1 Introduction

These notes contain a bunch of problems built around exploring what might be one of the simplest-sounding algebraic objects imaginable: systems of polynomial equations. There are many special cases, like solving a single quadratic equation in one variable or a system of linear equations, which are easy enough to be answered in most high school algebra classes. When all the equations are linear, you can usually learn everything you need to know by just solving the equations, writing each variable in terms of the others until you run out. But as we'll see, the general situation can get vastly more complicated.

As a first example, we'll look at a theorem in Euclidean geometry that can be solved using the techniques you'll learn by solving the problems in this packet. Consider the diagram in Figure 1.

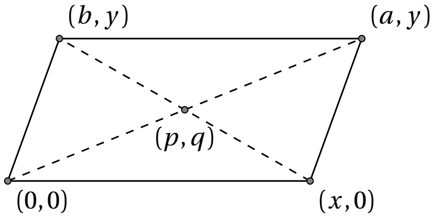

Figure 1: The diagram used in Problem 1.

If you've seen much Euclidean geometry, you know that the point labeled (p,q) is the midpoint of two diagonals. There are, of course, relatively simple proofs of this fact using coordinate-free methods, but what might be surprising is that we can prove it just by thinking about polynomial equations with almost no cleverness once we have the right tools.

Problem 1. We'll prove the theorem in a couple steps.

(a) Assuming that the diagram above actually depicts a parallelogram, show that the following equations are true:

- $a-x=b$
- $aq=yp$
- $(y-q)(b-x)=y(b-p)$

[Hint: these equations are expressing something about slopes of lines.]

(b) Suppose I tell you that

$$4(p^2+q^2)-(a^2+y^2)=\frac{1}{xy}(F\cdot(a-x-b)+G\cdot(aq-yp)+H\cdot((y-q)(b-x)-y(b-p))),$$

where

$$F=-xya-2xyp-y^2q+xaq-abq+2xpq-2bpq-2yq^2,$$
$$G=y^2-xa+ab-2xp+2bp+2yq,$$
$$H=y^2+xa+ab+2xp+2bp+2yq.$$

Conclude that (p,q) is the midpoint of the diagonal that goes from $(0,0)$ to (a,y). Argue by symmetry that it is also the midpoint of the other diagonal. [Note: you can verify that equation if it would make you feel better, but I'm certainly not asking you to!]

You might object that, even if the final answer only relied on adding and multiplying a bunch of polynomials, it must have taken a great deal of cleverness to come up with F, G, and H in the first place — they seem to have come out of nowhere. In fact, not only did it not involve any cleverness at all, but I found F, G, and H on a computer: there is an algorithm that takes the polynomials in this problem and either produces such an F, G, and H or tells you that they don't exist. The point of this packet is to understand how this works.

These problems were originally written for a one-week class at Canada/USA Mathcamp 2016. If you're reading them in that context, you should know that this is an "inquiry-based" class, which for us means that you'll be expected to work on these problems outside the classroom and take turns presenting the solutions at the board. I strongly encourage you to talk to each other and to me, both during TAU and otherwise. We might not get through this whole packet; if we don't and you'd like to keep working on it, I'm happy to keep talking to you about it.

If you're not reading these notes at Mathcamp, then I'm still happy to answer any questions you might have if you send me an e-mail at `njmford@gmail.com`.

Many problems from this packet are taken from the excellent *Ideals, Varieties, and Algorithms* by David Cox, John Little, and Donal O'Shea. The whole book should be accessible to anyone who can understand the problems in this packet, and it's definitely worth reading. We'll only be covering a tiny fraction of the material from that book here.

2 Ideals

First, let's explore a bit more about the structure of systems of polynomial equations.

Definition. A set of polynomials I is called an **ideal** if the following conditions are satisfied:

- $0 \in I$
- If $f, g \in I$, then $f + g \in I$.
- If $f \in I$ and p is any polynomial at all, then $pf \in I$.

  Given some finite set of polynomials $f_1, \dots, f_n$, we can form an ideal by multiplying the f_i's by other polynomials and adding them together in all possible ways. That is, the set of all polynomials of the form
  $$p_1 f_1 + p_2 f_2 + \cdots + p_n f_n.$$
  We call this the **ideal generated by** the f_i's, and we'll write it $(f_1, \dots, f_n)$.

Problem 2. Verify that $(f_1, \dots, f_n)$ is indeed an ideal.

Problem 3. Describe the following ideals as simply as you can.

(a) (0)

(b) (1)

(c) (7)

(d) (x), for polynomials of one variable.

(e) $(x-1)$, for polynomials of one variable.

(f) (x, y), for polynomials of two variables.

(g) (x, x^2), for polynomials of one variable.

