ICS 180: Introduction to Cryptography April 15, 2004

Lecture 4: One-Way Encryption vs. Indistinguishability

Lecturer: Stanislaw Jarecki

1 LECTURE SUMMARY

Last time we saw an example of an encryption scheme, the "textbook RSA" scheme, which can be one-way secure (that's exactly the belief expressed in the "RSA assumption") but is not secure in the sense of indistinguishability. Now we'll see that *any* one-way encryption might have some bad characteristics that make it not indistinguishably secure. With these arguments we'll try to convince you that the one-way security requirement on encryption is in fact not enough in practice.

2 One-Way Secure Encryption Can Leak Some Messages

We'll first show that an encryption scheme can be one-way secure and yet it can totally leak some messages. In fact, if an encryption scheme is one-way secure on some reasonable message space, for example $\mathcal{M}_\tau = \{0,1\}^\tau$ where τ is the security parameter, then it can very well be that there is a polynomially-sized subset $\mathcal{M}'_\tau \subset \{0,1\}^\tau$ of messages (i.e. $|\mathcal{M}'_\tau| \leq p(\tau)$ for some polynomial $p(\cdot)$)[1], s.t. when the encryption scheme is applied to any message $m \in \mathcal{M}'_\tau$, the adversary can immediately recover m from the ciphertext.

You might be tempted to think that since the size of this bad-message space $\mathcal{M}'$ is negligible compared to $\mathcal{M}_\tau$, because $\mathcal{M'}_\tau/\mathcal{M}_\tau = p(\tau)/2^\tau < \mathsf{negl}(\tau)$, maybe it follows that one is unlikely to encounter any m in this subset $\mathcal{M'}_\tau$? But that's not the right argument, because this encryption scheme can be bad for *any* set $\mathcal{M'}_\tau \subset \{0,1\}^\tau$, including the set of messages which are in fact the most likely ones that will get encrypted in a given application. For example, $\mathcal{M}'_\tau$ can contain "yes", "no", "nothing new", etc, and these might be what someone often wants to send.

Thus, an encryption scheme which reveals the plaintext of all messages in such subset would obviously not be a good one to use. And yet, we'll see that an encryption scheme can very well be one-way secure and still perform disastrously on some such message subset. What it means is that one-wayness of encryption is not a good enough reason to use this encryption scheme for general applications.

Fact 1 *If one-way encryption schemes exist at all then there are also encryption schemes which are one-way on some message space but which completely leak plaintexts of all messages belonging to some polynomially-sized subset of this message space (and hence in particular such one-way secure encryption schemes are not indistinguishably secure).*

Proof: Since we assume that one-way encryption schemes exist, take any encryption scheme $\Sigma = (KGen, Enc, Dec)$ which is one-way secure on message space $\{0,1\}^\tau$. [In fact you generalize this

[1] Note that the size of the message space $\mathcal{M}_\tau$ is $|\mathcal{M}_\tau| = 2^\tau$.

argument to *any* message spaces, but let's assume for simplicity that Σ is one-way secure on the message space of τ-bit long strings.]

Take any set $\mathcal{M}'$ s.t. for every τ, set $\mathcal{M}'$ contains at most $p(\tau)$ of τ-bit long messages. In other words, if we define $\mathcal{M}'_\tau$ as $\mathcal{M}' \cap \{0,1\}^\tau$, then there is a polynomial $p(\cdot)$ s.t. $|\mathcal{M}'_\tau| \leq p(\tau)$.

Using Σ, we'll construct an encryption scheme $\Sigma' = (KGen, Enc', Dec')$ s.t.:

1. If Σ is one-way secure on message space $\{0,1\}^\tau$ then Σ' must be one-way secure on message space $\{0,1\}^\tau$ as well.

2. Σ' will totally leak plaintexts of messages in set $\mathcal{M}'$ (and hence Σ' is not indistinguishably secure).

Why will that proof our claim? Because this will show that there exist one-way encryptions which are one-way secure for some message space, but totally leak messages in some polynomially-sized subset of that message space.

Constructing Σ' from Σ is actually very easy! Let $Enc'(k,m)$ output $c = Enc(k,m)$ if $m \notin \mathcal{M}_\tau$ and output simply $c = m$ in case $m \in \mathcal{M}_\tau$. The new decryption procedure $Dec'(k,c)$ will then invert this process. Just to make it super simple, we can ask Enc' to actually tag the ciphertext with tag 0 if it used Enc to encrypt m and with tag 1 if it just copied the plaintext m. Now decryption Dec' has an evevn easier job.

First, it's clear that Σ' is not indistinguishably secure: Just take any two $m_0, m_1 \in \mathcal{M}'_\tau$.

Second, we'll show that Σ' is one-way on $\{0,1\}^\tau$ if Σ is. The easiest way to show this is by contradiction: Assume that Σ' is *not* one-way secure, and show that Σ must be *not* one-way secure as a consequence. If Σ' were not one-way secure then there exists an efficient attack A which on input a random $c = Enc'(k,m)$, for random $k \leftarrow KGen(1^\tau)$ and random $m \leftarrow \{0,1\}^\tau$, outputs m with some non-negligible probability $\epsilon(\tau)$. How would the *very same* efficient algorithm A do if we used it to attack the encryption scheme Σ?

