CS 6815: Lecture 3

Instructor: Eshan Chattopadhyay Scribe: Linus Setiabrata

Aug 30, 2018

Announcements:
Homework 0 is due; Homework 1 will come out today.

On the week of September 10, Bobby Kleinberg will guest lecture on capsets, polynomial method, and Fourier analysis.

In the last lecture we were talking about PRGs. We had the pairwise independent generator $X = (X_1, \ldots, X_m)$ with $(X_i, X_j) \sim U_2$. We'll be generalizing this:

k-wise independent distributions

Let X be a distribution on $\{0,1\}^n$. We say X is a k-wise independent distribution if $S \subseteq [n], |S| = k$, we have $X_S \sim U_k$ (where $X_S = (X_i)_{i \in S}$). We'll construct one of these by first constructing one such X on $\mathbb{F}_q^n$, where q is large.

Indeed, consider a matrix M of size $n \times k$ such that any k rows are independent. If $q \geq n+1$, then one such matrix exists (take, for example, the Vandermonde matrix).

Randomly pick $y \in \mathbb{F}_q^k$ and output $My = x \in \mathbb{F}_q^n$. We want to show X is k-wise independent, that is, for $S \subseteq [n]$ and $|S| = k$ we have $X_S \sim U_{\mathbb{F}_q^k}$ [he remarks again that we are doing this trick to minimize the randomness needed]. But actually $x_S = M_S y$, where M_S is the matrix minor obtained by keeping only rows in S (and all the columns). But now M_S is invertible because it has k rows, and they're independent. So M_S a bijection, and X_S is now uniform! The randomness used was $k \log q$ (number of bits).

Exercise 1. *(Unofficial homework) Verify the fact that the Vandermonde matrix has the property that any k rows are independent. [my favorite proof is here, though this might be a standard proof by now]*

We use this to construct k-wise independence on $\mathbb{F}_2^n$. Repeat the above construction, for $s = \lceil \log(n+1) \rceil, q = 2^s$. We want to use this to construct a matrix, this time over $\mathbb{F}_2$, such that any k rows are independent. Recall that there is a natural map $\varphi \colon \mathbb{F}_{2^s} \to \mathbb{F}_2^s$. For example, you might recall that $\mathbb{F}_{2^s} \cong \mathbb{F}_2[x]/(x^s - 1)$, and φ will send $[f] \in \mathbb{F}_2[x]/(x^s-1)$ to the coefficients of f (where the representative f of $[f]$ is chosen so that it has degree at most $s-1$, of course...)

So this pushes to a map

$$\bar{\varphi} \colon \underbrace{M}_{n \times k} \mapsto \underbrace{M'}_{n \times (1+s(k-1))},$$

acting elementwise on M. We claim that M' also has the property that any k rows are linearly independent (prove the contrapositive.... why are you doing this). Here, too, the randomness used is $k \log q \approx k \log(n+1)$ [modulo floors/ceilings].

Improving seed length

The claim is that it suffices to use

$$M_1 = \begin{bmatrix} 1 & \alpha_1 & \alpha_1^3 & \alpha_1^5 & \dots & \alpha_1^{k-1} \\ 1 & \alpha_2 & \alpha_2^3 & \alpha_2^5 & \dots & \alpha_2^{k-1} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \end{bmatrix}$$

where we drop the columns with even powers. This will allow us to shave off some columns so that we can reduce randomness further. Indeed, we claim that any k rows are linearly independent over $\mathbb{F}_2$.

Proof. Say that there was a $T = \{t_1, \dots, t_k\} \subseteq [n]$. Recall that any k rows of M are linearly independent (over $\mathbb{F}_2$). Fix $\beta_1, \dots, \beta_k \in \mathbb{F}_2$ that are not all 0.

We have $\beta_1 (M)_{t_1} + \beta_2 (M)_{t_2} + \dots + \beta_k (M)_{t_k} \neq \vec{0}$, where $(M)_i$ denotes the ith row of M (written as a column vector, but never mind that). So for some ℓ, the ℓth component of this matrix is nonzero, ie., $\beta_1 \alpha_{t_1}^\ell + \beta_2 \alpha_{t_2}^\ell + \dots + \beta_k \alpha_{t_k}^\ell \neq 0$.

Observe that if, for example, $\alpha_1^{12} + \alpha_3^{12} + \alpha_7^{12} = (\alpha_1^3 + \alpha_3^3 + \alpha_7^3)^4 \neq 0$ (recall that this field has characteristic 2, and doesn't contain any nilpotents), and so $\alpha_1^3 + \alpha_3^3 + \alpha_7^3 \neq 0$.

