Lecture 10: Low degree testing

Topics in Pseudorandomness and Complexity Theory (Spring 2017)
Rutgers University
Swastik Kopparty
Scribes: Mrinal Kumar, Ana Echavarria

1 Plan

In the last class we showed that every language in NEXP has an exponential length proof which can be verified in randomized polynomial time. In this lecture we will prove the one tool which we used as a blackbox in the last class, namely *low degree testing*. More formally, the goal of this lecture is to prove the following theorem.

Theorem 1. *Let p be a prime power, d and m be any positive integers such that $d < p/2$, and $\epsilon \in (0,1)$ be a parameter. Then, there is a randomized algorithm $\mathcal{A}$, which when given access to a function $f : \mathbb{F}_p^m \to \mathbb{F}_p$, runs in time $\mathsf{Poly}(\log p, m, d, \frac{1}{\epsilon})$ and satisfies the following.*

1. *If f is a polynomial of degree at most d, then $\mathsf{Pr}[\mathcal{A}$ accepts$] = 1$.*
2. *If f is ϵ-far from every polynomial of degree at most d, then $\mathsf{Pr}[\mathcal{A}$ rejects$] \geq 0.9$.*

Recall that two functions f and g are said to be ϵ-far from each other if they differ on at least an ϵ fraction of all inputs.

2 The low degree test

Recall the following elementary fact, which we have already seen in some of the earlier lectures.

Fact 2. *For any $d+1$ distinct values $\alpha_1, \alpha_2, \ldots, \alpha_{d+1}$ in $\mathbb{F}_p$ and any (possibly non-distinct) $\beta_1, \beta_2, \ldots, \beta_{d+1}$ in $\mathbb{F}_p$, there is a unique univariate polynomial P of degree at most d such that*

$$\forall i \in \{1, 2, \ldots, d+1\} \quad P(\alpha_i) = \beta_i$$

Fact 2 motivates the following natural low degree test for univariate polynomials.

Test 1 : A candidate test for univariates

- Given $f : \mathbb{F}_p \to \mathbb{F}_p$.
- Pick any $d+1$ distinct values $\alpha_1, \alpha_2, \ldots, \alpha_{d+1}$ in $\mathbb{F}_p$.

- Pick a random $\gamma \in \mathbb{F}_p$.
- Find the unique polynomial P of degree at most d which passes through the points

$$(\alpha_1, f(\alpha_1)), (\alpha_2, f(\alpha_2)), \ldots, (\alpha_{d+1}, f(\alpha_{d+1}))$$

- If $P(\gamma) = f(\gamma)$, then accept else reject.

Observe that if f is a polynomial of degree at most d, then the above algorithm will always accept. It is also not hard to see that if f is at least ϵ-far from every degree d polynomial then the test rejects f with probability at least ϵ.

This algorithm has the following natural generalization to the case of larger number of variables.

Test 2 : A candidate test for multivariates

- Given $f : \mathbb{F}_p^m \to \mathbb{F}_p$.
- Query the values of f on any set $S \subseteq \mathbb{F}_p^m$ which is an interpolating set for degree d polynomials.
- Find the unique polynomial P of degree at most d which agrees with f everywhere on the set S.
- Accept if P and f agree at a random point in $\mathbb{F}_p^m$, reject otherwise.

Again, the test accepts all polynomials of degree at most d and rejects everything else with a probability at least ϵ. However, the algorithm makes at least $\binom{d+m}{m}$ queries, since $|S| \geq \binom{d+m}{m}$. Therefore, the running time of the algorithm is not necessarily $\mathsf{Poly}(m, d)$. Our goal is to come up with a test which makes far fewer queries to the evaluation table of f.

Test 3 : A more query efficient test for multivariates

- Given $f : \mathbb{F}_p^m \to \mathbb{F}_p$.
- Pick random $x \in \mathbb{F}_p^m$ and $a \in \mathbb{F}_p^m$ and consider the line

$$\ell = \{x + ta : t \in \mathbb{F}_p\}$$

- Check that the restriction of f on ℓ (denoted by $f|_\ell$) is of degree at most d using the univariate Test 1.

In the rest of the class, we will analyze and understand the behavior of Test 3 described above, and show that a minor variant of it does indeed give us Theorem 1. We start with the following easy observation.

