *for all $\alpha \in \mathbf{Z}^n_{\geq 0}$, it holds that $x^\alpha \succeq 1$.*

We will prove the following fact in Section 8

Fact 6 *Let $\succeq$ be an admissible monomial ordering. Then any sequence of monomials which is strictly decreasing under $\succeq$, i.e.*

$$x^{\alpha(1)} \succ x^{\alpha(2)} \succ x^{\alpha(3)} \succ \cdots$$

must be finite.

Here are some examples of admissible monomial orderings:

Example 1: lexcographic (lex) ordering. To compare x^α and x^β under lex ordering, let j be such that $\alpha_i = \beta_i$ for $0 \leq i < j$ and $\alpha_j \neq \beta_j$. Then if $\alpha_j > \beta_j$ we have $x^\alpha \succ x^\beta$, and if $\alpha_j < \beta_j$ we have $x^\alpha \prec x^\beta$.

We note that under lex ordering, the "earlier" variables $x_1, x_2, \ldots$ are more important than the "later" variables which have higher indices. As a result, lex ordering is useful in applications such as the following: let $f_1, \ldots, f_s \in K[x_1, \ldots, x_n]$ and suppose $m < n$. Consider the following statement:

$$\text{“}\exists x_1, \ldots, x_m \; f_1(x_1, \ldots, x_n) = \ldots = f_s(x_1, \ldots, x_n) = 0.\text{”}$$

The truth value of this statement will be a function of $x_{m+1}, \ldots, x_n$; we would like to find this function. As we will see later in the course, we will eliminate the variables $x_1, \ldots, x_m$ to obtain new functions $g_1(x_{m+1}, \ldots, x_n), \ldots, g_t(x_{m+1}, \ldots, x_n)$ which depend only on the variables $x_{m+1}, \ldots, x_n$. The lexicographic ordering with $x_i \succ x_j$ for $i > j$ will facilitate this elimination.

Example 2: graded lexicographic (grlex) ordering. We write $|\alpha|$ to denote $\sum_{i=1}^n \alpha_i$. Graded lexicographic ordering ranks monomials by their total degree and breaks ties by using lexicographic ordering; so if $|\alpha| > |\beta|$ then $x^\alpha \succ x^\beta$. Otherwise, if $|\alpha| = |\beta|$, the ranking of x^α and x^β is determined by the lexicographic ordering of x^α and x^β.

Revisiting our earlier example where $f = x^2$ and $f_1 = x^2 - y^2$, we see that under the lexicographic ordering on monomials over $\{x, y\}$ with $x \succ y$, the remainder y^2 is smaller than x^2. In the next section we will further refine our intuitive notion of what constitutes a good remainder.

6 Division by several polynomials

Example 3: Let $f = x^2y + x + y^2$, let $f_1 = x^2$, let $f_2 = y^2 - 3$, and let $\succeq$ denote lex order on monomials over $\{x, y\}$ with $x \succ y$. One attempt to divide f by f_1 and f_2 would be to express f as $y \cdot f_1 + (x + y^2)$, leaving a remainder of $x + y^2$. Since any remainder must contain x, this expression has succeeded in minimizing the highest-degree term; yet this does not seem like the best possible remainder because it still contains a y^2 term. A better remainder can be obtained by writing $f = y \cdot f_1 + 1 \cdot f_2 + (x + 3)$. This remainder is such that no term in the remainder is divisible by the leading term of any polynomial which we are dividing by.

Example 4: Let $f = x^2y + xy^2 + y^2$, let $f_1 = xy - 1$, and let $f_2 = y^2 - 1$. If we first extract the largest possible multiple of f_1 and then extract the largest possible multiple of f_2 from what's left, we obtain $f = (x+y) \cdot f_1 + 1 \cdot f_2 + (x+y+1)$, yielding a remainder of $x+y+1$. On the other hand, if we were to first extract the largest multiple of f_2 and then f_1, we would get $f = (x+1) \cdot f_2 + x \cdot f_1 + (2x+1)$ for a remainder of $2x + 1$. Yet another remainder can be obtained by writing $f = (2y + x) \cdot f_1 + (-x + 1) \cdot f_2 + (2y + 1)$.

In order to state our division algorithm, we require the following definitions. Let $\succeq$ be an admissible monomial ordering over the variables $x_1, \ldots, x_n$ and let $f \in K[x_1, \ldots, x_n]$.

- The *leading monomial* of f, written $LM(f)$, is the monomial which is ranked highest under $\succeq$ of all monomials which have nonzero coefficients in f.
- The *leading coefficient* of f, written $LC(f)$, is the coefficient of $LM(f)$ in f.
- The *leading term* of f, written $LT(f)$, is $LC(f) \cdot LM(f)$. Note that if f is a nonconstant polynomial, then $LM(f) \succ LM(f - LT(f))$.
- The *degree* of f, written $\deg(f)$, is α such that $x^\alpha = LM(f)$.