These definitions will probably seem less arbitrary after the following problem:

Problem 4. Suppose we have a system of n polynomial equations in k variables, say

$$f_1(x_1, \dots, x_k) = 0$$

$$f_2(x_1, \dots, x_k) = 0$$

$$\dots$$

$$f_n(x_1, \dots, x_k) = 0.$$

Show that some point $(a_1, \dots, a_k) \in \mathbb{C}^k$ is a solution of this system of equations if and only if $g(a_1, \dots, a_k) = 0$ for every g in the ideal $(f_1, \dots, f_n)$.

Problem 4 is, in some sense, exactly the reason for introducing the concept of ideals: the ideal generated by a collection of polynomials captures all the information about the set of solutions of the corresponding equations. As we'll see later, you can learn a lot about the solutions to a set of polynomial equations by passing to a different set of polynomials that generate the same ideal.

Problem 5. Show that the following three sets of polynomials all generate the same ideal:

- $(x, y, x^2 + y^2)$
- (x, y)
- $(x + y, x - y)$

Verify the result from 4 in this case.

Problem 6. Find two systems of polynomial equations which have the same set of solutions as each other, but which don't generate the same ideal, showing that the converse of the statement from Problem 4 is false. [Hint: it's possible to do this with two different ideals which are each generated by a single polynomial in one variable.]

3 Polynomials of One Variable

As a warm-up, we'll start with a situation you might have seen in a high school algebra class, even if you haven't seen all the proofs: the case of polynomials in one variable. The basic tool is a result called the *Division Theorem*:

Theorem. *Let g be a nonzero polynomial in one variable. Then for every other polynomial f, there is a unique q and r (called the* quotient *and* remainder *respectively) so that $f = gq + r$ and* $\deg r < \deg g$.

Note the similarity to what happens when dividing integers: here the condition on degrees is playing the same role as the fact that the remainder has to be smaller than the number you were dividing by.

Problem 7. Find q and r in the following cases (no need to prove uniqueness). [Hint: in some cases, it is helpful to build up q slowly by finding multiples of g that cancel the terms of f one at a time.]

(a) $g = x$; $f = 5x^3 - 3x^2 + 7x + 2$

(b) $g = x^2 + 1$; $f = 2x^4 + 3x^3 - x^2 - x - 1$

(c) $g = 5x^5 + 3x + 1$; $f = 2x^2 - 7$

(d) $g = 1$; $f = x^2 + 3x + 6$

Problem 8. Now we'll prove the Division Theorem in general. We'll do it by formalizing the procedure you probably used in part (b) above.

(a) If $\deg f < \deg g$, what should q and r be?

(b) If $\deg f \geq \deg g$, show that there's some a and d so that $\deg(f - g \cdot ax^d) < \deg f$.

(c) Using the previous two parts, show using induction on $\deg f$ that a q and r satisfying the requirements of the Division Theorem have to exist.

(d) Suppose we had two different quotients and remainders for f, that is, some q, q', r, r' with $\deg r, \deg r' < \deg g$ and $f = gq + r = gq' + r'$. Show that $q = q'$ and $r = r'$, and therefore the result of dividing polynomials is always unique, completing the proof. [Hint: the condition on the degrees of r and r' is crucial.]

Ideals of one-variable polynomials have a very simple structure:

Problem 9. Show that if I is an ideal of one-variable polynomials, then there's some g for which $I = (g)$, that is, I consists of exactly all the multiples of g. [Hint: take g to be any element of I of smallest possible degree, and use the Division Theorem.]

Problem 10. The Division Theorem lets you prove several facts about polynomials you probably recognize from your algebra class:

(a) Use the Division Theorem to show that for a one-variable polynomial f and a number a, we have $f(a) = 0$ if and only if f is a multiple of $x - a$.

(b) Suppose $f(a) = 0$. We say that a is a **root of f with multiplicity** m if f is a multiple of $(x-a)^m$, and m is the largest number for which this is true. Use the Division Theorem to show that f has at most $\deg f$ roots, counted with multiplicity.

You may be familiar with the Fundamental Theorem of Algebra, which says that every polynomial with coefficients in $\mathbb{C}$ has a root in $\mathbb{C}$. Sadly, we won't be proving it here. But if we assume it, we can get a stronger version of Problem 10:

Problem 11. Use the Fundamental Theorem of Algebra to show that any nonzero one-variable polynomial f with coefficients in $\mathbb{C}$ has exactly $\deg f$ roots in $\mathbb{C}$, counted with multiplicity.

So suppose we have a system of polynomial equations in one variable. Problem 9 guarantees that it's actually just as good to solve a single equation $g(x) = 0$. But how do we find g?

Problem 12. In this problem we'll find an explicit algorithm to produce the single generator of an ideal of polynomials in one variable.

