6.966 Algebra and Computation		November 9, 1998
	Lecture 18	
Lecturer: Madhu Sudan		*Scribe: Salil Vadhan*

Bézout's Theorem and Applications

In this lecture, we will prove Bézout's theorem for curves in the plane and describe a couple of its applications in coding theory.

Clarification on Dimension and Degree

In the last lecture, we defined the dimension and degree of ideals and varieties. There was an error in the definition we gave last time which we correct now.

Definition 1 *Let F be a field and I an ideal in $F[x_1, \ldots, x_n]$. Then the* Hilbert polynomial *of I is the function $\mathrm{HP}_I : \mathbb{N} \to \mathbb{N}$ given by $\mathrm{HP}_I(d)$ = the number of monomials of degree at most d* **not** *in $\mathrm{LT}(I)$.*

Fact 2 *For every I, HP_I is a eventually a polynomial, that is, there is a polynomial $p(x) \in \mathbb{Q}[x]$ such that for sufficiently large $d \in \mathbb{N}$, $\mathrm{HP}_I(d) = p(d)$.*

Definition 3 *The* dimension *of I, denoted $\dim(I)$, is the degree of HP_I. The* degree *of I, denoted $\deg(I)$, is the leading coefficient of HP_I times $(\dim(I)!)$. The* dimension *(resp.,* degree*) of a variety $V \subset F^n$ is defined to be the dimension (resp., degree) of $I(V)$.*

To motivate these definitions, it is helpful to think of monomial ideals, because $\mathrm{LT}(I)$ is equal to the set of monomials in I. One can easily verify the following examples:

Example 4 (union of two axes in F^2) *$I = (x^2y) \subset F[x, y]$. The monomials outside of I are exactly those that have degree 1 or 0 in x or have degree 0 in y. The number of these of degree at most d is easily seen to be $3d$, so $\dim(I) = 1$ and $\deg(I) = 3 \cdot 1! = 3$. Note that $V(I)$ is is the union of the two coordinate axes, so dimension 1 agrees with to our geometric intuition.*

Example 5 (the origin in F^2) *$I = (x, y^4) \subset F[x, y]$. The only monomials outside of $\mathrm{LT}(I)$ are $\{1, y, y^2, y^3\}$ so $\mathrm{HP}_I(d) = 4$ for all $d \geq 3$, $\dim(I) = 0$ (corresponding to the fact that $V(I)$ is a point), and $\deg(I) = 4 \cdot 0! = 4$.*

Example 6 (monomial ideals in $F[x, y]$) *For a monomial ideal I in $F[x, y]$, $V(I)$ is either (i) $\emptyset$, (ii) the origin, (iii) the x-axis, (iv) the y-axis, (v) the union of the x and y axes, or (vi) the entire plane (depending on which "types" of monomials occur among the generators of I — where a monomial can be of one of the following four types: (a) 1, (b) x^i, $i > 0$ (c) y^i, $i > 0$ or (d) x^iy^j $i, j > 0$). By calculating the rate of growth of HP_I in each of these cases, it is easy to see that $\dim(I)$ will be (i) $-\infty$ (by convention), (ii) 0, (iii) 1, (iv) 1, (v) 1, or (vi) 2, respectively, regardless of what monomials are used to generate the ideal. The degree, however, does depend on the particular generating monomials. If however, we take $I = I(V)$ for one of these varieties, the lowest degree generators will always be used (namely (i) (1), (ii) (x, y), (iii) (x), (iv) (y), (v) (xy), or (vi) (0), respectively) and the degree of these varieties will be (i) 0 (by convention), (ii) 1, (iii) 1, (iv) 1, (v) 2, or (vi) 1, respectively.*

Example 7 (hypersurfaces) *$I = (f) \subset F[x_1, \ldots, x_n]$. Then $\mathrm{LT}(I)$ is the set of monomials which are divisible by $\mathrm{LT}(f)$. The number of these of degree at most d is $\binom{d+n-\deg(f)}{n}$, so $\mathrm{HP}_I(d) = \binom{d+n}{n} - \binom{d+n-\deg(f)}{n}$, whose leading term is $\deg(f) \cdot d^{n-1}/(n-1)!$. Thus $\dim(I) = n-1$ (which corresponds to the fact that $V(I)$ is a hypersurface) and $\deg(I) = \deg(f)$.*

Before doing our final computation of dimension and degree, we give another characterization of the Hilbert polynomial. Let P_d be the set of all polynomials in $F[x_1, \ldots, x_n]$ of degree at most d, and for any ideal I, let I_d be the set of polynomials of degree at most d in I. Then P_d is a vector space over F of dimension $\binom{n+d}{d}$ and I_d is a subspace of P_d.

