PROBLEM SET 4

Due: Monday, Oct. 8, beginning of class

Homework policy: Please try to work on the homework by yourself; it isn't intended to be too difficult. Questions about the homework or other course material can be asked on Piazza.

---

1. Informally: a "one-way permutation" is a bijective function $f : \mathbb{F}_2^n \to \mathbb{F}_2^n$ which is easy to compute on all inputs but hard to invert on more than a negligible fraction of inputs; a "pseudorandom generator" is a function $g : \mathbb{F}_2^k \to \mathbb{F}_2^m$ for $m > k$ whose output on a random input "looks unpredictable" to any efficient algorithm. Goldreich and Levin proposed the following construction of the latter from the former: for $k = 2n$, $m = 2n+1$, define
$$g(r,s) = (r, f(s), r \cdot s),$$
where $r,s \in \mathbb{F}_2^n$. When g's input $(\boldsymbol{r},\boldsymbol{s})$ is uniformly random then so is the first $2n$ bits of its output (using the fact that f is a bijection). The key to the analysis is showing that the final bit, $\boldsymbol{r}\cdot\boldsymbol{s}$, is highly unpredictable to efficient algorithms even *given* the first $2n$ bits $(\boldsymbol{r}, f(\boldsymbol{s}))$. This is proved by contradiction.

   (a) Suppose that an adversary has a deterministic, efficient algorithm A good at predicting the bit $\boldsymbol{r}\cdot\boldsymbol{s}$:
$$\mathop{\mathbf{Pr}}_{\boldsymbol{r},\boldsymbol{s}\sim\mathbb{F}_2^n}[A(\boldsymbol{r}, f(\boldsymbol{s})) = \boldsymbol{r}\cdot\boldsymbol{s}] \geq \tfrac{1}{2} + \gamma.$$
   Show there exists $B \subseteq \mathbb{F}_2^n$ with $|B|/2^n \geq \frac{1}{2}\gamma$ such that for all $s \in B$,
$$\mathop{\mathbf{Pr}}_{\boldsymbol{r}\sim\mathbb{F}_2^n}[A(\boldsymbol{r}, f(s)) = \boldsymbol{r}\cdot s] \geq \tfrac{1}{2} + \tfrac{1}{2}\gamma.$$

   (b) Switching to ± 1 notation in the output, deduce $\widehat{A_{[n]|f(s)}}(s) \geq \gamma$ for all $s \in B$.

   (c) Show that the adversary can efficiently compute s given $f(s)$ (with high probability) for any $s \in B$. If γ is nonnegligible this contradicts the assumption that f is "one-way". (Hint: use the Goldreich–Levin algorithm.)

   (d) Deduce the same conclusion even if A is a randomized algorithm.

2. Given $f : \{-1,1\}^n \to \{-1,1\}$ and integer $k \geq 2$ let $A_k = \frac{1}{k}(\mathbf{W}^{\geq 1}[f] + \mathbf{W}^{\geq 2}[f] + \cdots + \mathbf{W}^{\geq k}[f])$, the "average of the first k tail weights". (Recall $\mathbf{W}^{\geq \ell}[f] = \sum_{|S|\geq\ell} \widehat{f}(S)^2$.) Show that $\mathbf{NS}_{1/k}[f]$ is the same as A_k up to universal constants. E.g., you might show $\frac{1-e^{-2}}{2} A_k \leq \mathbf{NS}_{1/k}[f] \leq A_k$.

3. Let $f : \{-1,1\}^n \to \mathbb{R}$ and let $\epsilon > 0$. Show that f is ϵ-concentrated on a collection $\mathscr{F} \subseteq 2^{[n]}$ with $|\mathscr{F}| \leq \hat{\|} f \hat{\|}_1^2/\epsilon$. (Recall the notation from Problem 1 on Homework 3.)

4. For this problem, recall Problem 3 from Homework 3.

   (a) Let $H \leq \mathbb{F}_2^n$ be a subspace and let $z \in \mathbb{F}_2^n$. Let $\varphi_{H+z} : \mathbb{F}_2^n \to \mathbb{R}$ be the probability density function associated to the uniform probability distribution on the affine subspace $H+z$. Write the Fourier expansion of φ_{H+z}.

(b) For $f : \mathbb{F}_2^n \to \mathbb{R}$ and $z \in \mathbb{F}_2^n$, define the function $f^{+z} : \mathbb{F}_2^n \to \mathbb{R}$ by $f^{+z}(x) = f(x+z)$. Show that $f^{+z} = \varphi_{\{z\}} * f$. (In writing $\varphi_{\{z\}}$ we are treating $\{z\}$ as a 0-dimensional affine subspace and using the notation of the previous problem.) Show also that $\widehat{f^{+z}}(\gamma) = (-1)^{\gamma \cdot z} \widehat{f}(\gamma)$.

(c) Prove the "Poisson Summation Formula",

$$\mathop{\mathbf{E}}_{\boldsymbol{h} \sim H}[f(\boldsymbol{h}+z)] = \sum_{\gamma \in H^\perp} \chi_\gamma(z) \widehat{f}(\gamma).$$

(Hint: use Plancherel on $\langle \varphi_H, f^{+z} \rangle$.)

5. Give a direct (Fourier-free) simple proof that if $f : \{-1,1\}^n \to \mathbb{R}$ and $(\boldsymbol{J} \mid \boldsymbol{z})$ is a δ-random restriction then $\mathbf{E}[\mathbf{Inf}_i[f_{\boldsymbol{J}|\boldsymbol{z}}]] = \delta \mathbf{Inf}_i[f]$ for any $i \in [n]$.

6. In this exercise you will prove the "Baby Switching Lemma": If $\phi = T_1 \vee T_2 \vee \cdots \vee T_s$ is a DNF of width $w \geq 1$ over variables $x_1, \dots, x_n$ and $(\boldsymbol{J} \mid \boldsymbol{z})$ is a δ-random restriction ($0 < \delta < 1/3$), then

$$\mathbf{Pr}[f_{\boldsymbol{J}|\boldsymbol{z}} \text{ is not a constant function}] \leq 3\delta w.$$

(a) Suppose $R = (J \mid z)$ is a "bad" restriction, meaning that $\phi_{J|z}$ is not a constant function. Let i be minimal such that $(T_i)_{J|z}$ is neither constantly True or False, and let j be minimal such that x_j or $\overline{x}_j$ appears in this restricted term. Show there is a unique restriction $R' = (J \setminus \{j\} \mid z')$ extending R which doesn't falsify T_i.

(b) Suppose we enumerate all bad restrictions R, and for each we write down the associated R' as in part (6a). Show that no restriction is written more than w times.

(c) If $(\boldsymbol{J} \mid \boldsymbol{z})$ is a δ-random restriction and R and R' are as in part (6a), show that $\mathbf{Pr}[(\boldsymbol{J} \mid \boldsymbol{z}) = R] = \frac{2\delta}{1-\delta} \mathbf{Pr}[(\boldsymbol{J} \mid \boldsymbol{z}) = R']$.

(d) Complete the proof by showing $\mathbf{Pr}[(\boldsymbol{J} \mid \boldsymbol{z}) \text{ is bad}] \leq 3\delta w$.