(a) Suppose we have a (finite) system of polynomial equations in one variable with at least two distinct polynomials, say f and g, with $\deg f \geq \deg g$. Show that if we replace f by its remainder under division by g, the resulting set of polynomials generates the same ideal as the original set.

Turn this observation into an algorithm for finding the single polynomial that generates the ideal in question. (The algorithm you end up with is called the **Euclidean algorithm**; the same procedure works for finding the greatest common divisor of a set of integers, which is the context in which it got its name.)

(b) Find the single polynomial that generates the ideal (x^6-1, x^4-1).

$\boxed{4}$ Dividing Multivariable Polynomials

Now we'll start trying to apply the ideas that worked so well before to the case of polynomials in more than one variable. The situation here is considerably more complicated — the results won't be as simple as what we were able to extract from the Division Theorem before. But, using a new tool called a Gröbner basis, we'll be able to find surprisingly explicit answers to both of these questions.

We'll tackle the first of these questions first: while it's arguably the less exciting of the two, it provides an excellent setting in which to develop the tools we'll be using for the rest of the problems in this packet.

Specifically, we'll address the following question: given a bunch of polynomial equations $g_1(x_1,\ldots,x_k)=0,\ldots,g_n(x_1,\ldots,x_k)=0$ and some other polynomial f in k variables, how can we tell whether f is in the ideal $(g_1,\ldots,g_n)$?

Problem 13. Convince yourself that this question actually has something to do with solving polynomial equations: Suppose f is in the ideal $(g_1,\ldots,g_n)$. What does this tell us about the relationship between the set of solutions of the original equations and the set of solutions of $f(x_1,\ldots,x_k)=0$?

We've in fact already answered this question in the one-variable case:

Problem 14. Suppose we're working with just one variable. Use the Division Theorem and the other tools from the previous section to describe an algorithm for answering the ideal membership question in this case. [Hint: belonging to an ideal that's generated by a single element g is the same as being a multiple of g. How can you use the tools from the last section to tell if a given polynomial f is a multiple of g?]

This suggests that we'd like a result like the Division Theorem for polynomials of several variables. This is what we'll explore over the next few problems.

Problem 15. Find an ideal that can't be generated by a single polynomial. [Hint: there are examples with just two variables.]

Given this fact, a division-based answer to the ideal membership question can't involve dividing by just one polynomial. We'd like an analogue of the Division Theorem that lets us "divide" by several polynomials at once. That is, given some nonzero polynomials $g_1,\ldots,g_k$ as above, along with any other polynomial f, we'd like some way to write

$$f = q_1g_1+\cdots+q_kg_k+r,$$

with some nice condition on the remainder analogous to the one we had on its degree in the one-variable case.

Problem 16. Although it's not necessary to follow the rest of the packet (since I don't know how to write the rest of the problems without spoiling this question for you) I encourage you at this point to stop reading for a bit and think about what will be necessary for such a division procedure to work.

More specifically, recall that the one-variable version of the algorithm involved repeatedly canceling the leading term of f by multiplying g by a monomial, that is, by some ax^d. How should this work now?

If you're reading this at Mathcamp, come talk to me about this question at TAU when you reach this point in the packet.

We'll now start the work of ruining Problem 16 for you. As we mentioned before, a key ingredient in the one-variable division algorithm was the notion of a "leading term," which was the term of f we were always interested in canceling. In several variables, it's not always even clear which term is the leading term, so we'll have to make a choice:

Definition. A **monomial** is a polynomial that you get by multiplying several of the variables together, so any polynomial is a sum of monomials multiplied by different complex numbers. (We'll use the words "monomial" and "term" differently: a term might have any coefficient, but a monomial has a coefficient of 1.)

We'll often abbreviate a monomial like $x_1^2 x_2^4 x_3^5$ with a symbol like x^α where (in this case) α is the triple $(2,4,5)$. We call α the **multidegree** of the monomial. Given a multidegree, the sum of its entries is called the **degree** of the monomial, and we'll write it $|\alpha|$. For example, $\deg(x_1^2 x_2^4 x_3^5) = |(2,4,5)| = 11$.

We'll sometimes talk about the sum of two multidegrees, by which we'll mean the sum of each component. For example, $(2,4,5)+(1,0,1)=(3,4,6)$. Note that this definition matches up with multiplication of monomials, that is, $x^\alpha \cdot x^\beta = x^{\alpha+\beta}$.

We define an ordering on multidegrees as follows: we'll say a multidegree α is **lexicographically larger** than β, and write $\alpha > \beta$, if in the first position in which they differ α has a larger entry. For example, $(1,3,7) > (1,2,8)$, since they first differ in the second position and $3 > 2$.

We'll often think of lexicographic as giving us an order on monomials as well, and we'll write $\alpha > \beta$ and $x^\alpha > x^\beta$ interchangeably.