Observation 3. *If f is a polynomial of degree at most d, then Test 3 accepts with a probability* 1.

The proof of the observation is straightforward since the restriction of any polynomial of degree at most d to any line is a polynomial of degree at most d. In the next subsection, we show that if f is not a polynomial of degree at most d, then Test 3 rejects with a non-zero probability. Although this does not suffice for proving Theorem 1 (where we need to reject with a high probability), it is an important step in that direction.

2.1 Test 3 - non-zero rejection probability

In this section we will show the following theorem.

Theorem 4. *If $d < p/2$, and for every line $\ell \subseteq \mathbb{F}_p^m$, $f|_\ell$ has degree at most d, then the degree of f is at most d.*

We start with the following lemma, which says something about the rejection probability in the very special case when f is a polynomial of degree D such that $d < D < p$. This turns out to be useful in the proof of Theorem 4.

Lemma 5. *Let $f : \mathbb{F}_p^m \to \mathbb{F}_p$ be a polynomial of degree equal to D such that $D < p$. Then, there exists a line $\ell \subseteq \mathbb{F}_p^m$ such that $f|_\ell$ has degree equal to D.*

Proof. For the sake of contradiction, let us assume that for every line $\ell \subseteq \mathbb{F}_p^m$, $f|_\ell$ has degree strictly less than D.

Since f is a polynomial of degree equal to $D < p$ in m variables, we know that f can be written uniquely as

$$f = \sum_{\overline{e} \text{ s.t} \sum_{i=1}^m e_i \leq D} c_{\overline{e}} \prod_{i=1}^m X_i^{e_i}$$

Note that here the sum is over all $\overline{e} = (e_1, e_2, \ldots, e_m)$ such that $\forall i, e_i \in \{0, 1, \ldots, p-1\}$. Now, let us consider the *formal* restriction of f to a line $\ell = B + TA$. Here, B denotes the m-tuple $(B_1, B_2, \ldots, B_m)$ and A denotes the m-tuple $(A_1, A_2, \ldots, A_m)$. So, by definition, we get

$$f|_\ell = \sum_{\overline{e} \text{ s.t} \sum_{i=1}^m e_i \leq D} c_{\overline{e}} \prod_{i=1}^m (B_i + TA_i)^{e_i}$$

Now, observe that the coefficient of T^D in $f|_\ell$ is given by

$$\sum_{\overline{e} \text{ s.t} \sum_{i=1}^m e_i = D} c_{\overline{e}} \prod_{i=1}^m A_i^{e_i}$$

Since $D < p$, and we have assumed that the restriction of f to every line in $\mathbb{F}_p^m$ is of degree strictly less than D, it must be the case that the coefficient of T^D in $f|_\ell$ is identically zero as a polynomial in $A_1, A_2, \ldots, A_m$ and $B_1, B_2, \ldots, B_m$. Thus,

$$\sum_{\overline{e} \text{ s.t} \sum_{i=1}^m e_i = D} c_{\overline{e}} \prod_{i=1}^m A_i^{e_i} \equiv 0$$

But then, this implies that $c_{\overline{e}} = 0$ for every $\overline{e} = (e_1, e_2, \ldots, e_m)$ such that $\sum_{i=1}^{m} e_i = D$. This contradicts the fact that degree of f is equal to D. Thus, f cannot be of degree strictly smaller than D on every line in $\mathbb{F}_p^m$. □

We now proceed to prove the following special case of Theorem 4 for the special case when $m = 2$. This will serve as another important step towards the proof of Theorem 4.

Theorem 6. *If $d < p/2$, and for every line $\ell \subseteq \mathbb{F}_p^2$, $f|_\ell$ has degree at most d, then the degree of f is at most d.*

Proof. We will first use the fact the restriction of f to every line has degree at most d to conclude that the degree of f can be no larger than $2d$. Since $d < p/2$, this implies that the degree of f is strictly less than p. The proof then follows from Lemma 5.

