HOMEWORK 5

Due: Tuesday, April 1

---

1. Fourier expansion uniqueness. Using *only* Parseval's Theorem (and arithmetic) show that each $f : \{-1,1\}^n \to \mathbb{R}$ is *uniquely* expressible as a multilinear polynomial.

2. Plancherel. Prove "Plancherel's identity": if $f, g : \{-1,1\}^n \to \mathbb{R}$, then

$$\mathbf{E}_{\boldsymbol{x}}[f(\boldsymbol{x})g(\boldsymbol{x})] = \sum_{S\subseteq[n]} \hat{f}(S)\hat{g}(S).$$

Disprove

$$\mathbf{E}_{\boldsymbol{x}}[f(\boldsymbol{x})g(\boldsymbol{x})h(\boldsymbol{x})] = \sum_{S\subseteq[n]} \hat{f}(S)\hat{g}(S)\hat{h}(S).$$

3. Fourier support. Given $f : \{-1,1\}^n \to \mathbb{R}$, the "Fourier support" of f is $\mathcal{S} = \{S \subseteq [n] : \hat{f}(S) \neq 0\}$.

a. Identify all boolean-valued $f : \{-1,1\}^n \to \{-1,1\}$ with Fourier support $\mathcal{S}$ satisfying $|\mathcal{S}| = 1$.

b. Repeat question (a) but for $|\mathcal{S}| = 2$ and for $|\mathcal{S}| = 3$.

c. If the Fourier support $\mathcal{S}$ of $f : \{-1,1\}^n \to \mathbb{R}$ has $|S|$ odd for all $S \in \mathcal{S}$, we say "f is odd". State and prove a *simple* characterization of the odd functions f — in terms of their values, rather than their Fourier coefficients.

4. Quasirandomness. We say $f : \{-1,1\}^n \to \mathbb{R}$ is "(ϵ,δ)-quasirandom" if $\hat{f}(S)^2 \le \epsilon$ for all $0 < |S| \le 1/\delta$. Clearly this definition becomes stricter as ϵ and δ get smaller. Recall also that the "bias" of f is $\mathbf{E}_{\boldsymbol{x}}[f(\boldsymbol{x})]$.

a. Suppose f is (ϵ,δ)-quasirandom. Show that the bias of g is within an additive $\pm\sqrt{\epsilon}\cdot 2^c$ of the bias of f whenever g is a subfunction of f gotten by fixing at most $c \le 1/\delta$ input bits.

b. Conversely, suppose that $f : \{-1,1\}^n \to \mathbb{R}$ has the property that for any subfunction g gotten by fixing at most $c \le 1/\delta$ input bits, the bias changes by at most an additive $\pm\sqrt{\epsilon}$. Show that f is (ϵ,δ)-quasirandom.

c. Let $f : \{-1,1\}^n \to \{-1,1\}$ be the Majority function (assuming n is odd, for simplicity.) Show that f is (ϵ,δ)-quasirandom for ϵ and δ as small as you can. (This problem is fairly open-ended.)

5. Isoperimetry. For $f : \{-1,1\}^n \to \{-1,1\}$, let us define the *influence of coordinate* i to be:

$$\mathrm{Inf}_i(f) = \mathbf{Pr}_{\boldsymbol{x}}[f(\boldsymbol{x}) \neq f(x^{\oplus i})]$$

where $x^{\oplus i}$ is x with the ith bit flipped. The *total influence* of f is $\mathrm{Inf}(f) = \sum_i \mathrm{Inf}_i(f)$.

a. Show that $\mathrm{Inf}_i(f) = \sum_{S:i\in S} \widehat{f}(S)^2$. Hence show that $\mathrm{Inf}(f) = \sum_S |S|\widehat{f}(S)^2$.

b. Show that $\mathrm{Var}[f] := \mathbf{E}[f^2] - \mathbf{E}[f]^2$ is equal to $\sum_{S\neq\emptyset} \widehat{f}(S)^2$.

c. Use the above parts to show that for any set $A \subseteq \{-1,1\}^n$ with $|A| \le 2^n/2$, the number of edges in ∂A (i.e., edges having one endpoint in A and the other outside A) is at least $2|A|(1 - |A|/2^n)$.

d. Conclude that the edge-expansion of the cube is exactly 1.

