MATH 642:587–Fall, 2011
Assignment 2(Version: November 21, 2011)

A few preliminary remarks.

1. Follow the general instructions for homework given in:

   http://www.math.rutgers.edu/∼saks/homework.html

2. Most of these problems are based on results that have appeared in research papers. Problems typically are broken into multiple steps that guide you through the solution.

3. For some problems hints are provided at the end of the assignment. This is noted by the statement "Hint given below." I recommend you try to solve the problem without the hint and refer to the hint if you are stuck. If you use a hint you should acknowledge using it in your acknwowledgements for the problem (as described in the general instructions for homework mentioned above.

4. Please be on the look out for errors in the problem statements. If something seems not to make sense, check with me before investing a lot of time on the problem. I would appreciate being notified of any typos (even minor ones).

PROBLEMS

1. The purpose of this problem is to prove the following bipartite graph packing theorem of Bollobas and Eldridge. Let G, H be (U, W)-bipartite graphs with $|U| = |W| = n$. If $d^U_{\max}(G)d^W_{\max}(H) + d^U_{\max}(H)d^W_{\max}(G) < n$ then the G_1, G_2 can be packed into the complete bipartite graph $K_{U,W}$. For permutations σ of U and τ of V, write $G[\sigma, \tau]$ be the (U, W)-graph with $uw \in G[\sigma, \tau]$ iff $\sigma(u)\tau(w) \in G$. Note that G, H pack iff there are permutations σ, τ such that $G[\sigma, \tau]$ is disjoint from H. Also, a $G - H$ path from vertex u to u' is a two edge path, whose first edge is in G and whose second edge is in H.

   (a) Consider a pair of permutations σ, τ that minimizes the number of common edges of $G' = G[\sigma, \tau]$ and H. Suppose $uw \in G' \cap H$. Show that for every vertex $u' \in U - u$, there is either a $G' - H$ path from u to u' or a $H - G'$ path from u to u'.

   (b) Prove the theorem.

2. Let $\Delta(G)$ be the maximum degree of G and $\bar{d}(G)$ be the average degree. In class we mentioned the following theorem of graph packing theorem of Sauer and Spencer (which is analogous to the bipartite theorem proved in the previous problem). Let G, H be n-vertex graphs. If $\Delta(G)\Delta(H) < n/2$ then G and H can be packed in K_n. Assuming this theorem (if it helps) prove: Suppose G and H are graphs on n vertices and suppose G has at most εn edges where $\varepsilon < 1/2$. Prove that if $\Delta(G)\bar{d}(H) < \varepsilon(1 - 2\varepsilon)n$ then G and H can be packed in K_n. (Hint provided below)

3. In class we talked about 0-error randomized decision trees. This problem deals with randomized decision trees that are allowed to make errors (2-sided error randomized decision trees). First, recall that a 0-error randomized decision tree for f is a probability distribution $\tilde{T}$ over decision trees that each compute f. The (worst case expected) cost of $\tilde{T}$, $cost(\tilde{T})$ is the maximum over all inputs x of the expected number of variables read on input x. Define $R(f)$ to be the minimum cost of any 0-error tree for f.

   A ε-error randomized decision tree for a function f is a probability distribution $\tilde{T}$ over decision trees with the property that for every input x, $\mathbf{Prob}[\tilde{\mathrm{T}}(\mathrm{x}) \neq \mathrm{f(x)}] \leq \varepsilon$. (Note: we do not require that the trees in the support of $\tilde{T}$ individualy compute f). We say $\tilde{T}$ has depth d if every tree in the support of the distribution has depth d. Define $R_\varepsilon(f)$ to be the smallest depth of a ε-error randomized decision tree for f.

   (a) Prove that $R_\varepsilon(f) \leq R(f)/\varepsilon$.

   (b) We say a polynomial $p(x_1, \ldots, x_n)$ ε-approximates f if $p(x) \in [0,1]$ and for all $x \in \{0,1\}^n$, $|p(x) - f(x)| \leq \varepsilon$. The ε-approximate degree of f, $\deg_\varepsilon(f)$ is the minimum degree of a polynomial that ε-approximates f. Show that $\deg_\varepsilon(f) \leq R_\varepsilon(f)$.