Since f is a function from $\mathbb{F}_p^2$ to $\mathbb{F}_p$, there is a unique polynomial of individual degree at most $p-1$ in $\mathbb{F}_p[X, Y]$ which is equal to f as a function over $\mathbb{F}_p^2$. Let the polynomial be given by

$$f = \sum_{j \leq p-1} \left(\sum_{i \leq p-1} a_{i,j} X^i \right) Y^j$$

For every j, let us denote by $R_j(X)$ the univariate polynomial $\left(\sum_{i \leq p-1} a_{i,j} X^i\right)$. So, we get

$$f = \sum_{j \leq p-1} R_j(X) Y^j$$

Let us look at the restriction of f to vertical and horizontal lines in $\mathbb{F}_p^2$, namely the lines of the form $\{(u, y) : y \in \mathbb{F}_p\}$ and $\{(x, v) : x \in \mathbb{F}_p\}$ for every u and v in $\mathbb{F}_p$. The formal restriction of f to the line $\ell_u = \{(u, y) : y \in \mathbb{F}_p\}$ is given by

$$f|_{\ell_u} = \sum_{j \leq p-1} R_j(u) Y^j$$

Now, $f|_{\ell_u}$ is a univariate polynomial of degree at most $p-1$, which by our hypothesis agrees with a polynomial $P(Y)$ of degree at most $d < p/2$ over all of F_p. Thus, by Fact 2 it follows that $f|_{\ell_u}$ must in fact be a polynomial of degree at most d (as a formal object). Thus, the coefficient of Y^j in $f|_{\ell_u}$ must be equal to zero for every $j > d$. So, for every $j > d$

$$R_j(u) = 0$$

Moreover, since $R_j(X)$ is itself a polynomial of degree at most $p-1$, and for every $u \in \mathbb{F}_p$, $R_j(u) = 0$, it follows from another application of Fact 2 that $R_j(X)$ is in fact identically zero as a formal polynomial. In summary, we conclude that the degree of Y in f is at most d.

By a similar argument, where we work with horizontal lines, we can conclude that the degree of X in f is at most d. So, the total degree of f is at most $2d$, which by our choice of d is strictly smaller than p.

The theorem now follows immediately via an application of Lemma 5. □

From the above proof we get a weaker version of Theorem 4 for the case when $d < p/m$, and this might be sufficient for some applications. To get to the case when $p/2 > d > p/m$, we apply an induction on m.

Proof of Theorem 4. We are already done when $m \leq 2$. So, we assume that $m > 2$ and the theorem is true for polynomials on $m-1$ variables. By writing f as a polynomial of degree at most $p-1$ in every variable $X_1, X_2, \ldots, X_m$ and grouping together monomials based on the degree of X_m, we get that f can be written as

$$f = \sum_{i \leq p-1} R_i(X_1, X_2, \ldots, X_{m-1}) Y^j$$

Let ℓ be a line in $\mathbb{F}_p^{m-1}$ given by

$$\ell = \{y + bt : t \in \mathbb{F}_p\}$$

Here, $y, b \in \mathbb{F}_p^{m-1}$. Let $x, a \in \mathbb{F}_p^m$ be given by

$$x = (y, \alpha) \text{ and } a = (b, 0)$$

and let us consider the line $\ell'_\alpha = \{x + at : t \in \mathbb{F}_p\}$. From our hypothesis, the degree of f restricted to line ℓ'_α is at most d. So, for every $\alpha \in \mathbb{F}_p$, there is a polynomial $S_\alpha(T)$ of degree at most d such that for every $t \in \mathbb{F}_p$

$$S_\alpha(t) = \sum_{i \leq p-1} R_i(y + tb)\alpha^i$$

Since, the vectors $(1, \alpha, \alpha^2, \ldots, \alpha^{p-1})$ are linearly independent for distinct non-zero values of $\alpha \in \mathbb{F}_p$ and span $\mathbb{F}_p^{p-1}$. Therefore, by taking an appropriate $\mathbb{F}_p$ linear combination of $\{S_\alpha(T) : \alpha \in \mathbb{F}_p\}$, we get that for every $i \leq p-1$, there is a polynomial P_i of degree at most d which agrees with $R_i|_\ell$. Since this is true for all lines $\ell \subseteq \mathbb{F}_p^{m-1}$, we can conclude that every R_i satisfies the hypothesis of the theorem and is a function in $m-1$ dimensions. From the induction hypothesis, it follows that the degree of every R_i is at most d. Moreover, from an argument which is identical to the one we used in the proof of Theorem 6, we can conclude that the individual degree of every variable in f is at most d. So, the total degree of f is at most $2d$, which is strictly less than p. The theorem now follows immediately from Lemma 5. □

2.2 Acceptance and closeness

Let us apply a small modification to test 3 to define test 4.