The algorithm to divide a polynomial f by an ordered s-tuple of polynomials $f_1, \ldots, f_s$ is as follows. The initial invocation should be $\textsc{Divide}(f; 0; f_1, \ldots, f_s; 0, \ldots, 0)$, and the return value will be $(r; q_1, \ldots, q_s)$ such that $f = q_1 f_1 + \cdots + q_s f_s + r$.

$\texttt{Divide}(f; r; f_1, \ldots, f_s; q_1, \ldots, q_s)$

1. If $f = 0$ then return $(r; q_1, \ldots, q_s)$.
2. Let i be least index such that $LM(f_i)$ divides $LM(f)$.
3. If such an i exists, let $q'_i = LT(f)/LT(f_i)$ and return

$$\texttt{Divide}(f - q_i f_i; r; f_1, \ldots, f_s; q_1, \ldots, q_{i-1}, q_i + q'_i, q_{i+1}, \ldots, q_s).$$

4. If no such i exists, return $\texttt{Divide}(f - LT(f), r + LT(f), f_1, \ldots, f_s, q_1, \ldots, q_s)$.

We can immediately make several observations about the algorithm. Perhaps the most important properties of the algorithm, both of which can be easily verified, are the following: if the algorithm outputs $(r; q_1, \ldots, q_s)$, then the identity $f = q_1 f_1 + \cdots + q_s f_s + r$ must indeed hold, and the remainder polynomial r has the property that no monomial term of r is divisible by any of $LT(f_1), \ldots, LT(f_s)$. We say that a remainder which satisfies this property is *legitimate.*

Another important property is that the algorithm must eventually terminate on any input. To see this, note that the degree of f drops with each invocation of `Divide` within the algorithm. Thus, the sequence of leading monomials of f which successively occur as the algorithm progresses form a strictly decreasing sequence under $\succeq$, and by Fact 6 such a sequence can be only finitely long.

Unfortunately, we cannot say that this algorithm runs in polynomial time in any meaningful sense. If each of $f, f_1, \ldots, f_s$ have at most t nonzero coefficients, the algorithm's runtime may exceed any fixed polynomial in s and t. If d_i is the maximum degree in each variable across all the polynomials, however, it can be shown that the run time will be $O(s \cdot \Pi_{i=1}^n (d_i + 1))$. This function grows extremely rapidly; even if each d_i is 1 it will be exponential in n.

Another less-than-desirable property of the algorithm is that it is sensitive to the order of its inputs $f_1, \ldots, f_s$; hence using it to divide f by (f_1, f_2, f_3) can produce different results than dividing f by (f_2, f_1, f_3). Thus, it is possible for two different polynomials r_1, r_2 to both be legitimate remainders.

If in some sense we could divide the polynomial f by the *infinite* collection of polynomials $I(f_1, \ldots, f_s)$, though, then there would be a unique legitimate remainder. To see this, suppose that $f = g_1 + r_1 = g_2 + r_2$ where $g_1, g_2 \in I(f_1, \ldots, f_s)$ and r_1, r_2 are both legitimate remainders, i.e. no term of r_i is divisible by $LT(h)$ for any $h \in I(f_1, \ldots, f_s)$. It follows that no term of $r_2 - r_1$ can be divisible by $LT(h)$ for any h in $I(f_1, \ldots, f_s)$; but since $r_2 - r_1 = g_1 - g_2 \in I(f_1, \ldots, f_s)$, this means that $LT(r_2 - r_1)$ is not divisible by $LT(r_2 - r_1)$, which is a contradiction.

7 Groebner bases

Of course, our division algorithm cannot be applied to divide f by an infinite collection of polynomials. We might hope, though, to find a finite collection $g_1, \ldots, g_t$ of polynomials which would be such that using our division algorithm to divide f by $(g_1, \ldots, g_t)$ would somehow give us the remainder that we would like to obtain if we could divide f by $I(f_1, \ldots, f_s)$. It is clear that the leading terms $LT(g_1, \ldots, g_t)$ of such a collection might have to contain more information than the leading terms of the polynomials $f_1, \ldots, f_s$. For example, under lex ordering with $x \succ y \succ z$, the leading terms of the polynomials $f_1 = x^2 - y$ and $f_2 = x^2 - z$ yield no information whatsoever about y and z.

