PCPs and HDX - Lecture 12

Lecturer: Inbal Livni Navon

January 31, 2017

In this lecture we will see a recent work by Amey Bhangale, Irit Dinur and I that was presented in ITCS2017 [BDN17], it is called cube vs. cube low degree test, and can be found here. We begin with some background and reminders from the previous lessons.

1 Cube vs. Cube Low Degree Test

In low degree testing, we are given a function $f : \mathbb{F}^m \to \mathbb{F}$, and we want to find if it is a degree d polynomial by querying it only on a few points.

1.1 Rubinfeld and Sudan's low degree test

In lesson 4, we have seen Rubinfeld and Sudan's low degree test and theorem [RS96], the test

1. Choose random $x, h \in \mathbb{F}^m$.
2. Compute via interpolation the degree d polynomial for which $p(t) = f(x+th)$ for $t \in [d+1]$, accept if $p(0) = f(0)$.

Theorem 1.1. *If the test passes with probability* $1 - \delta$ *for* $\delta < \frac{1}{2(d+2)^2}$*, then there exist a degree* d *polynomial that is* $1 - 2\delta$ *close to* f*.*

1.2 Low Soundness

In the previous lesson talked about the low acceptance regime, i.e. we want a low degree theorem that holds even when the test acceptance probability is very small, $\delta > 0$, and not only when it is $1-\delta$. There have been previous works on low soundness low degree tests, [RS97],[AS97],[MR08], reducing the threshold acceptance probability δ to $O\left(\frac{1}{|\mathbb{F}|^{\frac{1}{8}}}\right)$. In our work, we reduce this value further to $O\left(\frac{1}{\sqrt{|\mathbb{F}|}}\right)$ (ignoring dependencies on m, d).

In order get low soundness low degree test with a few queries we need to query the function on more than a single point on each query, which brings us to the setting of the previous lesson.

Agreement Test Recall the previous lesson agreement test, we had a set $\mathcal{S} = \{S \subset U | |S| = k\}$, and an assignment over $\mathcal{S}$ was $a = \{a_S\}_{S \in \mathcal{S}}$, when $a_S : S \to \Sigma^k$. An agreement test on a with respect to the distribution $\mathcal{D}$ is:

1. Choose $S_1, S_2 \sim \mathcal{D}$.
2. Accept if $\forall u \in S_1 \cap S_2$, $a_{S_1}(u) = a_{S_2}(u)$.

The agreement of a is defined as the test success probability, i.e.

$$agree_{\mathcal{D}}(a) = \Pr_{S_1, S_2 \sim \mathcal{D}}[\forall u \in S_1 \cap S_2, a_{S_1}(u) = a_{S_2}(u)].$$

Goal: find $\mathcal{D}$ such that $agree_{\mathcal{D}}(a) > \delta \implies \exists g : U \to \Sigma$ s.t $\Pr_{S \in \mathcal{S}}[a_S = g|_S] \geq \delta'$.

Cube vs cube agreement test The cube vs. cube test is an agreement test in a very similar setting,

- $U = \mathbb{F}^m$, $k = |\mathbb{F}|^3$.
- Instead of using the set $\mathcal{S}$ defined above, we use $\mathcal{C} \subset \mathcal{S}$, which is defined by

$$\mathcal{C} = \{C \subset \mathbb{F}^m ||C| = k, C \text{ is an affine subspace}\}.$$

- The assignment $a = \{a_C\}_{C \in \mathcal{C}}$ satisfies: for every C, $a_C : C \to \mathbb{F}$ is a degree d polynomial (recall in SAT we allowed only satisfying assignments, here we do the same).
- The distribution $\mathcal{D}$ that we analyze is $\mathcal{D}_1$ from the previous lesson - we pick two random cubes $C_1, C_2 \in \mathcal{C}$ that intersects on a point $x \in \mathbb{F}^m$.