Proposition 8 *The number of monomials of degree at most d in $LT(I)$ equals the dimension of I_d as a vector space over F. Equivalently, $\mathrm{HP}_I(d)$ equals the codimension[1] of I_d in P_d.*

Proof Suppose there are s monomials of degree at most d in $\mathrm{LT}(I)$. Then there are polynomials $p_1, \ldots, p_s$ in I whose leading terms are these monomials. These polynomials are all of degree at most d and are linearly independent because they have distinct leading monomials. Hence, the dimension of I_d is at least s.

Conversely, suppose the dimension of I_d is t. Let $p_1, \ldots, p_t$ be a basis for I_d. We will construct polynomials $q_1, \ldots, q_t$ in sequence satisfying

1. $\mathrm{Span}(q_1, \ldots, q_i) = \mathrm{Span}(p_1, \ldots, p_i)$ for every i, and
2. The q_i's have distinct leading monomials.

We can take q_1 to be p_1. Now inductively assume that we have found $q_1, \ldots, q_i$. If the leading monomial of p_{i+1} is different than all the leading monomials of $q_1, \ldots, q_i$, then we take $q_{i+1} = p_{i+1}$ and we're done. Otherwise, we can subtract a scalar multiple of one of the q_j's to obtain p'_{i+1} whose leading monomial is smaller (with respect to a fixed ordering on monmials) than that of p_i and still satisfies $\mathrm{Span}(q_1, \ldots, q_i, p'_{i+1}) = \mathrm{Span}(p_1, \ldots, p_{i+1})$. We keep eliminating the leading monomial until we obtain p^*_{i+1} whose leading monomial is distinct from all the q_j's, and then we can take $q_{i+1} = p_{i+1}$. This process must stop after finitely many steps because each time we reduce the leading monomial. Once we have finished, the leading monomials of $q_1, \ldots, q_t$ give us t distinct monomials of degree at most d in $\mathrm{LT}(I)$.

For the "equivalently" part of the proposition, note that $\mathrm{HP}_I(d)$ is defined to be $\binom{n+d}{d}$ minus the number of monomials of degree at most d in $\mathrm{LT}(I)$, and the codimension of I_d in P_d is $\binom{n+d}{d}$ minus the dimension of I_d as a vector space over F. ■

Now let $V \subset F^n$ be any variety, and let F^V denote the set of functions from V to F. For every d, there is a linear map $\psi_{V,d}$ from P_d to F^V which simply takes a polynomial $p \in P_d$ and maps it to $p|_V \in F^V$. The kernel of this map are those polynomials of degree at most d which vanish on all of V, i.e. the kernel is I_d for $I = I(V)$. Thus the dimension of the image (as a vector space over F) is exactly the codimension I_d in P_d, and by Proposition 8, we have

Corollary 9 *For $I = I(V)$, $\mathrm{HP}_I(d)$ equals the dimension of $\mathrm{Im}(\psi_{V,d})$ as a vector space over F.*

Now we can calculate the dimension and degree of finite sets.

Example 10 (finite sets) *Let V be a finite set of m distinct points in F^n. We will show that V has dimension 0 and degree m. F^V has dimension m, so $\mathrm{Im}(\psi_{V,d}) \subset F^V$ has dimension at most m. By Corollary 9, for $I = I(V)$, $\mathrm{HP}_I(d) \leq m$ for any d, so the dimension of V is zero and the degree of V is at most m. To show that the degree of V actually equals m, we need to show that for sufficiently large d, $\psi_{V,d}$ is surjective. To see this note that for every point $a \in V$, we can find a polynomial p_a of degree $m-1$ which is nonzero at a but vanishes on the all the $m-1$ points in $V \setminus \{a\}$. These polynomials $\{p_a|_V\}$ span the space of all functions from V to F, so $\psi_{V,d}$ is surjective for all $d \geq m-1$.*

[1] If V is a subspace of W then the *codimension* of V in W is $\dim(W) - \dim(V)$.