Test 4 : Polynomial encoding

- Given $f : \mathbb{F}_p^m \to \mathbb{F}_p$.
- Fix $\alpha_1, \alpha_2, \ldots, \alpha_{d+1}$ distinct and non-zero
- Pick random $x \in \mathbb{F}_p^m$ and $h \in \mathbb{F}_p^m$ uniformly.

- Let $P_{x,h}(T)$ be the unique univariate polynomial of degree $\leq d$ s.t. for each $i \in \{1, 2, \ldots, d+1\}$

$$P_{x,h}(\alpha_i) = f(x + \alpha_i h)$$

- Accept if $P_{x,h}(0) = f(x)$, reject otherwise.

Now, let us use this test to show that a function that gets accepted in the test, with high probability is close to degree d polynomials.

Theorem 7. *If a function $f : \mathbb{F}_p^m \to \mathbb{F}_p$ passes test 4 with probability $1 - \alpha$, then f is close to some degree d polynomial.*

To prove this theorem let us first define $g :: \mathbb{F}_p^m \to \mathbb{F}_p$ be the "self correction" of f where

$$g(x) = \underset{h \in \mathbb{F}_p^m}{\text{Maj}} \left(P_{x,h}(0) \right)$$

That is, the function that selects the most common value of $P_{x,h}(0)$ when running the test for all possible values of the slope h. Now, let's use this definition for the following lemma about the closeness between f and g.

Lemma 8. *If a function $f : \mathbb{F}_p^m \to \mathbb{F}_p$ passes test 4 with probability $1 - \alpha$, then*

$$\mathsf{Pr}_x \left[g(x) \neq f(x) \right] \leq 2\alpha$$

Proof. Let f be a function that passes test 4 with probability $1 - \alpha$, then

$$\mathsf{Pr}_{x,h} \left[f(x) \neq P_{x,h}(0) \right] \leq \alpha$$
$$\mathsf{Pr}_x \left[\mathsf{Pr}_h [f(x) \neq P_{x,h}(0)] \geq \frac{1}{2} \right] \leq 2\alpha$$

But if $\mathsf{Pr}_h[f(x) \neq] P_{x,h}(0)] < \frac{1}{2}$ then $g(x) = f(x)$ so $\mathsf{Pr}_x[g(x) \neq f(x)] \leq 2\alpha$ □

Now, we want to show that g has degree at most d using the following lemma

Lemma 9.

Proof.
Part 1: We will prove that for all x, $\mathsf{Pr}_{h_1,h_2} \left[P(x, h_1(0) = P_{x,h_2}(0) \right] \geq 1 - O(d\alpha)$

Let us think of a table of the following form

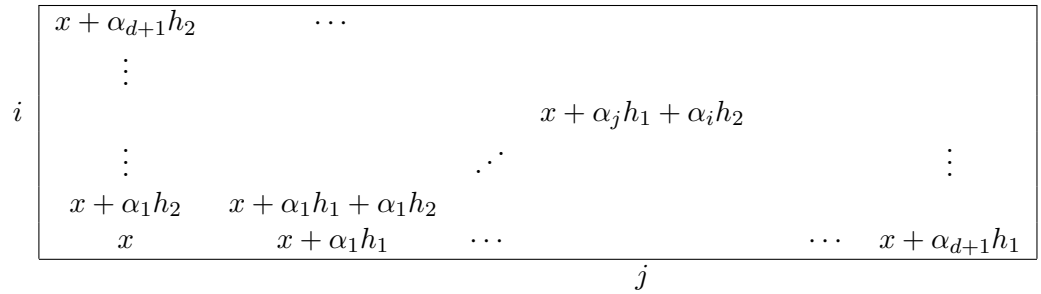

Note that we can run the test for each element of the first row and the first column, except the first element x, and obtain a degree d polynomial that fits each row and each column except the first row and column. Additionally, we obtain the following probabilities of acceptance of the test.

$$\forall j \ \mathsf{Pr}_{h_1,h_2}\left[P_{x+\alpha_j h_1,h_2}(0) = f(x+\alpha_j h_1)\right] \geq 1-\alpha$$
$$\forall i \ \mathsf{Pr}_{h_1,h_2}\left[P_{x+\alpha_i h_2,h_1}(0) = f(x+\alpha_i h_2)\right] \geq 1-\alpha$$

so with probability at least $(S \setminus \{0\})$ all elements of the first column and all elements of the first row, except the first element x, pass the test.