With this motivation, we make the following important definition:

Definition 7 *Under a fixed monomial ordering $\succeq$, a finite subset $(g_1, \ldots, g_t)$ of an ideal J forms a* Groebner basis *for J if*

1. *$I(g_1, \ldots, g_t) = J$ and*
2. *$I(LT(g_1), \ldots, LT(g_t)) = I(LT(J))$, where $LT(J) = \{LT(h) : h \in J\}$.*

Returning to the above example, we can immediately see that $\{f_1, f_2\}$ does not constitute a Groebner basis for $I(f_1, f_2)$, since $y \in I(LT(I(f_1, f_2)))$ but $y \notin I(x^2, x^2)$. Later in the course we will prove that given any ideal $J \subset k[x_1, \ldots, x_n]$ and any admissible monomial ordering $\succeq$, there is a Groebner basis for J under $\succeq$. As a corollary we will obtain Hilbert's basis theorem, which states that every ideal has a finite basis. We will also show that our division algorithm, when applied to a Groebner basis, yields an effective solution to the ideal membership problem. (Unfortunately, it will not be the case that our division algorithm applied to Groebner bases yields an efficient solution to the ideal membership problem, but there is strong evidence that no efficient algorithm exists for this problem – Mayr and Meyer have showed that the the ideal membership problem is EXPSPACE-hard.) We will also define a minimal Groebner basis and show that the minimal Groebner basis for an ideal is unique.

8 Monomial Ideals

We turn first to the study of monomial ideals. In this section we prove that every monomial ideal is finitely generated; this is a special case of Hilbert's basis theorem.

Definition 8 *An ideal $J \subseteq K[x_1, \ldots, x_n]$ is a* monomial ideal *if there is some set A (possibly infinite) of monomials such that*

$$J = \{ \sum_{x^\alpha \in A} q_\alpha x^\alpha \mid q_\alpha \in K[x_1, \ldots, x_n]\}.$$

We can immediately state the following claim:

Claim 9 *Let $J = I(x^\alpha : x^\alpha \in A)$ be a monomial ideal. A monomial x^β belongs to J iff x^β is divisible by x^α for some $x^\alpha \in A$.*

Proof Clearly if $x^\beta = x^\alpha \cdot x^\gamma$ for some $x^\alpha \in A$, then $x^\beta \in J$. For the other direction, suppose that $x^\beta \in J$; then $x^\beta = q_{\alpha_1} x^{\alpha(1)} + \cdots + q_{\alpha(s)} x^{\alpha(s)}$ for some $x^{\alpha(1)}, \ldots, x^{\alpha(s)} \in A$. The right hand side is a sum of monomials each of which is divisible by some element of A, and the left hand side must have the same property. ■

As an example, consider the ideal $I(x^2y^5, x^4y^3, x^5y) \subset K[x, y]$. The monomials in I correspond to the integer grid points which fall in the shaded region of Figure 1.

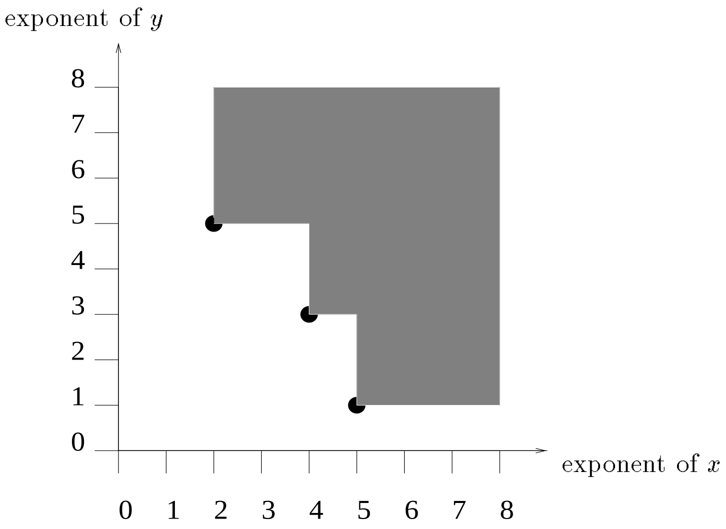

Figure 1

Now we prove that every monomial ideal has a finite basis:

Lemma 10 (Dickson's Lemma) *Let $J = I(x^\alpha : x^\alpha \in A) \subseteq K[x_1, \ldots, x_n]$. Then there is some finite set of monomials $\{x^{\alpha(1)}, \ldots, x^{\alpha(s)}\} \subseteq A$ such that $J = I(x^{\alpha(1)}, \ldots, x^{\alpha(s)})$.*

Proof The proof is by induction on n, the number of variables. If $n = 1$, then let $\beta = \min\{\alpha : x^\alpha \in A\}$; we have that x^β divides x^α for all $x^\alpha \in A$, and hence $J = I(x^\beta)$.

