Abstract Algebra in Theoretical Computer Science

Assignment 4

Lecturer: Gil Cohen

Problem 1 - Mergers without the uniformity assumption

In class we showed how to merge two (possibily correlated) n-bit random variables, one of which is uniform, to a $(1-\alpha,\varepsilon)$-random source. The seed length required for the merging process is $d = O(\frac{1}{\alpha} \cdot \log(n/\varepsilon))$. In this question you are asked to generalize the result and show how to merger two random variables even if the "good" one is not necessarily uniform but rather has some amount of entropy in it. The goal then is to "preserve" the entropy of the good source in the output.

Formally, devise an algorithm $M \colon \{0,1\}^n \times \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^n$, where d is as above, with the following property. Let X_1, X_2 be a pair of (possibly correlated) n-bit random variables. Assume that one of the X_is is uniformly distributed over a set $S \subseteq \{0,1\}^n$ of size $|S| \geq 2^{\gamma n}$ for some constant $0 < \gamma < 1$ (the identity of who this X_i is as well as the identitiy of the set S is unknow to you. Further, you do not know γ). Let Y be a random variable that is uniformly distributed over d-bit strings independently of (X_1, X_2). The property is that, under these assumptions, $M(X_1, X_2, Y)$ is $((1-\alpha)\gamma, \varepsilon)$-random.

Problem 2 - Secret sharing schemes

The goal of this problem is to construct an important primitive in cryptography called a *secret sharing scheme*. We will construct an elegant scheme due to Adi Shamir.

Assume you have a "secret" in the form of, say, an m-bit string s. Given integer parameters $1 \leq k < n$ we wish to divide the secret to n pieces $S_1, \ldots, S_n$, called *shares*. This division is going to be done using some randomness and so $S_1, \ldots, S_n$ are in fact random variables that are functions of the secret s (this is why we write them in capital). We want that:

- Knowning any k (or more) of the shares, one can efficiently reconstruct the secret s.
- Knowning less than k shares give no information whatsoever about s in the following sense. If $t < k$ and $i_1, \ldots, i_t \in \{1, 2, \ldots, n\}$ then given that $S_{i_1} = s_{i_1}, \ldots S_{i_t} = s_{i_t}$, the secret s can be any m-bit string with equal probability (i.e, 2^{-m}).

Shamir's construction is as follows. Let $\mathbb{F}$ be the field of 2^m elements, $2^m > n$. Sample $k-1$ elements $a_1, \ldots, a_{k-1}$ of $\mathbb{F}$ uniformly at random. Set $a_0 = s$, and consider the polynomial $f(x) \in \mathbb{F}[x]$ that is given by $f(x) = \sum_{i=0}^{k-1} a_i x^i$. Let $\alpha_1, \ldots, \alpha_n$ be (arbitrarily chosen) distinct nonzero elements of $\mathbb{F}$ (here we use that $2^m > n$). The n shares are then $S_i = (\alpha_i, f(\alpha_i))$ for $i = 1, \ldots, n$.

Prove the correctness of Shamir's scheme.

Problem 3 - The Schwartz-Zippel Lemma

Prove the Schwartz-Zippel lemma that was presented in class. That is, prove that any nonzero polynomial $f(x_1, \ldots, x_m) \in \mathbb{F}_q[x_1, \ldots, x_m]$ of total degree d has at most dq^{m-1} roots in $\mathbb{F}_q^m$. Hint: Use induction on the number of variables.

Problem 4 - Small-bias sets

For an integer $m \geq 1$ define the function $\mathrm{Tr}\colon \mathbb{F}_{2^\ell} \to \mathbb{F}_{2^\ell}$ by

$$\mathrm{Tr}(x) = x^{2^0} + x^{2^1} + \cdots + x^{2^{\ell-1}}$$

1. Prove that, in fact, the function Tr maps $\mathbb{F}_{2^\ell}$ to $\mathbb{F}_2$.
2. Prove that for every nonzero element $a \in \mathbb{F}_{2^\ell}$, the function $f\colon \mathbb{F}_{2^\ell} \to \mathbb{F}_2$ that is given by $f(x) = \mathrm{Tr}(ax)$ is an $\mathbb{F}_2$-linear function and that $\mathbf{E}_{x \sim \mathbb{F}_{2^\ell}}[f(x)] = 0$.

Consider the following construction of a set $S \subseteq \{0,1\}^n$. The elements of S are indexed by a triplet of elements $x, y, z \in \mathbb{F}_{2^\ell}$ for a parameter ℓ to be chosen later on. We index each entry of the string $s(x,y,z)$ by a pair of numbers $0 \leq i, j \leq c\sqrt{n}$ where the constant c is chosen so that there are n entries. For such x, y, z and i, j, we define $s(x,y,z)_{i,j} = \mathrm{Tr}(x^i y^j z)$.

3. Let n be as above and $\varepsilon > 0$. Show that for a proper choice of ℓ, the set S described above is an ε-biased set of size $O(n\sqrt{n}/\varepsilon^3)$.