The **leading term** of a nonzero polynomial is then just the term with the largest monomial in it. (We think of the zero polynomial as having no terms, and therefore no leading term.) If f is a nonzero polynomial, the **multidegree of** f, written multideg f, is the multidegree of its leading term.

Note. There are lots of other ways to order monomials besides lexicographically and they're useful for different things. We're making the choice to stick with just the one ordering only for simplicity, but the theory we develop here works equally well for other orderings.

Problem 17. Prove the following facts about lexicographic order:

(a) If α,β,γ are multidegrees and $\alpha < \beta$, then $\alpha+\gamma < \beta+\gamma$.

(b) If f and g are two nonzero polynomials, then $\operatorname{multideg}(fg) = \operatorname{multideg} f + \operatorname{multideg} g$.

(c) If furthermore $f+g \neq 0$, then $\operatorname{multideg}(f+g) \leq \max(\operatorname{multideg} f, \operatorname{multideg} g)$, with equality if $\operatorname{multideg} f \neq \operatorname{multideg} g$.

Finally, we're ready to state our beefed-up Division Theorem.

Theorem (Multivariable Division Theorem, Part 1)**.** *Let $g_1,\ldots,g_n$ be polynomials in k variables. Then any other k-variable polynomial f can be written in the form*

$$f = q_1 g_1 + \cdots + q_n g_n + r,$$

where none of the terms of r is divisible by the leading term of any of the g_i's. Also, whenever some $q_i \neq 0$, we further have that $\operatorname{multideg} f \geq \operatorname{multideg}(q_i g_i)$.

We'll explore how the algorithm ought to work through a couple examples, which are taken from *Ideals, Varieties, and Algorithms.*

Problem 18. Divide $f = xy^2 + 1$ by $g_1 = xy + 1$ and $g_2 = y + 1$, that is, find a q_1, q_2, and r satisfying the conditions in the Multivariable Division Theorem. [Hint: follow an algorithm kind of like the the one-variable division algorithm, repeatedly canceling the leading term of f with the leading term of either g_1 or g_2, preferring g_1 whenever possible.]

Problem 19. Now divide $f = x^2y + xy^2 + y^2$ by $g_1 = xy + 1$ and $g_2 = y^2 - 1$. What happens in this example that didn't happen in the previous one?

Problem 20. Write down the algorithm you used to solve the previous two problems and prove that it satisfies the conditions in the theorem.

So it would seem that we can answer the ideal membership question the same way we did before. Sadly, this is not the case.

Problem 21. Let r_1 be the remainder you got in Problem 19. Now switch the roles of g_1 and g_2 and do the division again. Call this remainder r_2.

You should get a different answer, proving that the Multivariable Division Theorem doesn't give us the same uniqueness of the remainder we got from the ordinary division theorem.

Define $r = r_1 - r_2$. Is r in the ideal (g_1, g_2)? If not, explain why not. If so, find polynomials a and b so that $r = ag_1 + bg_2$.

What happens when you divide r by g_1 and g_2? Could you have guessed the answer before you did it? What does this say about using the Multivariable Division Theorem to answer the ideal membership question? What property could we ask g_1 and g_2 to have that would stop something like this from happening?

5 Gröbner Bases

In Problem 21, you saw that it's possible for a polynomial to belong to the ideal generated by the g_i's even if the remainder after dividing by them is nonzero. (If this isn't what you found, you should probably work through the problem again!)

There is hope, though. Since it's just as good to replace any system of equations with ones that generate the same ideal, we can try to find a set of generators that doesn't have the problems that g_1 and g_2 did in Problem 21. Here is the exact condition we'll want:

Definition. A set of generators $g_1, \ldots, g_k$ of an ideal I is called a **Gröbner basis** if the leading term of every element of I is divisible by the leading term of one of the g_i's.

The excellent news is that this is enough to fix our Division Theorem:

Theorem (Multivariable Division Theorem, Part 2)**.** *If $g_1, \ldots, g_k$ is a Gröbner basis for the ideal that they generate and f is any other polynomial, then there is a unique polynomial r with the following properties:*

- *None of the terms of r is divisible by the leading term of any g_i.*
- *$f = g + r$ for some g in the ideal $(g_1, \ldots, g_k)$.*

Problem 22. Prove this. [Hint: If r and r' both satisfy the properties in the theorem for the same f, what can you say about the leading term of $r - r'$?]

Problem 23. Prove that any polynomial f is a Gröbner basis for the ideal (f). [Hint: use Problem 17.] What does this tell us about the Multivariable Division Theorem when we're dividing by a single polynomial?

It turns out that every ideal has a Gröbner basis and there is an algorithm to construct one, but we'll delay the proof until after we've seen some fun examples. For now, I'll show you how to find Gröbner bases using the computer algebra system Sage.