Statements of Bézout's Theorem

The most general version of Bézout's theorem is the following:

Theorem 11 (Bézout) *For any two varieties V and W (satisfying some "niceness" conditions that are beyond the scope of this course) over an algebraically closed field,*

$$\deg(V \cap W) \leq \deg(V) \cdot \deg(W).$$

Applying this to the special case of hypersurfaces and finite sets of points (Examples 7 and 10), we obtain

Theorem 12 (Bézout for hypersurfaces) *Let F be an algebraically closed field, and suppose $f_1, \ldots, f_k \in F[x_1, \ldots, x_n]$ are of degrees $d_1, \ldots, d_k$, respectively. Then the number of common zeroes of $f_1, f_2, \ldots, f_k$ is either infinite at most $\prod_i d_i$*

We will prove a special case of this for curves in the plane. We want a result that also holds for finite fields, so the case that the number of common zeroes is infinite will be replaced by something meaningful for finite fields.

Theorem 13 (Bézout in the plane) *Let F by any field (not necessarily algebraically closed) and let $f, g \in F[x, y]$ be polynomials of degree m and n, respectively. Then f and g either share a nontrivial common factor or have at most mn common zeroes in F^2*

Proof of Bézout's theorem in the plane

First, we observe that it suffices to prove Theorem 13 for infinite fields: If we're given two polynomials $f, g \in F[x, y]$ over a finite field F which have more than mn common zeroes in F^2, they they also have more than mn common zeroes in L^2 for any infinite field L containing F. Assuming Theorem 13 holds for infinite fields, we can deduce that f and g have a nontrivial common factor $h \in L[x, y]$. The degree of h in one of the variables, say x, must be greater than 0. This implies that $\mathrm{Res}_x(f, g) = 0$, where the resultant is computed in $L(y)$. But the resultant will be the same if we compute it in $F(y)$, so we conclude that f and g have a common factor in $F(y)[x]$. By Gauss' Lemma, we deduce that they have a common factor in $F[x, y]$

Now let's proceed to the proof of Theorem 13 for infinite fields. Suppose f and g have more than mn zeroes. Let S be any set of $mn + 1$ distinct zeroes of f and g. We first perform a random linear transformation to make the second coordinates of all the elements of S distinct. Specifically, choose α uniformly at random from some subset of L of size greater than $(mn+1)^2$ and let $f'(x, y) = f(x, y + \alpha x)$, $g'(x, y) = g(x, y + \alpha x)$, $S' = \{(u, v - \alpha u) \colon (u, v) \in S\}$. Then, for every α, every element of S' is common zero of f' and g', and f' and g' have a common factor iff f and g do. Moreover, we claim that, with nonzero probability over the choice of α, no two elements of S' have the same second coordinate. (For every pair of points $(u_1, v_1), (u_2, v_2) \in S$, the probability that $v_1 - \alpha u_1$ equals $v_2 - \alpha u_2$ is smaller than $1/(mn + 1)^2$. Since there are fewer than $(mn + 1)^2$ pairs to consider, the claim follows by a union bound.) So fix any α for which no two elements of S' have the same second coordinate.

Now consider $p(y) = \mathrm{Res}_x(f', g')$, which is a polynomial in y. We have argued in previous lectures that p is of degree at most $\max\{m, n\}^2$. In fact, we will show later that p is of degree at most mn; assuming this fact, we continue with the proof. For every point $(u, v) \in S'$, $p(v) = \mathrm{Res}_x(f'(x, v), g'(x, v)) = 0$, because $f'(x, v)$ and $g'(x, v)$ have the common zero $x = u$. Thus p has at least $mn + 1$ distinct roots, and must be the zero polynomial. In other words, $\mathrm{Res}_x(f', g') = 0$, so f' and g' have a common factor in $F(y)[x]$, implying a common factor in $F[x, y]$ by Gauss' Lemma.

This completes the proof of Theorem 13, except for the following fact about resultants.

