Homework 1

1 Problem 1 (15 points)

Let $x \in \{0,1\}^\lambda$, and let $H : \{0,1\}^\lambda \to \{0,1\}$ be a function such that $H(r) = \langle x, r \rangle$ for at least a fraction p its inputs r. Here, $\langle x, r \rangle$ means the inner product mod 2 of x and r: $\langle x, r \rangle = \sum_{i=1}^{\lambda} x_i r_i \bmod 2$.

In class, we showed that if $p \geq \frac{3}{4} + \epsilon$ for a non-negligible ϵ, then it is possible to determine x efficiently, given only polynomially-many queries to H. Here, you will show that this is essentially tight.

(a) Construct two inputs $x_0 \neq x_1$ and a function H such that $H(r) = \langle x_0, r \rangle$ for at least 3/4 of its inputs, and at the same time $H(r) = \langle x_1, r \rangle$ for at least 3/4 of its inputs. Note that the two sets of inputs may be different.

This is why, when moving to the regime where $p = \frac{1}{2} + \epsilon$, we could no longer give an algorithm that outputted a single x. Instead, we had to output multiple x values, one of which was the right answer.

(b) Generalize the above construction to more inputs. For any integer n, construct n distinct inputs $x_0, \ldots, x_{n-1}$ and a function H such that $H(r) = \langle x_i, r \rangle$ for at least p fraction of inputs simultaneously for all i, where $p = \frac{1}{2} + \frac{1}{2n}$. Here, you may assume n is a power of 2.

2 Problem 2 (20 points)

In class, we built a PRF with where the range was equal to the key length, and the domain was arbitrary. Here, we will show how to vary the key length, domain, and range.

(a) Let $\texttt{PRF} : \{0,1\}^\lambda \times \{0,1\}^n \to \{0,1\}^m$ be a PRF. Give a simple construction of a PRF $\texttt{PRF}' : \{0,1\}^\lambda \times \{0,1\}^{n'} \to \{0,1\}^{km}$, for a given value k (which may be polynomial in λ). Here, n' should only be slightly smaller than n. Prove that $\texttt{PRF}'$ is secure, assuming only the security of $\texttt{PRF}$.

(b) Let $\texttt{PRF} : \{0,1\}^\lambda \times \{0,1\}^n \to \{0,1\}^\lambda$ be a PRF where the range is the same as the key length. Let $\texttt{PRF}' : \{0,1\}^\lambda \times \{0,1\}^{kn} \to \{0,1\}^\lambda$ for a given integer k be defined as follows: on input $x \in \{0,1\}^{kn}$, write $x = (x_1, \ldots, x_k)$ for $x_i \in \{0,1\}^n$. Then run $\texttt{PRF}$ on k and x_1 to obtain a new PRF key $k_{x_1} = \texttt{PRF}(k, x_1)$. Then run $\texttt{PRF}$ again, this time with key k_{x_1} and input x_2 to derive a different PRF key $k_{x_1,x_2} = \texttt{PRF}(k_{x_1}, x_2)$. Repeat this process to derive PRF keys $k_{x_1,x_2,x_3}, k_{x_1,x_2,x_3,x_4}$, etc, until you have computed $k_{x_1,x_2,\ldots,x_k}$. Define $k_{x_1,x_2,\ldots,x_k}$ as the output of $\texttt{PRF}'$ on input $(x_1, \ldots, x_k)$.

Prove that $\texttt{PRF}'$ is a secure PRF. If it helps, you may assume that the queries the adversary makes are fixed and known in advance (like we did when we constructed PRFs from PRGs).

(c) Explain how the PRF construction from any PRG we saw in class is a special case of Part (b).

3 Problem 3 (25 points)

In class, we defined security for a message authentication code as follows. Let $(\texttt{MAC}, \texttt{Ver})$ be a MAC. Define EUF-CMA-Exp(A, λ) as the following experiment on A:

- The challenger Ch chooses a random key $k \in \{0,1\}^\lambda$
- A is allowed to make many queries on arbitrary messages m. In response, the Ch runs $\sigma \leftarrow \texttt{MAC}(k, m)$, and gives σ to A. These queries can be made adaptively in sequence, so for example the third query message may depend on the MACs obtained from the first two queries.
- Finally, A outputs a forgery candidate (m', σ'). Ch checks that (m', σ') was not the message/MAC pair in one of A's queries, and that $\texttt{Ver}(k, m', \sigma')$ accepts. If both checks pass, Ch outputs 1; otherwise it outputs 0.

We define security by saying, for all PPT adversaries A, the probability that EUF-CMA-Exp(A, λ) outputs 1 is negligible.

Here, we consider a more general variant. EUF-CMA-Exp$'(A, \lambda)$ is identical to the above, except that we allow A to additionally make verification queries, interleaved arbitrarily with the choosen message queries. Here, A makes a query on (m, σ), and Ch returns the result of $\texttt{Ver}(k, m, \sigma)$. Otherwise, the two experiments are the same. Security is defined analagously.

Show that the two definitions of security are equivalent. Namely, given an adversary A that breaks EUF-CMA security, construct an adversary that breaks EUF-CMA$'$ security, and vice versa.

4 Problem 4 (40 points)

Here, you will extend the Goldreich-Levin theorem to multiple hardcore bits.

Let $F : \{0,1\}^\lambda \to \{0,1\}^{n(\lambda)}$ be a one-way function. Let $F' : \{0,1\}^{k\lambda+\lambda} \to \{0,1\}^{k\lambda+n(\lambda)}$ be the function

$$F'(r_1, \ldots, r_k, x) = (r_1, \ldots, r_k, F(x))$$

Assume k is *logarithmic* in λ. Consider the functions $h_i(r_1, \ldots, r_k, x) = \langle r_i, x \rangle$. Show that $h_1, \ldots, h_k$ are all *simultaneously* hardcore bits for F'. This means that for any PPT adversary A, there exists a negligible ϵ such that

$$\begin{aligned}\Big|&\Pr[1 \leftarrow A(F'(x'), h_1(x'), \ldots, h_k(x')) : x' \leftarrow \{0,1\}^{k\lambda+\lambda}] \\ &\quad - \Pr[1 \leftarrow A(F'(x'), b_1, \ldots, b_k) : x' \leftarrow \{0,1\}^{k\lambda+\lambda},\ b_1, \ldots, b_k \leftarrow \{0,1\}]\Big| < \epsilon(\lambda)\end{aligned}$$

To prove this, you can use the basic Goldreich-Levin theorem as a black box (but perhaps for a slightly modified one-way function); you do not need to reprove GL from scratch in this more general setting.

5 Problem 5 (0 points)

Please let us know roughly how long you spent on this homework assignment (for calibrating future homework assignments).