6.841: Advanced Complexity Theory Fall 2012

Problem Set 4 – Pseudorandomness

Prof. Dana Moshkovitz/TA: Henry Yuen *Due Date: November 15, 2012*

Turn in your solution to each problem on a separate piece of paper. Mark the top of each sheet with the following: (1) your name, (2) the question number, (3) the names of any people you worked with on the problem. We encourage you to spend time on each problem individually before collaborating!

Problem 1 – Equivalence of pseudorandomness and hardness

Recall the following definition from class.

Definition 1. *A pseudorandom generator $G : \{0,1\}^s \to \{0,1\}^n$ is ϵ-pseudorandom against circuits of size T iff for all circuits C on n-bits of size at most T:*

$$\left| \Pr_{x \sim \{0,1\}^s}[C(G(x)) = 1] - \Pr_{y \sim \{0,1\}^n}[C(y) = 1] \right| \leq \epsilon.$$

You saw in lecture that *hardness* implies *pseudorandomness*: under the assumption that there are functions that are hard to compute on average, there are pseudorandom generators that can be used to derandomize BPP to P. In this problem, you will see that the converse is also true: *pseudorandomness* implies *hardness*.

Let $\alpha \in (0,1)$. Show that if there exists a $2^{\alpha n}$-time computable pseudorandom generator $G : \{0,1\}^s \to \{0,1\}^n$ that is $(1/3)$-pseudorandom against circuits of size $T < 2^{\alpha n}$, then there exists a $2^{O(n)}$-time computable function $f : \{0,1\}^{O(n)} \to \{0,1\}$ that cannot be computed by circuits of size T. Here, assume that $s = O(\log n)$.

Problem 2 – Constant Sized Expanders

Recall that for an undirected graph G, $\lambda(G)$ is the second largest eigenvalue (in absolute value) of the random walk matrix $A(G)$.

Let $\mathbb{F}$ be a finite field. Consider the following graph G: its vertex set is $\mathbb{F}^2$, and its edge set consists of all pairs $((a,b),(c,d))$ such that $ac = b + d$. Equivalently, we connect every vertex (a,b) to all points on the line $y = ax - b$.

Prove that G is $|\mathbb{F}|$-regular and that $\lambda(G) \leq 1/\sqrt{|\mathbb{F}|}$.

[Note: In the SL = L analysis, one starts with a constant-sized expander and uses a combination of the tensor product and the replacement product to produce larger expanders (to say, build the expander graph H that's used in the replacement product). This problem gives one easy way of producing an initial graph.]

Problem 3 – Expander graph pseudorandomness and ExpanderVille dynamics

(a) An important feature of expander graphs is that they look very much like randomly generated graphs. The following result captures this notion. Suppose G is a d-regular graph on vertex set V, edge set E, and

has second eigenvalue $\lambda(G)$. Then, the following holds: for all sets $A, B \subseteq V$,

$$\left| e(A,B) - d \cdot \frac{|A| \cdot |B|}{n} \right| \leq \lambda \cdot d \cdot \sqrt{|A| \cdot |B|},$$

where $e(A,B) = |\{(u,v) \mid u \in A, v \in B\}|$. What this says is for all sets A, B, the deviation of $e(A,B)$ from what you would expect with a random graph is small.

Prove this result.

[Hint: Show that you can express $e(A,B) = \chi_A^\top M \chi_B$, where M is the adjacency matrix for G, and $\chi_A(v) = 1$ iff $v \in A$. Then, like in the analysis of Lemma 21.3 in the book, split χ_A and χ_B into a uniform part and an orthogonal part.]

(b) Now let's see an application of this result you just proved. Imagine the following situation: there are group of n people who live on a d-regular expander graph (ExpanderVille) with n vertices and eigenvalue $\lambda < 1/13$. There's a group $S \subseteq [n]$ of the people who are diehard iPhone users, and the rest are ambivalent fans of Android. The iPhone fans are extremely persuasive: at each time step, anybody who has at least $d/3$ iPhone fans as neighbors will immediately purchase an iPhone herself. Otherwise, if less than $d/3$ neighbors are iPhone fans, people will succumb to peer pressure and become an Android user instead. However, the iPhone camp at the beginning is small – $|S| < n/4$. Show that, despite their persuasive powers, after $O(\log n)$ time steps there will be no more iPhone users in ExpanderVille.

Problem 4 – Introduction to Fourier Analysis

Consider the mapping $0 \to 1$ and $1 \to -1$: this allows us to write Boolean functions $f : \{0,1\}^n \to \{0,1\}$ as real-valued functions $f : \{-1,1\}^n \to \{-1,1\}$ instead. Why is this useful? It turns out that functions of the form $f : \{-1,1\}^n \to \mathbb{R}$ can be analyzed via the *Fourier Transform over the Boolean hypercube*. It is a theorem that every such f can be written as

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} x_i,$$

where the $\widehat{f}(S)$ are real, and are called the *Fourier coefficients* of f. This is called the *Fourier expansion* of f. Let $\chi_S(x)$ denote $\prod_{i \in S} x_i$ (by convention $\chi_\emptyset(x) = 1$). For any two functions $f, g : \{-1,1\}^n \to \mathbb{R}$, define their inner product $\langle f, g \rangle = \mathbb{E}_{x \in \{-1,1\}^n}[f(x)g(x)]$.

You may find Chapter Section 22.5 in the textbook useful.

(a). Show that (i) for all S, S', $\langle \chi_S, \chi_{S'} \rangle = \delta_{S,S'}$ (which is equal to 1 iff $S = S'$, 0 otherwise), and (ii) $\langle \chi_S, f \rangle = \widehat{f}(S)$.

(b). Prove that for all $S \neq S'$, $\chi_S(x) \neq \chi'_S(x)$ on exactly 2^{n-1} x's.

(c). Define the convolution of two functions $f, g : \{-1,1\}^n \to \mathbb{R}$ to be $(f * g)(x) = \mathbb{E}_{y \in \{-1,1\}^n}[f(y)g(x \cdot y)]$, where $x \cdot y$ denotes entry-wise multiplication. Prove that $\widehat{f * g}(S) = \widehat{f}(S)\widehat{g}(S)$ for all $S \subseteq [n]$. You may recognize this as the Convolution Theorem from classical Fourier analysis.