(If you're reading this at Mathcamp, come talk to me at this point; I'll have a laptop set up for you to do this all on. If not, you can go to `cloud.sagemath.com` and set up an account, then in a new project, create a "new SageMath worksheet"; check SageMathCloud's documentation for more details.)

Problem 24. We'll compute a Gröbner basis to use in the example from Problem 1. Open Sage and run the following code (note the spelling of `groebner`):

```
R.<x,y,a,b,p,q,x_inv,y_inv> = PolynomialRing(QQ, order='lex');
I = ideal(a-x-b, a*q - y*p, (y-q)*(b-x) - y*(b-p),
          x*x_inv - 1, y*y_inv - 1);
I.groebner_basis()
```

(The variables `x_inv` and `y_inv` stand for the multiplicative inverses of x and y; we'll explore why this was necessary in a moment.) You should get the following Gröbner basis:

$$(x + b - 2p, y - 2q, a - 2p, bx^{-1} - 2px^{-1} + 1, qy^{-1} - \frac{1}{2}).$$

(a) Use this Gröbner basis and the multivariable division algorithm to solve Problem 1.

(b) We used `x_inv` and `y_inv` to force x and y to be nonzero. Why is this necessary? [Hint: our original equations forced (p,q) to be the intersection point of the two diagonals. What happens when x or y is 0?]

We can in fact skip doing the division ourselves and get Sage to tell us directly whether a given polynomial is in the ideal we care about. If you enter:

```
4*(p^2 + q^2) - (a^2 + y^2) in I
```

Sage will output `True`.

A surprisingly large amount of geometry can be done this way. Here are a couple more examples.

Problem 25. Consider the diagram in Figure 2.

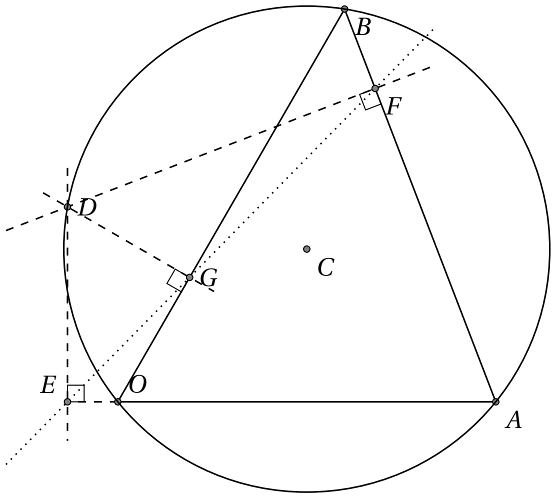

Figure 2: The diagram used in Problem 25.

Here C is the center of the circumcircle of triangle OAB (that is, the circle that goes through the points O, A, and B) and D is an arbitrary point on that circle. The points E, F, and G are formed by dropping perpendicular lines to the edges of the triangle. (These might meet the edge outside the triangle, as we see with E in the diagram.) We're going to prove that E, F, and G always lie on a line. (This is called the *Simson line.*)

(a) Let's begin by giving coordinates to the points. We might as well assume that $O=(0,0)$ and that A is on the x-axis, so $A=(a,0)$ for some a. Label the rest of the diagram with coordinates accordingly. [Hint: you know something about the coordinates of D and E.]

(b) Come up with equations that express the fact that G lies on OB and F lies on AB. Why am I not asking you to do this for E?

(c) Come up with equations that express the fact that DG is perpendicular to OB and DF is perpendicular to AB. Do you have to assert that DE is perpendicular to OA?

(d) Write equations that express the fact that O, A, B, and D lie on a circle through C.

(e) What quantities have to be nonzero here? Add in equations like the ones with `x_inv` and `y_inv` before to express this.

(f) Finally, write the conclusion we're looking for: that E, F, and G are collinear.

(g) Check in Sage that you did everything right. (You can either get a Gröbner basis for your ideal and perform the division yourself or just ask Sage whether if your answer to part (f) is in your ideal.)

Problem 26. Use this method to prove that in a right triangle ABC, if H is the point on BC which meets the altitude from A, then A and H lie on a circle together with the midpoints of the three edges. (See Figure 3. This is a special case the the *nine-point circle*; in a right triangle several of the points coincide, leaving only five distinct points.)

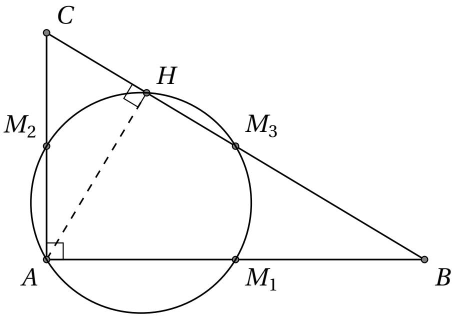

Figure 3: The diagram used in Problem 26.