   (c) Prove that $R_\varepsilon(f) \geq bs(f)(1 - 2\varepsilon)$.

4. Recall that a partial assignment α for boolean variables $x_1, \ldots, x_n$ is a mapping from $[n]$ to $\{0, 1, *\}$. $Fixed_0(\alpha) = \{j \in [n] : \alpha_j - 0\}$, $Fixed_1(\alpha) = \{j \in [n] : \alpha_j = 1\}$, $Fixed(\alpha) = Fixed_0(\alpha) \cup Fixed_1(\alpha)$, and $Unfixed(\alpha) = [n] - Fixed(\alpha)$. We associate to α the subcube $C(\alpha) = \{x \in \{0,1\}^n : \forall i \in Fixed(\alpha), x_i = \alpha_i\}$. A set P of partial assignments is a *subcube partition* if every point in $\{0,1\}^n$ belongs to $C(\alpha)$ for a unique $\alpha \in P$. For a subcube partition and $x \in \{0,1\}^n$ we denote by $\alpha(x) = \alpha_P(x)$ the unique $\alpha \in P$ containing x. For $i \in \{0,1\}$, let $F_i(x)$ for the set $Fixed_i(\alpha(x))$ for $i \in \{0,1\}$.

   Fix a subcube partition P. Let $p \in (0,1)$ and consider the distribution on $\{0,1\}^n$ in which each variable is set to 1 with probability p. For an input X chosen according to this distribution, prove that $\mathbb{E}[|F_1(X)|]/\mathbb{E}[|F_0(X)|] = p/(1-p)$. (Hint provided below.)

5. The purpose of this problem is to prove the following theorem of Friedgut, Kahn, Wigderson: Fix a monotone graph property $\mathcal{P}$ on n-vertex graphs. let p be the critical probability,i.e. the unique $p \in (0.1)$ such that the random graph $G_{n,p}$ has $\mathcal{P}$ with probability exactly 1/2. Theorem: $R(\mathcal{P})$ is at least $n^2/A(n,p)$ where $A(n,p) = C\max(p(1-p)n, \log(n))$, for a suitably large constant C. (Observe that this is not such a good bound for $p = \Theta(1)$, but is often quite good. For example (not to hand in) derive a lower bound on $R(\mathcal{P})$ in the case that $\mathcal{P}$ is "contains a triangle")

   Consider the random graph $G_{n,p}$. Fix a deterministic decision tree T. By Yao's minimax theorem it suffices to prove that the expected number of variables read by T on $G_{n,p}$ is at least $n^2/A(n,p)$.

   (a) In what follows we assume $p \leq 1/2$. Why is this without loss of generality?

   (b) Let Y_0, Y_1 be random variables where Y_0 is the number of edges that were asked and received a "No" answer, and Y_1 is the number of edges asked that received a "Yes" answer. Fix a contant $K > 0$. Prove that if $\mathbb{E}[Y_1] \geq n/K$ then the desired lower bound holds. (Hint: Use the result of problem 4.)

(c) Prove that there is a graph G that satisfies $\mathcal{P}$ having maximum degree at most $C(np + \log(n))$ for a suitable constant $C > 0$ and at most $4\mathbb{E}[Y_1]$ edges. (Hint provided below).

(d) For a suitable constant $K > 0$, prove that if $\mathbb{E}[Y_1] \leq n/K$ then any 0-certificate H for $\mathcal{P}$, $|H| \geq n^2/K(np + \log(n))$. (Hint: Use the result of Problem 2).

(e) Prove the theorem.

6. Let $f : \{T, F\}^n \longrightarrow \mathbb{R}$. The fourier degree of f is the maximum J such that $\hat{f}(S) \neq 0$. (Not to hand in: the fourier degree of f does not depend on the product probability distribution, and is the same as the usual degree of f.) We write $deg(f)$ for the degree. In this problem we consider only unbiased fourier analysis (with respect to the uniform distribution over $\{T, F\}^n$.