Now suppose the claim holds for $1, \ldots, n-1$, and let $J \subseteq K[x_1, \ldots, x_{n-1}, y]$ be a monomial ideal over n variables. We write a monomial in $K[x_1, \ldots, x_{n-1}, y]$ as $x^\alpha y^m$ where $\alpha \in \mathbf{Z}_{\geq 0}^{n-1}$ and $m \in \mathbf{Z}_{\geq 0}$. Let $J' \subseteq K[x_1, \ldots, x_{n-1}]$ be the monomial ideal generated by those monomials x^α which are such that $x^\alpha y^m \in J$ for some $m \geq 0$. Clearly J' is a monomial ideal; by the induction hypothesis, we have that J' is finitely generated, i.e. $J' = I(x^{\alpha(1)}, \ldots, x^{\alpha(s)})$.

Let m be the smallest integer such that $x^{\alpha(i)} y^m \in J$ for $i = 1, \ldots, s$. For each $k \in \{0, \ldots, k-1\}$ let $J'_k \subseteq K[x_1, \ldots, x_{n-1}]$ be the ideal generated by those monomials x^β which are such that $x^\beta y^k \in J$. Using the inductive hypothesis again, each J'_k is finitely generated, i.e. $J'_k = I(x^{\alpha_k(1)}, \ldots, x^{\alpha_k(s_k)})$.

Now consider the following list of monomials over $\{x_1, \ldots, x_{n-1}, y\}$:

from J' : $\quad x^{\alpha(1)}, \ldots, x^{\alpha(s)}$

$$\begin{array}{rcl} \text{from } J_0' & : & x^{\alpha_0(1)}, \ldots, x^{\alpha_0(s_0)} \\ & \vdots & \\ \text{from } J_k' & : & x^{\alpha_k(1)}, \ldots, x^{\alpha_k(s_k)} \\ & \vdots & \\ \text{from } J_{m-1}' & : & x^{\alpha_{m-1}(1)}, \ldots, x^{\alpha_{m-1}(s_{m-1})} \end{array}$$

Let A' be the (finite) set of all monomials which occur in this list; we claim that $J = I(A')$. Clearly $I(A') \subseteq J$, we must show that $J \subseteq I(A')$. Let $x^\alpha y^\ell$ be an arbitrary monomial which belongs to J. If $\ell \geq m$, then by our construction of J' we have that $x^\alpha y^\ell$ must be divisible by some $x^{\alpha(i)} y^m$. Otherwise, if $\ell < m$, then $x^\alpha y^\ell$ must be divisible by some $x_k^\alpha(i) y^\ell$ by the construction of J_ℓ'. By Claim 9, any monomial in J must also belong to $I(A')$, and it is easy to verify that this implies that $J \subseteq I(A')$, and hence $J = I(A')$. With a little additional work one can show that the finite set of monomials which generate J can actually be chosen from A (exercise for the reader). ■

As a corollary of Dickson's Lemma we can prove the following fact about admissible monomial orderings, which we used earlier to show that our division algorithm will always run in finite time.

Fact 6 *Let $\succeq$ be an admissible monomial ordering. Then any sequence of monomials which is strictly decreasing under $\succeq$, i.e.*

$$x^{\alpha(1)} \succ x^{\alpha(2)} \succ x^{\alpha(3)} \succ \cdots$$

must be finite.

Proof Let $x^{\alpha(1)} \succ x^{\alpha(2)} \succ x^{\alpha(3)} \succ \cdots$ be a sequence of monomials which is strictly decreasing under $\succeq$. Let $A = \{x^{\alpha(1)}, x^{\alpha(2)}, \ldots\}$ and let $J = I(x^\alpha : x^\alpha \in A)$. Since J is a monomial ideal, Dickson's Lemma implies that for some $x^{\alpha(i_1)} \succ \cdots \succ x^{\alpha(i_s)} \in A$, we have that $J = I(x^{\alpha(i_1)}, \ldots, x^{\alpha(i_s)})$. We claim now that $x^{\alpha(i_s)}$ must be the smallest element of A, and thus that A must be finite. To see this, let x^α be any element of A. Since $x^\alpha \in J = I(x^{\alpha(i_1)}, \ldots, x^{\alpha(i_s)})$, it follows from Claim 9 that $x^\alpha = x^{\alpha(i_j)} \cdot x^\gamma$ for some $j \in \{1, \ldots, s\}$ and some $\gamma \in \mathbf{Z}^n_{\geq 0}$. Since $\succeq$ is an admissible monomial ordering, though, we have that $x^\alpha = x^{\alpha(i_j)} \cdot x^\gamma \succeq x^{\alpha(i_j)} \succeq x^{\alpha(i_s)}$ and hence $x^{\alpha(i_s)}$ is indeed the smallest element of A. ■