Problem 27. Think of another fact in Euclidean geometry that seems like it could be proved using this method and prove it. Talk to me if you need ideas.

6 Finding Gröbner Bases

We'll conclude by describing an algorithm that produces Gröbner bases. While there are more efficient algorithms than the one we describe here, ours has the advantage of being very simple to understand.

We'll start by asking the question of how to tell whether a set of generators for an ideal is a Gröbner basis at all; this will end up being a critical part of the algorithm. Note that we don't actually currently know how to do this in a finite number of steps: the definition itself involves a condition that has to be satisfied by every element of the ideal.

Definition. Given two monomials x^α and x^β, we define their **least common multiple**, written $\operatorname{lcm}(x^\alpha, x^\beta)$, to be the smallest monomial divisible by both of them. It's computed by taking as the power of each variable x_i whichever of α_i or β_i is larger. For example, $\operatorname{lcm}(x_1^2x_2^4, x_1^3x_2^2) = x_1^3x_2^4$.

Given two polynomials f and g, write x^γ for the least common multiple of their leading terms. Then the **S-polynomial** of f and g is defined by

$$S(f,g) = \frac{x^\gamma}{\operatorname{lt} f} f - \frac{x^\gamma}{\operatorname{lt} g} g,$$

where $\operatorname{lt} f$ and $\operatorname{lt} g$ are the leading terms of f and g. Note, in particular, that x^γ is actually divisible by $\operatorname{lt} f$ and $\operatorname{lt} g$, so this division makes sense.

For example, let $f = x_1^2 + x_1 + x_2$ and $g = 2x_1x_2 + 5$. Then $x^\gamma = x_1^2x_2$, $\operatorname{lt} f = x_1^2$, and $\operatorname{lt} g = 2x_1x_2$, so $S(f,g) = x_1x_2 + x_2^2 - \frac{5}{2}x_1$. In other words, the S-polynomial is the simplest expression of the form $Af + Bg$ in which the leading terms cancel.

Problem 28. What does an S-polynomial look like when the two polynomials have the same multidegree?

Problem 29. Let's formalize the intuition that the S-polynomial is the simplest way to make leading terms cancel: suppose we have some polynomials $h_1, \ldots, h_m$ which all have the same multidegree, say δ. Suppose we have some expression of the form

$$p = a_1h_1 + \cdots + a_mh_m$$

where the a_i's are numbers and with $\operatorname{multideg} p < \delta$, so some cancellation happened among the leading terms. Show that p actually belongs to the ideal generated by the S-polynomials $S(h_i, h_j)$. [Hint: use Problem 28. The notation is simpler if you divide h_i by its leading coefficient, so that you can assume that that coefficient is 1.]

S-polynomials will be what give us our method for checking if something is a Gröbner basis:

Theorem (Buchberger's Criterion)**.** *The polynomials $g_1, \ldots, g_k$ form a Gröbner basis for the ideal they generate if and only if, for each S-polynomial $S(g_i, g_j)$, the remainder after diving by $g_1, \ldots, g_k$ is zero.*

Problem 30. We'll prove Buchberger's Criterion in several steps.

(a) First, prove that if the g's for a Gröbner basis, then the condition on S-polynomials is true.

(b) Now suppose we know the condition on S-polynomials. Pick any polynomial f in the ideal. Say we have some expression of the form

$$f = q_1 g_1 + \cdots + q_k g_k.$$

Recall from Problem 17 that this means that

$$\operatorname{multideg} f \leq \max(\operatorname{multideg}(q_i g_i)).$$

I claim that it's enough to show that for some way of writing f in this form, we actually have $\operatorname{multideg} f = \operatorname{multideg}(q_i g_i)$ for some i. Why?

(c) Suppose that this is not the case. That is, suppose that $\operatorname{multideg} f$ is strictly smaller than the largest $\operatorname{multideg}(q_i g_i)$. Call this largest multidegree δ. Split the sum into the terms with multidegree δ and those with smaller multidegree:

$$f = \sum_{\operatorname{multideg}(q_i g_i)=\delta} q_i g_i + \sum_{\operatorname{multideg}(q_i g_i)<\delta} q_i g_i.$$

Writing

$$p = \sum_{\operatorname{multideg}(q_i g_i)=\delta} \operatorname{lt}(q_i) g_i,$$

prove that the multidegree of p is itself less than δ.

(d) Apply Problem 29 to conclude that p is in the ideal generated by the S-polynomials $S(\operatorname{lt}(q_i)g_i, \operatorname{lt}(q_j)g_j)$.