    (a) Prove: if $deg(f) \leq d$ and f is not identically 0 then $f(\alpha) = 0$ for at most $2^n - 2^{n-d}$ inputs. (Hint provided).

    (b) Prove: if $deg(f) \leq d$ and $Range(f) \in [-1, 1]$ then $Inf(f)$ (the total influence, which is $\sum_i Inf_i(f)$) is at most d.

    (c) Prove: if $deg(f) \leq d$ and $Range(f) \in \{-1, 1\}$ then any relevant variable has influence at least 2^{-d}.

    (d) Prove: If $deg(f) \leq d$ and $Range(f) \in \{-1.1\}$ then f has at most $d2^d$ relevant variables.

    (e) Show that for each d there is a function of degree d that has $2^{d-1} + d - 1$ relevant variables.

7. In this problem, we consider only unbiased fourier analysis (with respect to the uniform distribution over $\{T, F\}^n$.

    For $\rho \in [0, 1]$, we say $y \sim N(\varepsilon)$ if for each coordinate $i \in [n]$, $y_i = True$ with probability ε and $False$ with probability $1 - \varepsilon$. We say that y is a random ε-noise vector.

    For $\alpha \in \{T, F\}^n$ and $\rho \in [0, 1]$, we say that β is a ρ-accurate copy of α if $\beta = \alpha \oplus y$ where $y \sim N((1 - \rho)/2)$. (Thus when $\rho = 1$, $\beta = \alpha$ and when $\rho = 0$, β is uniformly random.)

    We define the noise operator T_ρ (also called the Bonami-Beckner operator) which maps $f \in \mathbb{R}^{B^n}$ to the function $T_\rho[f]$ given by $T_\rho[f](\alpha) = \mathbb{E}[f(\beta)]$, where β is a ρ-accurate copy of f.

    Define $\mathbb{S}_\rho(f, g) = \langle f, T_\rho[g] \rangle$ and $\mathbb{S}_\rho(f) = \mathbb{S})_\rho(f, f)$.

    Also, for $i \in [n]$ define the ith derivative operator D_i that maps f to $D_i[f]$ given by $D_i[f](\alpha) = (f(\alpha^{(i=F)}) - f(\alpha^{(i=T)}))/2$ (here $\alpha^{(i=T)}$ is the input obtained from α be setting $\alpha_i = T$.).

    (a) Determine the fourier coefficients of $T_\rho[f]$ in terms of the fourier coefficients of f.

    (b) Show that $\mathbb{S}_\rho(f, g) = \mathbb{S}_\rho(g, f)$.

    (c) Determine the fourier coefficients of $D_i[f]$ in terms of the fourier coefficients of f.

    (d) Show that for any function f and $\rho \in [0, 1]$, $Inf_i(T_{\sqrt{\rho}}[f]) = \rho\mathbb{S}_\rho(D_i[f])$ and $Inf(T_{\sqrt{\rho}}[f]) = \rho\frac{\partial}{\partial\rho}\mathbb{S}_\rho(f)$.

8. Let f be a boolean function and let T be a decision tree for f. Let $p \in [0,1]$ and let μ be the distribution on inputs to f in which each bit is set to T with probability p. Let Δ be the expected cost of T (number of variables read by T) when inputs are distributed according to μ. Prove: $\sum_i \hat{f}^{(p)}(i) \leq \sqrt{\Delta}$. (Hints provided below)

Hints

Problem 2 . How many isolated vertices must G have?

Problem 4 Use reverse induction on the number of partial assignments in P.

Problem 5 Part 5c. Analyze $G_{n,p}$ with an appropriate set of bad events.

Problem 6a Use the map $T \longrightarrow -1$ and $F \longrightarrow 1$ to view f as a polynomial. Proceed by induction on n.

Problem 8 Suggestion: Look at hints one at a time and try to solve this using a minimum number of the hints.

1. First hint: express f as a sum over root-to-leaf paths P of some function g_P associated to path P.
2. Second hint: Write $\sum \hat{f}^{(p)}(i)$ using the definition as an expected value and use the expression of the previous hint to simplify.
3. Third hint: Apply the Cauchy-Schwartz inequality (possibly twice).