Proposition 14 *If $f, g \in F[x, y]$ are polynomials of total degree at most m and n, respectively, then $\mathrm{Res}_x(f, g)$ is a polynomial of degree at most mn in y.*

Proof We can write $f(x, y) = a_0(y) + a_1(y)x + \cdots + a_m(y)x^m$ and $g(x, y) = b_0(y) + b_1(y)x + \cdots + b_n(y)x^n$, where a_j (resp., b_j) is a polynomial of degree at most $m - j$ (resp., $n - j$). Recall that $\mathrm{Res}_x(f, g)$ is the determinant of the $(m + n) \times (m + n)$ matrix M given by

$$M_{ij} = \begin{cases} a_{i-j}(y) & \text{if } j \leq n \\ b_{n+i-j}(y) & \text{if } j > n, \end{cases}$$

where the $a_k(y)$ and $b_k(y)$ are defined to be 0 when k is out of bounds. Thus,

$$\deg(M_{ij}) \leq \begin{cases} m - (i - j) = m - i + j & \text{if } j \leq n \\ n - (n + i - j) = j - i & \text{if } j > n, \end{cases}$$

Since $\det(M) = \sum_\sigma \left(\pm \prod_{j=1}^{m+n} M_{\sigma(j)j} \right)$, where the sum is over all permutations $\sigma \colon [m+n] \to [m+n]$, we have

$$\begin{aligned} \deg(\det(M)) &\leq \max_\sigma \left(\sum_{j=1}^{m+n} \deg(M_{\sigma(j)j}) \right) \\ &\leq \max_\sigma \left(\sum_{j=1}^{n} (m + j - \sigma(j)) + \sum_{j=n+1}^{n+m} (j - \sigma(j)) \right) \\ &= \max_\sigma \left(mn + \sum_{j=1}^{n+m} j - \sum_{j=1}^{n+m} \sigma(j) \right) \\ &= mn, \end{aligned}$$

where in the last equality we use the fact that σ is a permutation. ■

Application 1 — Goppa Codes

Bézout's theorem has many applications in complexity and coding theory. It can be viewed as a generalization of the fact that a degree d polynomial has at most d roots (take $f(x, y) = y - p(x)$ and $g(x, y) = y$ in Theorem 13). This leads to a natural generalization of Reed–Solomon codes, called Goppa codes.

Recall that the Reed–Solomon code identifies "messages" with polynomials p of degree at most $k-1$ over some finite field $\mathbb{F}_q$. The codeword associated with a message/polynomial p is $(p(a_1), \ldots, p(a_n))$ where $\{a_1, \ldots, a_n\}$ is a fixed subset of $\mathbb{F}_q$. n is called the *block length* of the code. The important property of RS codes is that every two distinct codewords disagree in at least $d = n - k + 1$ places; d is called the *distance* of the code. Our messages (polynomials of degree at most $k - 1$) can we thought of as strings of length k over $\mathbb{F}_q$, so k is called the *rate* of the code. Thus, the relationship among the rate, distance, and block length of RS codes can be summarized as $k = n - d + 1$. The main disadvantage of RS codes is that the alphabet size q must be at least as large as the block length n. (because $\{a_1, \ldots, a_n\}$ must be a subset of $\mathbb{F}_q$.) This means that if one wants to encode long messages, one has to use a large alphabet. Large alphabets have the disadvantage that one can corrupt an entire symbol by only corrupting one bit in a binary representation of the symbol, and hence by corrupting only a $1/(\log_2 q)$ fraction of the bits in a binary representation of a codeword, every symbol in the codeword can be corrupted and error-correction is impossible.

Reed-Solomon codes can be thought of as evaluating low-degree polynomials on a line (namely, the elements of $\mathbb{F}_q$), which has only q points. An idea for increasing the block-length relative to the alphabet size would be to evaluate low-degree polynomials on a curve C with more than q points. We will still need the properties that no two distinct low-degree polynomials agree on C in many places and that there are many distinct low-degree polynomials on C — Bézout's theorem will enable us to prove these properties.

Fix an irreducible bivariate polynomial $f(x, y)$ of degree $\ell < q$ over $\mathbb{F}_q$, and let $C \subset \mathbb{F}_q^2$ be the zero-locus of f (i.e., $C = V((f))$). There exist such curves containing nearly $q\ell$ points (since f is of degree ℓ, the most we can hope for is ℓ choices for y for every value of x), and it is known how to find such curves efficiently. So let us assume that we have a curve with $n = |C| \approx q\ell$. Our messages will be interpreted as bivariate polynomials $g(x, y)$ of degree at most m (where $\ell \leq m < n/\ell$) over $\mathbb{F}_q$, and the corresponding codeword will be the evaluation of g on all points of C. We need to be careful because two polynomials g_1 and g_2 can agree everywhere on C. However, this means that the polynomials $g_1 - g_2$ and f have at least $n > m\ell$ common zeroes, so by Bézout's theorem $f|(g_1 - g_2)$.[2] Thus the number of distinct codewords we obtain is the number of bivariate polynomials of degree m ($= q^{\binom{m+2}{2}}$) divided by the number of bivariate polynomials of degree m which are divisible by f ($q^{\binom{m-\ell+2}{2}}$). In other words, the dimension of the space of messages we can encode, or *rate* of the code, is