Theorem 1.2 (Cube vs. cube low degree theorem)**.** *There exist constants c_1, c_2, such that if $agree_{\mathcal{D}}(a) > \frac{c_1 d^4}{\sqrt{|\mathbb{F}|}}$, then there is a degree d polynomial g, such that $\Pr_{C \in \mathcal{C}}[a_C = g_{|C}] \geq c_2 \cdot agree_{\mathcal{D}}(a)$.*

2 Proof Sketch

In the proof we take an assignment $a = \{a_C\}_{C \in \mathcal{C}}$ such that $\epsilon = agree_{\mathcal{D}}(a) \geq \frac{c_1 d^4}{\sqrt{|\mathbb{F}|}}$, and find a low degree polynomial $g : \mathbb{F}^m \to \mathbb{F}$ that satisfies the theorem requirements.

The proof has three parts, in the first we look for candidate functions for g, in the second we pick one such function, and in the third show that it is close to a degree d polynomial.

2.1 Local structure

We are looking for a function $g : \mathbb{F}^m \to \mathbb{F}$ that **agrees** with many of the assignments a_C for cubes $c \in \mathcal{C}$. A very natural approach is to pick the most frequent value, i.e for every $x \in \mathbb{F}^m$, to look on all $C \in \mathcal{C}$ that contains x and take the most frequent $a_C(x)$. This approach is not going to work, see the examples bellow.

Example Let $P_1 \neq P_2 : \mathbb{F}^m \to \mathbb{F}$ be two different degree d polynomials, and let $a = \{a_C\}_{C \in \mathcal{C}}$ be the assignment such that for half of the cubes, $a_C = P_1|_C$ and for the second half, a_C is a restriction of P_2. This a has agreement $\geq \frac{1}{2}$ (because if C_1, C_2 are a restriction of the same polynomial the test will pass). a also has an agreement with a global function, for half of the cubes $C \in \mathcal{C}$, $a_C = P_1|_C$. In this section we want to find the function P_1.

How do we find P_1? obviously, taking the most frequent value is not going to work, as it equals P_1 on half of the points and P_2 on the other half. If someone tells us which cubes are a restriction of P_1 and which of P_2 we could find P_1 by taking the most frequent value, but we don't know which entries belong to which polynomial.

We use an idea that was used in direct product testing, which is to look on some x, and then choose only the cubes C such that $x \in C$ and $a_C(x) = \sigma = P_1(x)$. For almost all $x \in \mathbb{F}^m$, all these cubes satisfies $a_C(x) = P_1(x)$. Taking the most frequent value among these cubes gives us P_1.

More explicitly, we define $f_x : \mathbb{F}^m \to \mathbb{F}$ as follows: for every $y \in \mathbb{F}^m$, $f_x(y)$ is the most frequent $a_C(y)$ among all cubes such that,

1. $x, y \in C$.

2. $a_C(x) = \sigma = P_1(x)$.

Since the cubes $C \in \mathcal{C}$ are dimension 3 affine subspaces, for almost every $y \in \mathbb{F}^m$ there exist many cubes that satisfy the two requirement. If there are none, we define $f_x(y)$ arbitrarily.

In the general case, where the assignment a is not a restriction of P_1, P_2, we define f_x in the same way, but the analysis is more complicated. The main steps of the general case,

1. Two different degree d polynomials disagree on $1 - \frac{d}{|\mathbb{F}|}$ of their domain, so it is unlikely for a random $C_1, C_2 \in \mathcal{C}$ such that $C_1 \cap C_2 = \ell$ and $x \in \ell$ to satisfy both $a_{C_1}(x) = a_{C_2}(x)$ and $a_{C_1}|_\ell \neq a_{C_2}|_\ell$.

2. Two random cubes $C_1, C_2 \in \mathcal{C}$ that intersects on x satisfy $a_{C_1}(x) = a_{C_2}(x)$ w.p. ϵ.

3. For ϵ fraction of $x \in \mathbb{F}^m$, exists σ such that a function f_x defined when we require $a_c(x) = \sigma$ is good, i.e. $a_C \approx f_x|_\sigma$ for ϵ fraction of $C \ni x$, when $\approx$ means equal on almost all of the coordinates.