6. An Algorithm for Sparsest Cut. Recall the *(generalized) sparsest cut* problem defined on a graph $G = (V, E)$ with edge capacities/costs c_e, and terminal pairs $T = \{(s_1, t_1), (s_2, t_2), \ldots, (s_k, t_k)\}$, with "demands" D_i for each (s_i, t_i) pair. In this problem, you need not separate all pairs; however, if deleting E' separates pairs $\{(s_j, t_j) \mid j \in J \subseteq [k]\}$, the objective function is $\phi(E') = \frac{c(E')}{\sum_{j \in J} d_j}$. The LP relaxation is $\min \sum_e c_e x_e$ over $\{x$ is a metric, $\sum_i D_i \cdot x_{(s_i,t_i)} = 1, x \geq 0\}$.

a. Given a solution x for the sparsest cut LP with LP value λ^*: for each $j = 0, 1, \ldots, (\log \sum_i D_i)$, define $T_j = \{i \mid x_{(s_i,t_i)} \in [2^{-j-1}, 2^{-j})$ to be the terminal pairs at distance $\approx 2^{-j}$ in the LP.

   Show that, for one of these values of j, (a scaled version of) the LP solution x is a feasible LP solution to the sparsest cut instance (G, T_j), with the LP value at most $\widehat{\lambda^*} = \lambda^* \times O(\log \sum_i D_i)$. (Be careful: do all the terminal pairs belong to some T_j defined above?)

b. Since the distances between all (s_i, t_i) pairs in this T_j are almost the same, use the low-diameter decompsition procedure given in class to separate these pairs. Show that you get a solution E' to sparsest cut with $\phi(E') \leq \widehat{\lambda^*} \times O(\log n)$. Conclude that you get an $O(\log n \log \sum_i D_i)$ approximation to sparsest cut.

7. Integrality Gaps for Multicut. A (d, c, α)-expander is a graph with n vertices and the property that each vertex has degree d, and for any set S with $|S| \leq cn$, the edge-expansion $\frac{|\partial S|}{|S|} \geq \alpha$. Say, a random 10-regular graph is an expander for $c = 0.01, \alpha = 1.01$ with high probability—of course, randomized constructions with much better parameters are known, as are explicit constructions.

a. Suppose we delete a set of edges in such an expander G so that each connected component has diameter at most $\frac{1}{2}\log_{10} n$. Show that we must have deleted $\Omega(n)$ edges.

b. Use this observation to construct an instance of multicut on these expanders which have an integrality gap of $\Omega(\log n)$ with respect to the LP used in class. (Hint: how can you force each connected component in the feasible solution to have small diameter?)

Now consider the cube: a stronger isoperimetric inequality than the one above shows that for any $A \subseteq \{-1, 1\}^d$, $|\partial A| \geq |A|(d - \log |A|)$. We will consider the cube with $d = \log_2 n$ (and hence having n vertices.).

d. Suppose we delete edges in the cube so that each connected component has diameter at most $\frac{1}{10}\log_2 n$. Show that the number of edges cut is at least $\Omega(n \log n)$.

e. Give a multicut instance on the cube with an $\Omega(\log n)$ integrality gap. In your instance, setting the length of each hypercube edge to $\epsilon \log n$ for some $\epsilon > 0$ should be a feasible LP solution.

f. Show that the (non-linear) constraint "the metric x_e embeds into ℓ_1" is a valid constraint for the multicut problem.

   Conclude from the previous part that even if we could augment the multicut LP with these constraints, the resulting integrality gap would be $\Omega(\log n)$.