Computer Science 294 Lecture 13 Notes

Daniel Raban

February 28, 2023

1 Pseudorandomness

1.1 Pseudorandom distributions and generators

Suppose we have some function f, which we can think of as a randomized algorithm. We give f n uniformly random bits, and it takes in an input x. Now we want to consider another situation in which we give f n bits which are not uniformly random but may have some dependence. We want these bits to be **pseudorandom** in the sense that the output of f on input x should be similar to if we were using uniformly random bits.

Definition 1.1. A distribution $\mathcal{D}$ on $\{\pm 1\}^n$ is **ε-pseudorandom** against a class $\mathcal{C}$ of tests (or **ε-fools $\mathcal{C}$**) if for any $f \in \mathcal{C}$,

$$|\mathbb{E}_{Y \sim U_n}[f(Y)] - \mathbb{E}_{X \sim \mathcal{D}}[f(X)]| \leq \varepsilon.$$

The uniform distribution is pseudorandom by this definition, but we really care about pseudorandom distributions which depend on some small random seed of s bits.

Definition 1.2. A **pseudorandom generator** against $\mathcal{C}$ with error ε and seed-length s is an explicit (deterministic) function $F : \{\pm 1\}^s \to \{\pm 1\}^n$ such that for all $f \in \mathcal{C}$,

$$|\mathbb{E}_{Y \sim U_n}[f(Y)] - \mathbb{E}_{X \sim U_s}[f(X)]| \leq \varepsilon.$$

In other words, we want $G(U_s)$ to be ε-pseudorandom against $\mathcal{C}$.
Today we will care about two classes of functions:

- k-juntas: These correspond to k-wise independent distributions.
- Parity tests: These correspond to ε-biased distributions.

1.2 Small-biased and (k, δ)-wise independent distribtions

Here is a definition due to Naor and Naor.

Definition 1.3. A distribution $\mathcal{D}$ over $\{\pm 1\}^n$ is called ε**-biased** if for all nonempty sets $S \subseteq [n]$,

$$|\mathbb{E}_{X\sim\mathcal{D}}[\chi_S(X)]| \leq \varepsilon.$$

Equivalently,

$$\frac{1-\varepsilon}{2} \leq \mathbb{P}_{X\sim\mathcal{D}}\left(\prod_{i\in S} X_i = 1\right) \leq \frac{1+\varepsilon}{2}.$$

Definition 1.4. A distribution $\mathcal{D}$ is (k, δ)**-independent** if for all k-juntas $f : \{\pm 1\}^n \to \{\pm 1\}$,

$$|\mathbb{E}_{Y\sim U_n}[f(Y)] - \mathbb{E}_{X\sim\mathcal{D}}[f(X)]| \leq \delta.$$

Remark 1.1. $(k, 0)$ independence is sometimes called k-wise independence, which can be sampled using $O(k \log n)$ seed-length. This is optimal, as it can be shown that k-wise independence can't be sampled with $< \frac{k \log n}{2}$ bits.

Our goal today is to prove the following theorem.

Theorem 1.1 (Naor-Naor)**.** *There is an explicit $G : \{\pm 1\}^s \to \{\pm 1\}^n$ such that $G(U_s)$ is a ε-biased distribution and $s = O(\log n + \log(1/\varepsilon))$.*

We will be following the proof of AGHP, rather than the original proof. The following lemma will help us connect the ideas of small-biased and (k, δ)-independent distributions.