2.2 Global structure

The local structure gave us many good functions f_x, each satisfies $f_x|C \approx a_C$ for many cubes $C \ni x$. Our goal in this section if to find a single x such that its f_x is globally good, i.e. $f_x \approx a_C$ for ϵ fraction of **all cubes**, and not only of cubes containing x.

We do it by showing that for many $x, y \in \mathbb{F}^m$, $f_x \approx f_y$, in the lesson we don't go quantify what exactly $\approx$ means. Notice that since f_x, f_y are not degree d polynomials, they can be approximately equal without being equal.

Definition 2.1. *For every $x \in \mathbb{F}^m$ with a function f_x, let $F_x = \{C | x \in C, f_{x|_C} \approx a_C\}$.*

The global structure proof has two main steps.

1. For many $x, y \in \mathbb{F}^m$, $\Pr_C[C \in F_x \cap F_y | x, y \in C] \geq \Omega(\epsilon^2)$.

2. If $\Pr_C[C \in F_x \cap F_y | x, y \in C] \geq \Omega(\epsilon^2)$, then $f_x \approx f_y$.

3. There must be $x \in \mathbb{F}^m$ such that $f_x \approx f_y$ for many $y \in \mathbb{F}^m$.

In this lesson we will only see the second step, the first one is proven using similar technique and the third is implied by the first two.

Fix $x, y \in \mathbb{F}^m$ such that $\Pr_C[C \in F_x \cap F_y | x, y \in C] \geq \Omega(\epsilon^2)$, and let ℓ be the line connecting x and y. Let $G = (A \cup B, E)$ be the following bipartite graph, with $A = \mathbb{F}^m \setminus \ell$, $B = C \ni x, y$ and $(z, C) \in E$ if $z \in C$, see Figure 1.

The second largest eigenvalue of the graph G is $\lambda(G) = (1 + o(1))\frac{1}{\sqrt{q}}$ (since it is bipartite, $\lambda_0 = 1, \lambda_{n-1} = -1$), so G is a very good spectral expander. Bounding $\lambda(G)$ is done by a 2 step random walk on G, starting from $\mathbb{F}^m \setminus \ell$.

In the article, we show that for every subset $B' \subset B$ larger than $\lambda(G)^2$, the distribution of picking $b \in B'$ and a random neighbor $x \in A$, is close to the distribution of picking a uniform $x \in A$ and a random neighbor in B'.

Figure 1: The graph G

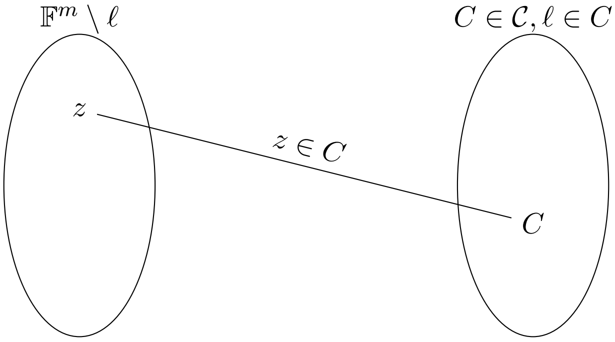
$\mathbb{F}^m \setminus \ell$
$C \in \mathcal{C}, \ell \in C$
z
$z \in C$
C

Figure 2: B' in G

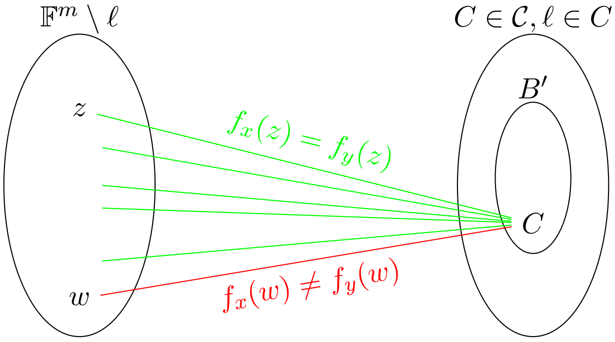
$\mathbb{F}^m \setminus \ell$
$C \in \mathcal{C}, \ell \in C$
B'
z
$f_x(z) = f_y(z)$
C
w
$f_x(w) \neq f_y(w)$