(e) Prove that in fact p is in the ideal generated by the S-polynomials $S(g_i, g_j)$.

(f) Use the condition on dividing S-polynomials into the g_i's to show that there is an expression for p of the form

$$p = r_1 g_1 + \cdots + r_k g_k$$

where the largest $\operatorname{multideg}(r_i g_i)$ is smaller than δ.

(g) Use this to show there is a way to rewrite f like

$$f = \tilde{q}_1 g_1 + \cdots + \tilde{q}_k g_k,$$

where the largest $\operatorname{multideg} \tilde{q}_k g_k$ is smaller than δ. [Hint: what can you say about $f - p$ using the sum from part (c)?]

(h) Repeating this procedure, starting with such an expression for f, we either eventually end up with one where $\operatorname{multideg} f = \max(\operatorname{multideg}(q_i g_i))$ or we can continue forever, making $\max(\operatorname{multideg}(q_i g_i))$ smaller each time. Show that this second situation is actually impossible, and conclude that we've proved the theorem.

Given Buchberger's Criterion, the algorithm we'll use is very simple to describe:

Theorem (Buchberger's Algorithm)**.** *Start with some polynomials $g_1, \ldots, g_k$. If there's some $S(g_i, g_j)$ which has a nonzero remainder after dividing by the g's, add it to the list. If not, stop. At the end of this procedure, you have a Gröbner basis.*

Sadly, the tools we've developed here are not enough for a proof that this algorithm actually terminates. (There is a nice proof in *Ideals, Varieties, and Algorithms* which I encourage you to read; it relies on a fact called the Hilbert Basis Theorem which we haven't shown here.) But we do have the tools to prove the following:

Problem 31. Show that if Buchberger's Algorithm terminates, it has produced a Gröbner basis.

7 The Hilbert Basis Theorem

In this section, we'll prove a result called the Hilbert Basis Theorem. In addition to being interesting in its own right, it also allows us to fill the hole in the last section and show that Buchberger's Algorithm terminates.

There is a distinction that we haven't been very careful about up to this point that it's important to pay attention to now. When we introduced ideals, there were two similar definitions: what it means for a set of polynomials to be an ideal, and what it means to take the ideal generated by some polynomials $g_1, \dots, g_k$. You showed in Problem 2 that $(g_1, \dots, g_k)$ satisfies the definition of an ideal, but we never established the other implication, that every ideal is generated by a finite set of polynomials. This is what the Hilbert Basis Theorem is about.

Theorem (Hilbert Basis Theorem)**.** *For any ideal I of polynomials, there are finitely many polynomials $g_1, \dots, g_k$ so that $I = (g_1, \dots, g_k)$.*

By the end of this section, we'll have a proof of the Hilbert Basis Theorem. As a first step, we'll consider a special case which will be useful in the general proof.

Definition. A *monomial ideal* is an ideal generated by a (possibly infinite) set of monomials. (When we say that an ideal is generated by an infinite set S of polynomials, we mean that every element is of the form $p_1 h_1 + \cdots + p_k h_k$ for some $h_1, \dots, h_k \in S$. In particular, we don't try to take an infinite sum of polynomials.)

Problem 32. Prove that if I is a monomial ideal generated by a set S of monomials, then a monomial x^β belongs to I if and only if it is divisible by some monomial $x^\alpha \in S$.

Problem 32 tells us what all the monomials belonging to a monomial ideal look like. Sometimes it's helpful to represent the monomials belonging to a monomial ideal with a picture like the one in Figure 4. Every point represents a monomial — for example, $(6,2)$ stands for the monomial $x^6 y^2$ — and the points in the shaded region correspond to the monomials that belong to the ideal.

In fact, we have a similar description of every element of I:

Problem 33. Prove that if I is a monomial ideal generated by a set S of monomials, any element $f \in I$ can be written in the form

$$f = \lambda_1 x^{\beta_1} + \lambda_2 x^{\beta_2} + \cdots + \lambda_k x^{\beta_k},$$

where each λ_i is a number (not a polynomial) and each x^{β_i} is a monomial which is a multiple of one of the monomials in S.

Problem 34. Now we'll prove that every monomial ideal is generated by a finite set of monomials. Specifically, suppose I is a monomial ideal generated by a set S of monomials. We'll show that there is a finite subset of S which also generates I.

(a) Show that if I is an ideal of one-variable polynomials, then I is actually generated by a single element of S. (Note that it's not enough just to cite the fact that ideals of one-variable polynomials are generated by a single polynomial; we specifically want the generator to be one of the monomials in S!)