Lemma 1.1. *Suppose $\mathcal{D}$ is a ε-biased distribution. Then*

$$|\mathbb{E}_{X\sim\mathcal{D}}[f(X)] - \mathbb{E}_{Y\sim U_n}[f(Y)]| \leq \varepsilon L_1(f),$$

where

$$L_1(f) = \sum_{S\subseteq[n]} |\widehat{f}(S)|.$$

Proof of Lemma. Use the Fourier expansion for both expectations.

$$\begin{aligned}
|\mathbb{E}_{X\sim\mathcal{D}}[f(X)] - \mathbb{E}_{Y\sim U_n}[f(Y)]| &= \left|\sum_{S\subseteq[n]} \widehat{f}(S)(\mathbb{E}_{X\sim\mathcal{D}}[\chi_S(X)] - \mathbb{E}_{Y\sim U_n}[\chi_S(Y)])\right| \\
&= \left|\sum_{\varnothing\neq S\subseteq[n]} \widehat{f}(S)\,\mathbb{E}_{X\sim\mathcal{D}}[\chi_S(X)]\right| \\
&\leq \sum_{\varnothing\neq S\subseteq[n]} |\widehat{f}(S)| \cdot |\mathbb{E}_{X\sim\mathcal{D}}[\chi_S(X)]| \\
&\leq \varepsilon.
\end{aligned}$$

□

Corollary 1.1 (Vazirani's XOR Lemma)**.** *Any ε-biased distribution is also a (k, δ)-wise distribution with $\delta = \varepsilon 2^{k/2}$.*

Before proving this corollary, let's see what it, along with the Naor-Naor theorem, tells us.

Corollary 1.2. *There is an explicit $G : \{\pm 1\}^s \to \{\pm 1\}^n$ such that $G(U_s)$ is a (k, δ)-wise independent distribution and $s = O(\log n + \log(1/\varepsilon))$.*

Proof of Corollary. Just apply the Naor-Naor theorem with $\varepsilon = \delta/2^{k/2}$ and then apply Vazirani's XOR lemma. Then

$$s = O(\log n + \log(1/\varepsilon)) = O(\log n + \log(1/\delta) + k/2). \qquad \square$$

Remark 1.2. This can be improved to finding a (k, δ)-wise independent distribution with seed-length $O(k + \log\log n + \log(1/\delta))$.

Proof of Vazirani's XOR Lemma. Let $g : \{\pm 1\}^n \to \{\pm 1\}$ be any k-junta. So $g(x) = h(x_{i_1}, \ldots, x_{i_k})$ for some $i_1 < \cdots < i_k$ and $h : \{\pm 1\}^k \to \{\pm 1\}$. Then

$$\begin{aligned} L_1(g) &= \sum_{S \subseteq [n]} |\widehat{g}(S)| \\ &= \sum_{S' \subseteq [k]} |\widehat{h}(S)| \end{aligned}$$

Using Cauchy-Schwarz,

$$\begin{aligned} &\leq \sqrt{\sum_{S' \subseteq [k]} \widehat{h}(S')^2 \cdot 2^k} \\ &= \sqrt{2^k} \\ &= 2^{k/2}. \end{aligned}$$

Now apply the previous lemma. $\square$

1.3 Small-biased distributions and error-correcting codes

The intuition for our construction for the Naor-Naor theorem will be based on composing the Reed-Solomon and Hadamard error-correcting codes.

Definition 1.5. A multiset $A \subseteq \{\pm 1\}^n$ is called an **ε-biased set** if the uniform distribution over A, denoted U_A, is ε-biased.

Think of this as a long matrix with $m = |A|$ rows and n columns. We want the XOR of every subset of columns to be $(1/2 \pm \varepsilon)$-balanced. In $\mathbb{F}_2$ notation, we want any column

spanned by the columns of this matrix to have $(1/2 \pm \varepsilon)m$ 1s. This is saying that this is an error-correcting code over $\mathbb{F}_2^m$ with dimension n and distance $\geq m(1/2 - \varepsilon)$.

