

Problem Set 6

Harvard SEAS - Spring 2015

Due: Fri. May 1, 2015

Your problem set solutions must be typed (in e.g. L^AT_EX) and submitted electronically to `cs225-hw@seas.harvard.edu`. You are allowed 12 late days for the semester, of which at most 5 can be used on any individual problem set. (1 late day = 24 hours exactly). Please name your file `PS6-lastname.*`.

The problem sets may require a lot of thought, so be sure to start them early. You are encouraged to discuss the course material and the homework problems with each other in small groups (2-3 people). Identify your collaborators on your submission. Discussion of homework problems may include brainstorming and verbally walking through possible solutions, but should not include one person telling the others how to solve the problem. In addition, each person must write up their solutions independently, and these write-ups should not be checked against each other or passed around.

Strive for clarity and conciseness in your solutions, emphasizing the main ideas over low-level details. Do not despair if you cannot solve all the problems! Difficult problems are included to stimulate your thinking and for your enjoyment, not to overwork you. *ed problems are extra credit.

Problem 7.1. PRGs imply hard functions Suppose that for every m , there exists a mildly explicit $(m, 1/m)$ pseudorandom generator $G_m : \{0, 1\}^{d(m)} \rightarrow \{0, 1\}^m$. Show that \mathbf{E} has a function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ with nonuniform worst-case hardness $t(\ell) = \Omega(d^{-1}(\ell - 1))$. In particular, if $d(m) = O(\log m)$, then $t(\ell) = 2^{\Omega(\ell)}$ (Hint: look at a prefix of G 's output.)

Problem 7.4. Deterministic Approximate Counting Using the PRG for constant-depth circuits of Theorem 7.29, give deterministic quasipolynomial-time algorithms for the problems below. (The running time of your algorithms should be $2^{\text{poly}(\log n, \log(1/\varepsilon))}$, where n is the size of the circuit/formula given and ε is the accuracy parameter mentioned.)

1. Given a constant-depth circuit C and $\varepsilon > 0$, approximate the fraction of inputs x such that $C(x) = 1$ to within an additive error of ε .
2. Given a DNF formula φ and $\varepsilon > 0$, approximate the number of assignments x such that $\varphi(x) = 1$ to within a *multiplicative* fraction of $(1 + \varepsilon)$. You may restrict your attention to φ in which all clauses contain the same number of literals. (Hint: Study the randomized DNF counting algorithm of Theorem 2.34.)

Note that these are *not* decision problems, whereas classes such as \mathbf{BPP} and \mathbf{BPAC}_0 are classes of decision problems. One of the points of this problem is to show how derandomization can be used for other types of problems.

Problem 7.6. Private Information Retrieval The goal of *private information retrieval* is for a user to be able to retrieve an entry of a remote database in such a way that the server

holding the database *learns nothing* about which database entry was requested. A trivial solution is for the server to send the user the entire database, in which case the user does not need to reveal anything about the entry desired. We are interested in solutions that involve much less communication. One way to achieve this is through replication.¹ Formally, in a *q-server private information-retrieval (PIR) scheme*, an arbitrary database $D \in \{0, 1\}^n$ is duplicated at q non-communicating servers. On input an index $i \in [n]$, the *user algorithm* U tosses some coins r and outputs queries $(x_1, \dots, x_q) = U(i, r)$, and sends x_j to the j 'th server. The j 'th server algorithm S_j returns an answer $y_j = S_j(x_j, D)$. The user then computes its output $U(i, r, x_1, \dots, x_q)$, which should equal D_i , the i 'th bit of the database. For privacy, we require that the distribution of each query x_j (over the choice of the random coin tosses r) is the same regardless of the index i being queried.

It turns out that q -query locally decodable codes and q -server PIR are essentially equivalent. This equivalence is proven using the notion of *smooth codes*. A code $\text{Enc} : \{0, 1\}^n \rightarrow \Sigma^{\hat{n}}$ is a *q-query smooth code* if there is a probabilistic oracle algorithm Dec such that for every message x and every $i \in [n]$, we have $\Pr[\text{Dec}^{\text{Enc}(x)}(i) = x_i] = 1$ and Dec makes q nonadaptive queries to its oracle, each of which is uniformly distributed in $[\hat{n}]$. Note that the oracle in this definition is a valid codeword, with no corruptions. Below you will show that smooth codes imply locally decodable codes and PIR schemes; converses are also known (after making some slight relaxations to the definitions).

1. Show that the decoder for a q -query smooth code is also a local $(1/3q)$ -decoder for Enc .
2. Show that every q -query smooth code $\text{Enc} : \{0, 1\}^n \rightarrow \Sigma^{\hat{n}}$ gives rise to a q -server PIR scheme in which the user and servers communicate at most $q \cdot (\log \hat{n} + \log |\Sigma|)$ bits for each database entry requested.
3. Using the Reed-Muller code, show that there is a $\text{polylog}(n)$ -server PIR scheme with communication complexity $\text{polylog}(n)$ for n -bit databases. That is, the user and servers communicate at most $\text{polylog}(n)$ bits for each database entry requested. (For constant q , the Reed-Muller code with an optimal systematic encoding as in Problem 5.4 yields a q -server PIR with communication complexity $O(n^{1/(q-1)})$.)

Problem 7.13. Hardcore Predicates A *hardcore predicate* for a one-way function $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ is a $\text{poly}(\ell)$ -time computable function $b : \{0, 1\}^\ell \rightarrow \{0, 1\}$ such that for every constant c , every nonuniform algorithm A running in time ℓ^c , we have:

$$\Pr[A(f(U_\ell)) = b(U_\ell)] \leq \frac{1}{2} + \frac{1}{\ell^c},$$

for all sufficiently large ℓ . Thus, while the one-wayness of f only guarantees that it is hard to compute *all* the bits of f 's input from its output, b specifies a particular bit of information about the input that is very hard to compute (one can't do noticeably better than random guessing).

1. Let $\text{Enc} : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\hat{L}}$ be a code such that given $x \in \{0, 1\}^\ell$ and $y \in [\hat{L}]$, $\text{Enc}(x)_y$ can be computed in time $\text{poly}(\ell)$. Suppose that for every constant c and all sufficiently large ℓ , Enc has a $(1/2 - 1/\ell^c)$ local list-decoding algorithm $(\text{Dec}_1, \text{Dec}_2)$ in which both Dec_1 and Dec_2 run in time $\text{poly}(\ell)$. Prove that if $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ is a one-way function, then $b(x, y) = \text{Enc}(x)_y$ is a hardcore predicate for the one-way function $f'(x, y) = (f(x), y)$.

¹Another way is through computational security, where we only require that it be *computationally infeasible* for the database to learn something about the entry requested.

2. Show that if $b : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ is a hardcore predicate for a one-way *permutation* $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$, then for every $m = \text{poly}(\ell)$, the following function $G : \{0, 1\}^\ell \rightarrow \{0, 1\}^m$ is a cryptographic pseudorandom generator:

$$G(x) = (b(x), b(f(x)), b(f(f(x))), \dots, b(f^{(m-1)}(x))).$$

(Hint: show that G is “previous-bit unpredictable”.)