Now, a way to generalize the idea above is to define a set $S = \{0, \alpha_1, \ldots, \alpha_d + 1\}$ and a function $A : S \times S \to \mathbb{F}_p$ where

$$A(u,v) = f(x + uh_1 + vh_2)$$

Now, for all $u \neq 0$, $A(u,\cdot)$ has degree $\leq d$ and for all $v \neq 0$, $A(\cdot,v)$ has degree $\leq d$.

Let $Q(X,Y)$ be the unique degree (d,d) polynomial agreeing with $A(u,v)$ on $(S \setminus \{0\}) \times (S \setminus \{0\})$ then

$$A(0,v) = Q(0,v) \ \forall v \in (S \setminus \{0\})$$
$$A(u,0) = Q(u,0) \ \forall u \in (S \setminus \{0\})$$

which implies that $P_{x,h_1}(0) = Q(0,0) = P_{x,h_2}$. And this occurs when all the test pass which we showed happens with probability at least $1 - O(d\alpha)$.

Part 2: We will prove that for all x, $\mathsf{Pr}_h\left[P_{h,x}(0) = g(x)\right] \geq 1 - O(d\alpha)$

We know that $\mathsf{Pr}_{h_1,h_2}\left[P(x, h_1(0) = P_{x,h_2}(0)\right] \geq 1 - O(d\alpha)$.
Now

$$\begin{aligned}\mathsf{Pr}_{h_1,h_2}\left[P(x, h_1(0) = P_{x,h_2}(0)\right] &= \sum p_i^2 \\ &\leq p_{\max} \sum p_i \\ &= p_{\max} \\ &= \mathsf{Pr}_h\left[P_{h,x}(0) = g(x)\right]\end{aligned}$$

Where each of the p_i are the values of the probabilities for choices of h and $p_{\max}$ is the maximum over the p_i. Then $\mathsf{Pr}_h\left[P_{h,x}(0) = g(x)\right] \geq 1 - O(d\alpha)$.

Part 3: We will use the result in part 2 to prove that g has degree $\leq d$. More specifically, we will prove that if g passes the test with probability 1 and we proved earlier that this happens if and only if g has degree at most d.

For all x, h let a, b be picked uniformly at random. Consider the table M where each row is a test and each column is a test. It is completely defined when we set the first two columns to be in direction a and b respectively and the first row to be in direction h.

	$x+\alpha_{d+1}a$	$x+\alpha_1 h+\alpha_{d+1}b$	$\cdots$		
	$\vdots$	$\vdots$	$\ddots$		$\vdots$
	$x+\alpha_1 a$	$x+\alpha_1 h+\alpha_1 b$	$\cdots$		
dir h	x	$x+\alpha_1 h$	$x+\alpha_2 h$	$\cdots$	$x+\alpha_{d+1}h$
	dir a	dir b			

There exists a unique way to fill up this table such that for all $u, v \in S$

$$M(U,V) = A_u + VB_u$$
$$M(U,V) = A_v + UB_v$$

That way, every row is a random test and every column is a random test centered at a given point $x + \alpha_i h$. Now,

$$\mathsf{Pr}(\text{all row tests pass}) \geq 1 - O(d\alpha)$$
$$\mathsf{Pr}(\text{all column tests pass}) \geq 1 - O(d^3\alpha)) = \mathsf{Pr}(\text{all tests pass for } g) > 0$$

So all tests pass for g which means that g has degree $\leq$ d. □

All that remains to prove theorem 7 is to see that by lemma 8 we have that if f passes test 4 then it is close to g which we know had degree at most d by lemma 9.

Next lecture, we will prove this these same results but we will increase the rejection probability.