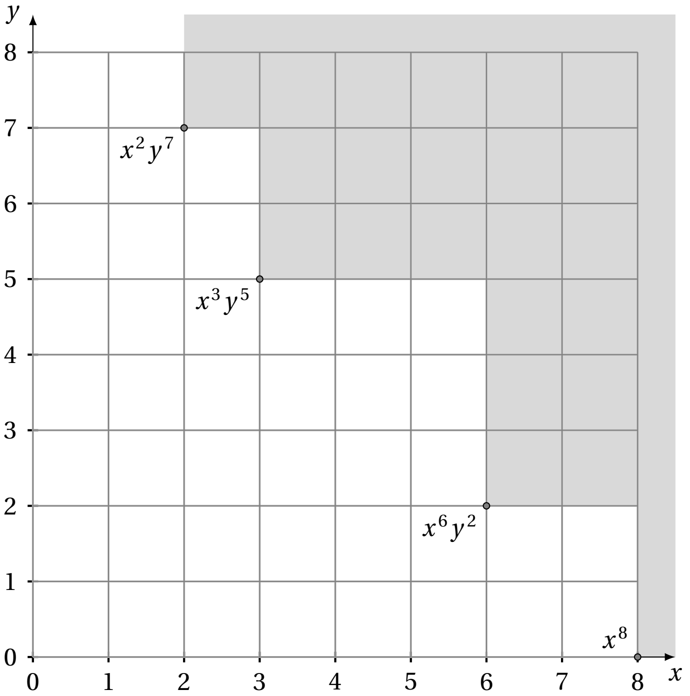

Figure 4: A graphical representation of the monomials in the monomial ideal generated by x^2y^7, x^3y^5, x^6y^2, and x^8.

(b) We'll go from here by induction on the number of variables. Suppose we have the claim for polynomials of $n-1$ variables, and we want to show it for polynomials of n variables.

Write S' for the set of monomials x^α in the variables $x_1, \dots, x_{n-1}$ with the property that $x^\alpha x_n^d \in I$ for some $d \geq 0$, and write J for the ideal of $(n-1)$-variable polynomials generated by the monomials in S'. Using the inductive hypothesis, we know that J is generated by finitely many of the monomials in S'. Call these $x^{\alpha_1}, x^{\alpha_2}, \dots, x^{\alpha_s}$.

Show that there is some m for which each $x^{\alpha_i} x_n^m \in I$. What is J for the ideal in Figure 4. What is m?

(c) For any integer j with $0 \leq j \leq m$, we'll define S'_j to be the set of those monomials x^α in S' for which $x^\alpha x_n^j \in I$, and we'll write J_j for the ideal generated by S'_j. (With this notation, the previous part is saying that $J = J_m$.) These ideals J_j are again generated by finitely many monomials in S'_j. Call these generators $x^{\alpha_{j1}}, \dots, x^{\alpha_{js_j}}$.

What are the J_j's for the ideal in Figure 4?

Show that every monomial in I is divisible by some monomial of the form $x^{\alpha_{jk}} x_n^j$ where $x^{\alpha_{jk}}$ is one of the generators of J_j constructed above. [Hint: take some monomial in I and look at the power of x_n. What can you say if this power is equal to some j with $0 \leq j \leq m$? What if it's more than m?]

(d) Use Problem 32 and Problem 33 to conclude from this that the monomials constructed in the previous part generate I.

(e) Finally, complete the proof of the theorem by showing that we can restrict our generating set to only consist of monomials that come from S.

This tells us that the Hilbert Basis Theorem is true for monomial ideals. In fact, that was the hardest part; we can use this to get the theorem in general.

Definition. Let I be an ideal. The *ideal of leading terms of I*, written $\mathrm{lt}(I)$, is the ideal generated by the leading terms of all the elements of I. Note that $\mathrm{lt}(I)$ is a monomial ideal.

Problem 35. Now we'll finish the proof of Hilbert's Basis Theorem.

(a) Using Problem 34, prove that $\mathrm{lt}(I)$ is can be generated by finitely many monomials, each of which is a leading term of some element of I.

(b) Call these elements $g_1, \ldots, g_k$. (That is, the g_i's are elements of I whose leading terms generate $\mathrm{lt}(I)$.) Suppose we have some $f \in I$. Prove that dividing f by the g_i's gives a remainder of 0, and that therefore the g_i's generate I, proving the Hilbert Basis Theorem.

Problem 36. Suppose we're given an *ascending chain* of ideals, that is, an infinite chain of ideals where each is contained in the next:

$$I_1 \subseteq I_2 \subseteq I_3 \subseteq \cdots.$$

Prove that there is some n so that all the ideals from I_n forward are equal. [Hint: prove that the union $\bigcup_{i=1}^{\infty} I_i$ is an ideal, and then use the Hilbert Basis Theorem.]

Problem 37. Show that if Buchberger's Algorithm didn't terminate, we would be able to construct an infinite ascending chain of ideals that never stabilizes in this way, contradicting Problem 36.