We prove item 2 by defining $B' = F_x \cap F_y$, which means that $|B'| \geq \epsilon^2|B| \geq \lambda(G)^2|B|$. For every $C \in B'$, by definition $a_C \approx f_{x|_C}$ and $a_C \approx f_{y|_C}$, this means that almost all of the outgoing edges from B', the edge (z, C) satisfies $f_x(z) = f_y(z)$, see Figure 2. From the expansion properties of G, we know that since $|B'| \geq \lambda(G)^2|B|$, the outgoing edges of B' covers almost all of A, so for almost all points $z \in A$, $f_x(z) = f_y(z)$ and we are done.

2.3 Low Degree

From the global structure we have a function $f_x : \mathbb{F}^m \to \mathbb{F}$ such that $f_x|_C \approx a_C$ for ϵ of the cubes $C \in \mathcal{C}$, and we want to show that f_x is close to a low degree polynomial. We do it by a reduction to Rubinfeld Sudan, we show that for almost all $d+2$ points on a line, f_x is a degree d polynomial on these points.

For every $d+2$ points $z_t = z + th$, $t \in \{0, \ldots d+1\}$, if there exist a cube $C \in \mathcal{C}$ such that $a_C(z_t) = f_x(z_t)$ for every t, then f_x is a degree d polynomial on these points (since a_C is a degree d polynomial).

We show that this happens with probability $1-\delta$ using an argument similar to the expansion of the global structure. We look on the bipartite graph $G' = (\mathcal{L} \cup \mathcal{C}, E')$, where the left side is all affine lines in $\mathbb{F}^m$, and the right side contains all cubes, and $(\ell, C) \in E'$ if $\ell \subset C$, see Figure 3.

Let $\mathcal{C}'$ be the set of cubes such that $f_x|_C \approx a_C$. This graph is also an expander, and has similar properties to G. This means that for almost all lines $\ell \in \mathcal{L}$, there exists a cube C such that $a_C|_\ell \approx f_x|_\ell$, so for almost all $d+2$ points on the line, f_x, a_C are equal on these points. Then, we use Rubinfeld Sudan theorem, Theorem 1.1, to conclude that f_x is close to a degree d polynomial.

Figure 3: G'

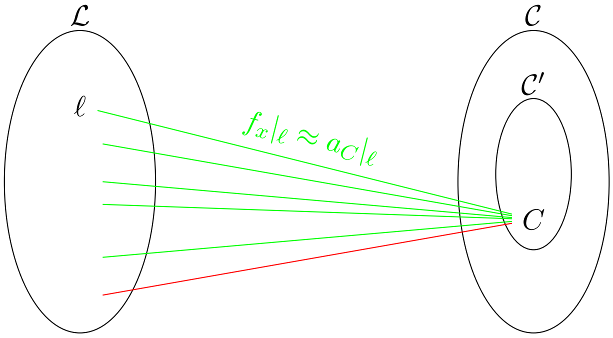

References

[AS97] Sanjeev Arora and Madhu Sudan. Improved low-degree testing and its applications. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 485–495. ACM, 1997.

[BDN17] Amey Bhangale, Irit Dinur, and Inbal Livni Navon. Cube vs. cube low degree test. In *Proceedings of the 2017 Conference on Innovations in Theoretical Computer Science, ITCS*, 2017.

[MR08] Dana Moshkovitz and Ran Raz. Sub-constant error low degree test of almost-linear size. *SIAM J. Computing*, 38(1):140–180, 2008.

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.

[RS97] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability pcp characterization of np. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 475–484. ACM, 1997.