Pick a random multiset $A \subseteq \{\pm 1\}^n$ by picking $m = O(n/\varepsilon^2)$ elements independently and uniformly at random $a^{(1)}, a^{(2)}, \ldots, a^{(m)} \in \{\pm 1\}^n$. Then $A = \{a^{(1)}, \ldots, a^{(m)}\}$ is an ε-biased set with high probability. This is because of the following claim:

Proposition 1.1. *For all nonempty $S \subseteq [n]$,*

$$\mathbb{P}_{a^{(1)},\ldots,a^{(m)}} \left(\left| \frac{1}{m} \sum_{i=1}^{m} \chi_S(a^{(i)}) \right| > \varepsilon \right) \leq 2e^{-m\varepsilon^2/2}.$$

Proof. This follows from the Chernoff bound, since the $a^{(i)}$ are independent. □

If we pick $m = 2n/\varepsilon^2$, then this probability is $\ll 2^{-n}$. This however, is non-constructive, so we will use a constructive argument.

1.4 AGHP construction of a pseudorandom generator for small-biased distributions

AGHP gave an explicit function $G : \{\pm 1\}^s \to \{\pm 1\}^n$ such that $G(U_s)$ is ε-biased and $s = 2\log_2(n/\varepsilon)$.

Proof of Naor-Naor theorem. Here is the construction: Take $\ell = \lceil \log_2(n/\varepsilon) \rceil$. We will define $G : \{\pm 1\}^{2\ell} \to \{\pm 1\}^n$. Identify the first ℓ bits with an element $x \in \mathbb{F}_{2^\ell}$ and the second ℓ bits with an element $y \in (\mathbb{F}_2)^\ell$. Let $\mathrm{bin} : \mathbb{F}_{2^\ell} \to (\mathbb{F}_2)^\ell$ be 1 to 1 and linear over $\mathbb{F}_2$; that is,

$$\mathrm{bin}(x + y) = \mathrm{bin}(x) \oplus \mathrm{bin}(y), \qquad \mathrm{bin}(0) = 0^\ell.$$

Sample $X \sim \mathbb{F}_{2^\ell}$ uniformly at random and $Y \sim (\mathbb{F}_2)^\ell$ uniformly at random. Output $(Z_0, \ldots, Z_{n-1})$, where

$$Z_i = \langle \mathrm{bin}(X^i), Y \rangle_2, \qquad \langle a, b \rangle_2 = \sum_{j=1}^{\ell} a_j b_j \pmod 2.$$

Here is the analysis: Let $\alpha \in \{0,1\}^n$ be nonzero. Then it suffices to show that $\mathbb{E}_Z[(-1)^{\sum_{i=0}^{n-1} \alpha_i Z_i}] \leq \varepsilon$.

$$\mathbb{E}_Z \left[(-1)^{\sum_{i=0}^{n-1} \alpha_i Z_i} \right] = \mathbb{E}_{X,Y} \left[(-1)^{\sum_{i=0}^{n-1} \alpha_i \langle \mathrm{bin}(X^i), Y \rangle_2} \right]$$

By the linearity of the inner product,

$$= \mathbb{E}_{X,Y} \left[(-1)^{\langle \sum_{i=0}^{n-1} \alpha_i \mathrm{bin}(X^i), Y \rangle_2} \right]$$

By the linearity of bin,

$$= \mathbb{E}_{X,Y}\left[(-1)^{\langle \mathrm{bin}(\sum_{i=0}^{n-1}\alpha_i X^i), Y\rangle_2}\right]$$

$P_\alpha(t) = \sum_{i=0}^{n-1}\alpha_i t^i$ is a polynomial in one variable over $\mathbb{F}_{2^\ell}$, so it has at most $n-1$ roots.

$$= \mathbb{E}_X\left[\mathbb{E}_Y\left[(-1)^{\langle \mathrm{bin}(P_\alpha(X)), Y\rangle_2}\right]\right]$$

For a fixed x, if $P_\alpha(x) = 0$, then $\mathbb{E}_Y[(-1)^{\langle P_\alpha(x), Y\rangle_2}] = 1$. Otherwise, $\mathbb{E}_Y[(-1)^{\langle P_\alpha(x), Y\rangle_2}] = 0$.

$$\begin{aligned} &= \mathbb{P}_X(P_\alpha(X) = 0) \\ &\leq \frac{n-1}{2^\ell} \\ &\leq \varepsilon. \end{aligned}$$

□