

Problem Set 1

November 26, 2008

Due: Friday Dec 5 in class

1. Concrete parameters for hardness amplification:

- (a) You wish to obtain a function f' that is $(0.01, 2^{10})$ -one-way (namely, any adversary that runs at most 2^{10} steps can invert f' on at most 0.01 of the points. For this purpose you have a supplier that is willing to sell you, for any e , a function that is $(e, 2^{30})$ -one way at a price of $1/e$ shekels. How much will you have to pay, using the hardness amplification construction in class?
Can you save money by switching to a vendor that is willing to sell, for any e, t , a function that is (e, t) -one way at a price of $2^{-30} \cdot t/e$ shekels? How would this affect your domain size, assuming that both vendors have functions where the domain size is $c \cdot t/e$ for some constant c ? How will the answer change if the domain size is $c \cdot t$?
- (b) You have in hand a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$ that is guaranteed to be $\epsilon(t)$ -one-way as defined in class. You wish to have a function $f' : \{0, 1\}^m \rightarrow \{0, 1\}^*$ that is $\epsilon^k(t)$ -one-way. How large should m be, as a function of n, ϵ, k , using the parameters from the proof in class?

(Note that in this question there are no asymptotics, all the numbers are absolute.)

2. **Hardness of the Discrete Log function:** Let $\mathbf{P} = \{(p_n, g_n)\}_{n \in \mathbf{N}}$ be such that $2^n < p_n < 2^{n+1}$ is a prime and g_n is a generator of the group Z_p^* . Let the the Discrete Log function be $DL_{\mathbf{P}}(x) = g_{|x|}^x \pmod{p_{|x|}}$, and assume that $DL_{\mathbf{P}}(x)$ is weakly one way. Show that $DL_{\mathbf{P}}(x)$ is one way. (For the purpose of this question, assume that $p_{|x|}, g_{|x|}$ can be computed efficiently given x . Here $|x|$ denotes the number of bits in the binary description of x .)

Bonus question: Can you show that $DL_{\mathbf{P}}(x)$ is one way under an even weaker assumption on $DL_{\mathbf{P}}(x)$?

3. **An alternative definition of statistical distance:** Let $\mathcal{D}^1 = \{D_n^1\}_{n \in \mathbf{N}}$ and $\mathcal{D}^2 = \{D_n^2\}_{n \in \mathbf{N}}$ be two distribution ensembles, and let Δ_n be the (common) support of D_n^1 and D_n^2 . Show that $\mathcal{D}_1 \approx_s \mathcal{D}_2$ (as defined in class) iff

$$\frac{1}{2} \sum_{\delta \in \Delta_n} |\text{Prob}_{d \leftarrow D_n^1}[d = \delta] - \text{Prob}_{d \leftarrow D_n^2}[d = \delta]| \quad (1)$$

is a negligible function of n .

(Note: The formulation in (1) is the common definition of the statistical distance between distributions. The formulation in class was used as a way to motivate the notion of computational indistinguishability.)

4. **Preservation of computational indistinguishability under efficient transformations:** For a distribution D and a function f , let $f(D)$ denote the distribution obtained by sampling a value d from D and applying $f(d)$. Let $\mathcal{D}^1 = \{D_n^1\}_{n \in \mathbf{N}}$ and $\mathcal{D}^2 = \{D_n^2\}_{n \in \mathbf{N}}$ be two computationally indistinguishable distribution ensembles, and let f be a deterministic function computable in polynomial time. Show that the ensembles $\{f(D_n^1)\}_{n \in \mathbf{N}}$ and $\{f(D_n^2)\}_{n \in \mathbf{N}}$ are computationally indistinguishable. Does this result hold if f can be probabilistic? Computable in time that's exponential in n ?