CS276 Midterm 1

March 8, 2004

Solutions due on paper at the start of class on **Wednesday, March 17.**
You must work on your own. No collaboration or discussion with others.
This is open-book, open-notes. You're allowed to use all the resources on the CS276 web page, including the scribe notes and the lecture notes linked therein. However, please don't use any other sources.

1 A hybrid encryption scheme

Let (E, D) be a public-key encryption scheme that can be used for encrypting n-bit messages. Let $G : \{0,1\}^n \to \{0,1\}^\ell$ be a pseudorandom generator. Define a public-key encryption scheme $(\mathcal{E}, \mathcal{D})$ intended to be used for encrypting ℓ-bit messages as follows. First, $\mathcal{E}_{\text{pk}}(m) = \mathcal{E}[E_{\text{pk}}, G](m)$, where

$\mathcal{E}[E_{\text{pk}}, G](m)$:
1. Pick $k \leftarrow \{0,1\}^n$ randomly.
2. Return $(E_{\text{pk}}(k), G(k) \oplus m)$.

Decryption is defined straightforwardly: $\mathcal{D}_{\text{tk}}(c) = \mathcal{D}[D_{\text{tk}}, G](c)$, where

$\mathcal{D}[D_{\text{tk}}, G](c)$:
1. Parse c as (y, z).
2. Return $G(D_{\text{tk}}(y)) \oplus z$.

Prove: If (E, D) is (t, q, ϵ_E)-RR-secure and if G is a (t, ϵ_G)-secure PRG, then $(\mathcal{E}, \mathcal{D})$ is $(t - cq, q, \epsilon_E + q\epsilon_G)$-RR-secure for some constant $c > 0$.

Reminder: RR-security refers to security of a public-key cryptosystem against chosen-plaintext attacks, in the real-or-random sense. In other words, (E, D) is (t, q, ϵ)-RR-secure if $\text{Adv}^{\text{rr}}(A) \le \epsilon$ for all A running in time $\le t$ and making $\le q$ queries, where

$$\text{Adv}^{\text{rr}}(A) \quad = \quad \Pr_{(\text{pk},\text{tk}) \leftarrow \text{Gen}()}[A^{E_{\text{pk}}}(\text{pk}) = 1] \quad - \quad \Pr_{(\text{pk},\text{tk}) \leftarrow \text{Gen}()}[A^{E_{\text{pk}} \circ \$}(\text{pk}) = 1],$$

where Gen denotes the key generation algorithm associated with (E, D) and $\$(\cdot)$ is a stateless[1] algorithm that returns a random string of the same length as its input

[1] In other words, calling $\$(\cdot)$ twice on the same input gives two independent random strings.

2 Implications and separations

I'm going to specify three possible definitions for security of symmetric-key encryption against chosen-plaintext attack. Let (E, D) be a symmetric-key encryption scheme with keyspace $\mathcal{K}$. Define functions $\mathbf{\$}, \mathbf{0} : \{0,1\}^* \to \{0,1\}^*$ so that $\mathbf{0}(x)$ represents the all-zeros string of the same length as x and $\mathbf{\$}(x)$ returns a random string of the same length as x. Define

$$\begin{aligned}
\mathrm{Adv}^{\mathrm{rr}}(A) &= \Pr_{k \leftarrow \mathcal{K}}[A^{E_k} = 1] - \Pr_{k \leftarrow \mathcal{K}}[A^{E_k \circ \mathbf{\$}} = 1] \\
\mathrm{Adv}^{\mathrm{rz}}(A) &= \Pr_{k \leftarrow \mathcal{K}}[A^{E_k} = 1] - \Pr_{k \leftarrow \mathcal{K}}[A^{E_k \circ \mathbf{0}} = 1] \\
\mathrm{Adv}^{\mathrm{rnd}}(A) &= \Pr_{k \leftarrow \mathcal{K}}[A^{E_k} = 1] - \Pr_{k \leftarrow \mathcal{K}}[A^{\mathbf{\$} \circ E_k} = 1].
\end{aligned}$$

Say that (E, D) is (t, q, ϵ)-XXX-secure if $\mathrm{Adv}^{\mathrm{xxx}}(A) \le \epsilon$ for all algorithms running in time $\le t$ and making $\le q$ queries. You've seen RR-security, but RZ and RND are new. Mnemonic: RR = "real-or-random", RZ = "real-or-zero", RND = "randomness of ciphertexts".

Your goal is to compare these definitions by proving some implications and/or separations. When I ask you "Does XXX $\Rightarrow$ YYY?", you should pick one of the following two claims (state which you're picking), fill in the blanks in the claim, and then prove it:

1. **Yes, it does.** (an *implication*)
   Claim: For all symmetric-key encryption schemes (E, D), if (E, D) is (t, q, ϵ)-XXX-secure, then (E, D) is also $(\_\_, \_\_, \_\_)$-YYY-secure.
   (Obviously, the blanks should be polynomially related to t, q, ϵ—don't try to get cute by telling me that (E, D) is guaranteed to be $(0, 0, 1)$-YYY-secure.)

2. **No, it doesn't.** (a *separation*)
   Claim: If there exists any[2] XXX-secure symmetric-key encryption scheme (E, D), then there exists a XXX-secure symmetric-key encryption scheme (E', D') that is not YYY-secure.
   (You may use asymptotic notions of security in this case, if you like.)

Obviously, these two choices are mutually contradictory. Pick only one.

For instance, if I asked "Does RR $\Rightarrow$ LR?", the correct answer would be "Yes, it does", since we saw in class that real-or-random security is equivalent to left-or-right security.

With that background, answer the following three questions:

(a) Does RR $\Rightarrow$ RZ?

(b) Does RZ $\Rightarrow$ RR?

(c) Does RR $\Rightarrow$ RND?

[2] Why is this clause needed? It's just a technicality, but it's necessary. For all we know, it could be that $P = NP$ and no encryption scheme is secure, and in this case the implication would be vacuously true. Therefore, without the assumption that there is *some* XXX-secure cryptosystem, proving a separation would be at least as hard as proving $P \neq NP$.