CS 120/CSCI E-177: Introduction to Cryptography

Problem Set 2

Assigned: Oct. 5, 2006 | Due: Oct. 11, 2006 (1:10 PM)

Justify all of your answers. See the syllabus for collaboration and lateness policies. You can submit by email to `ciocan@eecs` (please include source files) or by hardcopy to Carol Harlow in MD 343.

Problem 1. (Factorization is "NP-easy")

1. Let $L = \{(x, y) \in \mathbb{N} \times \mathbb{N} : x \text{ has a factor between 2 and } y\}$. Show that the language L is in **NP**.
2. Show that if L is in **P**, then there is a polynomial-time algorithm for integer factorization. Thus, if $\mathbf{P} = \mathbf{NP}$, then factorization is easy.

Problem 2. (Reducing the error of randomized algorithms) Suppose we have randomized algorithm for computing a function f which gives an incorrect answer with probability $\leq 1/3$, and we want to reduce its error by repeating it several times and taking a majority vote. Use the Chernoff Bound to estimate how many repetitions suffice to reduce the error probability to $1/1000$. And to 2^{-k}?

Problem 3. (Statistical Security) Recall that (G, E, D) has *statistically ε-indistinguishable encryptions* if for every two $m_1, m_2 \in \mathcal{P}$ and every $T \subseteq \mathcal{C}$,

$$|\Pr[E_K(m_1) \in T] - \Pr[E_K(m_2) \in T]| \leq \varepsilon,$$

where the probabilities are taken over $K \stackrel{\text{R}}{\leftarrow} G$ and the coin tosses of E.

1. Show that statistical 0-indistinguishability is equivalent to perfect indistinguishability.

For the remaining parts, suppose (G, E, D) has statistically ε-indistinguishable encryptions for message space $\mathcal{P}$. Below you will prove that the number of keys must be at least $(1 - \varepsilon) \cdot |\mathcal{P}|$, so statistical security doesn't help much to overcome the limitations of perfect secrecy.

2. Call a ciphertext c *decryptable* to $m \in \mathcal{P}$ if there is a key k such that $D_k(c) = m$. Prove that for every two messages $m, m' \in \mathcal{P}$,

   $$\Pr[E_K(m) \text{ is decryptable to } m'] \geq 1 - \varepsilon,$$

   where the probability is taken over $K \stackrel{\text{R}}{\leftarrow} G$ and the coin tosses of E.

3. Show that for every message $m \in \mathcal{P}$,

$$\mathrm{E}\left[\#\{m' : E_K(m) \text{ is decryptable to } m'\right] \geq (1-\varepsilon) \cdot |\mathcal{P}|,$$

where again the probability is taken over K and the coin tosses of E. (Hint: for each m', define a random variable $X_{m'}$ that equals 1 if $E_K(m)$ is decryptable to m', and equals 0 otherwise.)

4. Conclude that the number of keys must be at least $(1-\varepsilon) \cdot |\mathcal{P}|$.

5. Explain where this proof fails for computational security.