This takes some gymnastics but it's all very simple (all probabilities are chosen over random keys returned by $KGen(1^\tau)$, so for brevity we'll not make it explicit in the calculations below). Let's also denote $\mathcal{M}_\tau = \{0,1\}^\tau$ by $\mathcal{M}$, $\mathcal{M}'_\tau$ by $\mathcal{M}'$, and denote the complement of $\mathcal{M}'$ as $\overline{\mathcal{M}}' = \mathcal{M} \setminus \mathcal{M}'$.

$$
\begin{aligned}
&Prob[A(Enc(k,m)) = m \mid m \in \mathcal{M}] \;=\\
&Prob[A(Enc(k,m)) = m \mid m \in \mathcal{M}'] * Prob[m \in \mathcal{M}' \mid m \in \mathcal{M}] \;+\\
&Prob[A(Enc(k,m)) = m \mid m \in \overline{\mathcal{M}}'] * Prob[m \in \overline{\mathcal{M}}' \mid m \in \mathcal{M}] \;\geq\\
&Prob[A(Enc(k,m)) = m \mid m \in \mathcal{M}'] * 0 \;+\\
&Prob[A(Enc(k,m)) = m \mid m \in \overline{\mathcal{M}}'] * \left(1 - \frac{p(\tau)}{2^\tau}\right) \;\geq\\
&Prob[A(Enc(k,m)) = m \mid m \in \overline{\mathcal{M}}'] - \frac{p(\tau)}{2^\tau} \;\geq\\
&Prob[A(Enc(k,m)) = m \mid m \in \overline{\mathcal{M}}'] - \mathsf{negl}(\tau) \;=\\
&Prob[A(Enc'(k,m)) = m \mid m \in \overline{\mathcal{M}}'] - \mathsf{negl}(\tau)
\end{aligned}
$$

(the last equation is because Enc and Enc' work the same for $m \in \overline{\mathcal{M}}'$.)

So we see that A's invertion success against Σ on $\mathcal{M}$ can be only negligibly lower than A's invertion success against Σ' on $\overline{\mathcal{M}}'$. But what's that last success? It's easy to see that since the difference $\mathcal{M}' = \mathcal{M} \setminus \overline{\mathcal{M}}'$ between sets $\mathcal{M}$ and $\overline{\mathcal{M}}'$ is negligible in size, A's success against Σ' on $\overline{\mathcal{M}}'$ can be only negligibly different from A's success against Σ' on $\mathcal{M}$:

$$
\begin{aligned}
&Prob[A(Enc'(k,m)) = m \mid m \in \mathcal{M}] \;\; = \\
&Prob[A(Enc'(k,m)) = m \mid m \in \mathcal{M}'] * Prob[m \in \mathcal{M}' \mid m \in \mathcal{M}] \;\; + \\
&Prob[A(Enc'(k,m)) = m \mid m \in \overline{\mathcal{M}}'] * Prob[m \in \overline{\mathcal{M}}' \mid m \in \mathcal{M}] \;\; \leq \\
&1 * \left(\frac{p(\tau)}{2^\tau}\right) \;\; + \\
&Prob[A(Enc'(k,m)) = m \mid m \in \overline{\mathcal{M}}'] * 1 \;\; = \\
&Prob[A(Enc'(k,m)) = m \mid m \in \overline{\mathcal{M}}'] + \frac{p(\tau)}{2^\tau} \;\; \leq \\
&Prob[A(Enc'(k,m)) = m \mid m \in \overline{\mathcal{M}}'] + \mathsf{negl}(\tau)
\end{aligned}
$$

Putting the two together, we get that A's success against the original encryption scheme Σ is not negligible because it is at least:

$$Prob[A(Enc(k,m)) = m \mid m \in \mathcal{M}] \geq Prob[A(Enc'(k,m)) = m \mid m \in \mathcal{M}] - \mathsf{negl}(\tau) \geq \epsilon(\tau) - \mathsf{negl}(\tau)$$

And since $\epsilon(\tau)$ is not a negligible function of τ then neither is $\epsilon(\tau) - \mathsf{negl}(\tau)$. □

3 One-Way Secure Encryption Can Leak Some Message Bits

We can also show that one-way encryption schemes can be one-way secure while leaking some very particular bits of *every* plaintext. In fact, they can leak any constant fraction of the plaintext. Similarly as in the previous section, we can transform any one-way secure encryption $\Sigma = (KGen, Enc, Dec)$ into an encryption scheme $\Sigma' = (KGen, Enc', Dec')$ which remains one-way secure as long as Σ is. The new encryption scheme could leak half the bits of the plaintext, simply by encrypting any message $m = m_{[1,\ldots,n]}$ where $n = |m|$ as

$$c' = Enc'(k,m) = (m_{[1,\ldots,n/2]}, Enc(k, m_{[n/2+1,\ldots,n]}))$$

In this way clearly Enc' leaks half the plaintext bits, and hence cannot be indistinguishable: Just take any m_0, m_1 which differ on any of these leaked bits. And yet at the same time we can prove, similarly to the way we did in the example above, that if Σ' is *not* one-way on a message space $\{0,1\}^n$ then Σ cannot be one-way secure on message space $\{0,1\}^{n/2}$. By taking the counterpositive this shows that if Σ is one-way secure on $\{0,1\}^{n/2}$ then Σ' is one-way secure on $\{0,1\}^n$.