In general, let $\ell = 2^a b$ with b odd. Whenever $\ell \neq 0$, we can write ℓ uniquely in this way; assume $\ell \neq 0$ for now and we'll treat the special case later. Before, we had concluded that $\beta_1 \alpha_{t_1}^\ell + \beta_2 \alpha_{t_2}^\ell + \dots + \beta_k \alpha_{t_k}^\ell \neq 0$. Observe now that $\beta_1 \alpha_{t_1}^\ell + \beta_2 \alpha_{t_2}^\ell + \dots + \beta_k \alpha_{t_k}^\ell = (\beta_1 \alpha_{t_1}^b + \beta_2 \alpha_{t_2}^b + \dots + \beta_k \alpha_{t_k}^b)^{2^a} \neq 0$. Again, fields don't have nilpotents, so $\beta_1 \alpha_{t_1}^b + \dots + \beta_k \alpha_{t_k}^b \neq 0$, and b is odd.

This argument works for any β_i not all 0, so it follows that any k rows of M_1 are linearly independent; we've proven that if $T = \{t_1, \dots, t_k\} \subseteq [n]$ and $\beta_1, \dots, \beta_k \in \mathbb{F}_2$ are not all 0, there is an odd number b such that $\beta_1 \alpha_{t_1}^b + \dots + \beta_k \alpha_{t_k}^b \neq 0$.

We shouldn't forget about the case where $\ell = 0$. But in that case, it is already a column of M_1. [In particular, I don't see right now how to get rid of the column of 1's. I think in class we asked whether one could get rid of them to save a random bit] □

Let's switch gears:

Almost k-wise independence and small biased distributions

There are two definitions.

There's an L_∞ definition: Let X be a random variable on $\{0,1\}^n$. Let $T \subseteq [n]$, with $|T| = k$. We say that X is almost k-wise independent if for a fixed parameter $\varepsilon > 0$ (where ε quantifies "almost"), we have for all all $\alpha \in \{0,1\}^k$, the inequality

$$\left| \Pr[X_T = \alpha] - \frac{1}{2^k} \right| \leq \varepsilon.$$

There's also an L_1 definition: For all $T \subseteq [n]$, with $|T| = k$, we say X is almost k-wise independent if for a fixed parameter $\varepsilon > 0$ (quantifying "almost"), we have for all

$$\sum_{\alpha \in \{0,1\}^k} \left| \Pr[X_T = \alpha] - \frac{1}{2^k} \right| \leq \varepsilon.$$

There is the usual trickery: if you are ε-almost k-wise independent in the L_1 sense, then you are ε-almost k-wise independent in the L_∞ sense. For a partial converse, if you are ε-almost k-wise independent in the L_∞ sense, then you are $2^k\varepsilon$-almost k-wise independent in the L_1 sense.

We also define ε-biased distribution in the following way:

Let X be a distribution on $\{0,1\}^n$, and $T \subseteq [n]$. Define $\oplus X_T = \sum_{i \in T} X_i \pmod 2$ to be the parity of X_T.

Define

$$\text{Bias}(\oplus X_T) = \left| \Pr[\oplus X_i = 1] - \Pr[\oplus X_i = 0] \right|.$$

Now we say that X is a ε-biased distribution if for all nonempty subsets $T \subseteq [n]$, we have $\text{Bias}(\oplus X_T) \leq \varepsilon$.

Notice that we have

$$\left| \mathbb{E}[\oplus U_T] - \mathbb{E}[\oplus X_T] \right| = \left| 1/2 - \mathbb{E}[\oplus X_T] \right|$$

so if X is a ε-biased distribution then $|\Pr[\oplus X_i = 1] - \Pr[\oplus X_i = 0]| \leq \varepsilon$ and $(1-\varepsilon)/2 \leq \Pr[\oplus X_i = 1] \leq (1+\varepsilon)/2$ (there's a boring computation here that I'm hiding). Thus $(1-\varepsilon)/2 \leq \mathbb{E}[\oplus X_T] \leq (1+\varepsilon)/2$ and $|1/2 - \mathbb{E}[\oplus X_T]| \leq \varepsilon/2$. So ε-biased distributions are pseudoranom generators for the class of parity functions.

Construction of ε-biased spaces

Let $r = \lceil \log_2(n/\varepsilon) \rceil$, and let $q = 2^r \approx n/\varepsilon$.

Pick a random $y, z \in \mathbb{F}_q$. For $i \in \{0, 1, \ldots, r-1\}$, we map $y \to \vec{y}$ as a vector in $\mathbb{F}_2^r$ (the same trick as above; think of $\mathbb{F}_{2^r}$ as polynomials and map to the coefficients), and $z \mapsto \vec{z}$. Define $x_i = < y^i, z >$, where y^i is a product first taken in $\mathbb{F}_q$ and then interpreted as a vector in $\mathbb{F}_2^r$. Now consider $X = (x_0, x_1, \ldots, x_{r-1})$. We won't prove that this is a ε-biased space (that'll happen next lecture), but we can say that the randomness used is $2\log(n/\varepsilon)$.