$$k = \binom{m+2}{2} - \binom{m-\ell+2}{2} = \ell m + 1 - \binom{\ell-1}{2}.$$

Now we analyze the distance of the code. The same Bézout analysis we did above shows that no two polynomials g_1 and g_2 of degree m can agree on more than ℓm points of C unless $f|(g_1 - g_2)$ (in which case g_1 and g_2 agree everywhere on C and we consider them g_1 and g_2 to be the same codeword). Thus, the distance of the code is at least $d = n - \ell m$.

Summarizing, we relationship between rate, block length, and distance is now

$$k = n - d + 1 - \binom{\ell-1}{2}.$$

This is worse than what RS codes achieve by the defect $\binom{\ell-1}{2}$, but we have gained in that the alphabet size need only be about n/ℓ. If we take $\ell = \Theta(\sqrt[3]{n})$, q will be $\Theta(n^{2/3})$ (a substantial improvement over RS codes), and the defect is only $O(n^{2/3}) = o(n)$.

Generalizing these ideas (quite a bit further) eventually leads to *algebraic geometric codes*, which have no restriction on the alphabet size and satisfy

$$k \approx \left(1 - \frac{1}{\sqrt{q}}\right) n - d.$$

Application 2 — Decoding Reed–Solomon Codes

In general, the *decoding* problem for an error-correcting code is: Given a codeword which has been corrupted in some number of positions, find the message corresponding to the codeword (equivalently, find the uncorrupted codeword). In case there are so many errors that several messages could have produced the given string, the *list-decoding* problem is to find all possible messages that are consistent with the number of errors.

For Reed–Solomon codes, this can be reformulated as the following curve-fitting question: Given n points $(x_1, y_1), \ldots, (x_n, y_n)$, find a degree k polynomial p such that $p(x_i) = y_i$ for as many i possible. A widely used approach to solving this problem is algebraic-geometric in nature:

[2] If we were over an algebraically closed field, this would follow from the Nullstellensatz: $g_1 - g_2$ is identically zero on $C = V((f))$, so $(g_1 - g_2) \in I(V((f))) = \mathrm{Rad}((f)) = (f)$.

1. Find a low-degree algebraic curve $f(x, y) = 0$ which goes through all these points (x_i, y_i),
2. Decompose the curve into its irreducible components (by factoring f), and
3. Hope that the one of these components is the graph of the polynomial p we're looking for (i.e. $y - p(x)$ appears as a factor of f).

Let's prove (using Bézout's theorem) that this approach works for some settings of parameters. We will show that it can find any polynomial p which agrees with at least $2\sqrt{n} \cdot k$ of the points. To do this, we will look for C described by a polynomial $f(x, y)$ of total degree at most $2\sqrt{n}$. A nonzero polynomial f of this degree such that $f(x_i, y_i) = 0$ for all i always exists, simply because each pair (x_i, y_i) imposes a homogeneous linear constraint on the coefficients of f and f has at $\binom{2\sqrt{n}+2}{2} > n$ coefficients. To find f, we can just solve this linear system. Now we only need to prove that any polynomial p such that $p(x_i) = y_i$ for enough of the i's will appear in a factor of f of the form $y - p(x)$. Well, f and $y - p(x)$ are two algebraic curves of degree $2\sqrt{n}$ and k, respectively, that agree in many points. By Bézout's theorem, they share a common factor if they agree in at least $2\sqrt{n} \cdot k$ points. Since $y - p(x)$ is irreducible, this implies $y - p(x)$ will occur among the irreducible factors of f, and the algorithm works.

Note that for $k \ll \sqrt{n}$, this is list-decoding far beyond the range in which the answer is unique — there is guaranteed to be a unique decoding only when the the number of uncorrupted positions is greater than $(n + k)/2$. The best known result for this problem, due to Guruswami and Sudan (1998), allows list-decoding in polynomial-time when the number of uncorrupted positions is greater than $\sqrt{nk}$.