Secret-key encryption from a PRF

Noah Stephens-Davidowitz

June 8, 2023

1 SKE review

First, we recall the definition of a many-message semantically secure secret-key encryption scheme, since it has been a while. We change the definition slightly from the first lecture to allow for the possibility that the message space $\mathcal{M}_n$, key space $\mathcal{K}_n$, and ciphertext space $\mathcal{C}_n$ can depend on the security parameter n. E.g., our final scheme will encrypt n-bit messages when the key has length n.

Definition 1.1 (Encryption scheme)**.** *An* encryption scheme *consists of a plaintext space $\mathcal{M}_n$, a ciphertext space $\mathcal{C}_n$, and a key space $\mathcal{K}_n$ together with three PPT algorithms* $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *that follow the satisfying basic properties.*

1. *The* key generation *algorithm* Gen *takes as input 1^n and outputs a key $k \in \mathcal{K}_n$, i.e., $k \leftarrow \mathsf{Gen}(1^n)$.*
2. *The* encryption *algorithm* Enc *takes as input a key $k \in \mathcal{K}_n$ and a plaintext $m \in \mathcal{M}_n$ and outputs a ciphertext $c \in \mathcal{C}_n$, i.e., $c \leftarrow \mathsf{Enc}(k, m)$.*
3. *The* decryption *algorithm* Dec *takes as input a key $k \in \mathcal{K}_n$ and a ciphertext $c \in \mathcal{C}_n$ and outputs a plaintext $m \in \mathcal{M}_n$, i.e., $m \leftarrow \mathsf{Dec}(k, c)$.*
4. **Correctness:** *For any $k \in \mathcal{K}_n$ and $m \in \mathcal{M}_n$,*

$$\mathsf{Dec}(k, \mathsf{Enc}(k, m)) = m \ .$$

An encryption scheme is many-message semantically secure *if for every PPT adversary, there exists negligible ε such that*

$$\Pr[\mathcal{A} \textit{ wins the many-message semantic security game}] \leq 1/2 + \varepsilon(n) \ ,$$

where the many-message semantic security game is as shown below.

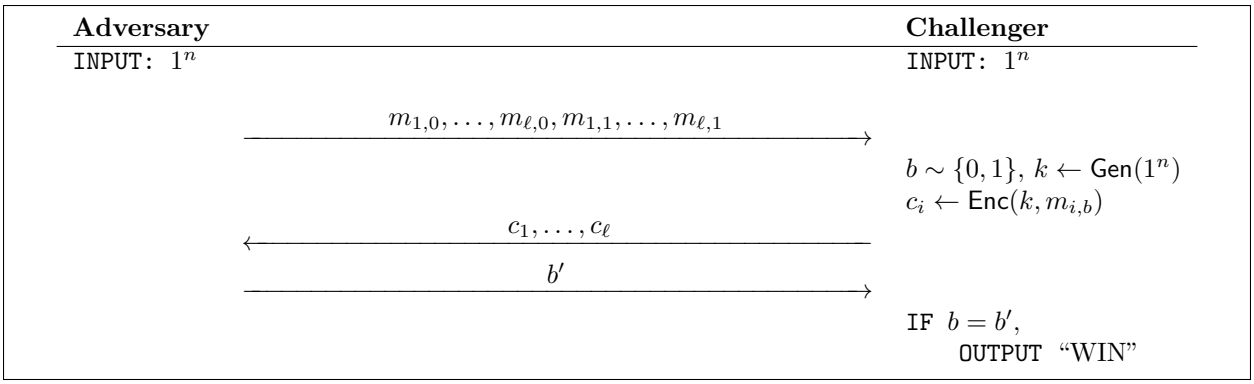

2 Semantic security from a PRF

Let F_k be a PRF (e.g., the GGM PRF that we constructed in the previous lecture [GGM86]). Then, consider the following encryption scheme, whose plaintext space is $\mathcal{M}_n := \{0,1\}^n$.

- $\mathsf{Gen}(1^n)$: Output $k \sim \{0,1\}^n$.
- $\mathsf{Enc}(k, m)$: Sample $r \sim \{0,1\}^n$ and output $(r, c := F_k(r) \oplus m)$.
- $\mathsf{Dec}(k, (r, c))$: Output $F_k(r) \oplus c$.

It is clear that these algorithms are efficient and that the scheme is correct. The tricky bit is, of course, security.

Theorem 2.1. *This scheme is many-message semantically secure if F_k is a PRF.*

The intuition behind the proof is as follows. The adversary in the many-message semantic security game receives the ciphertexts $(r_1, F_k(r_1) \oplus m_{1,b}), \ldots, (r_\ell, F_k(r_\ell) \oplus m_{\ell,b})$. Since the PRF is secure, *intuitively* these ciphertexts should be indistinguishable from $(r_1, y_1 \oplus m_{1,b}), \ldots, (r_\ell, y_\ell \oplus m_{\ell,b})$, where $y_i \sim \{0,1\}^n$ are sampled independently of everything else. But, these new "ciphertexts" are just uniformly random strings, independent of b. So, up to some negligible advantage in distinguishing the PRF from random, the adversary should not be able to win the game with probability better than $1/2$.

Making this precise takes more effort than you might expect. The standard way to do it uses a hybrid argument—although in this context I prefer to refer to it as a *sequence of games*. I.e., we will define a sequence of games Game 1, Game 2, Game 3. Game 1 will be our original game. Game 3 will be defined in a way that makes it obvious that no adversary can win it with probability better than $1/2$. And, for each i, we will argue that no adversary can have non-negligibly larger winning probability in Game i than in Game $i+1$. We will then conclude that no adversary has non-negligible advantage in Game 1. (Different authors differ on how they use the terms "hybrid argument" and "sequence of games." I prefer to reserve the term "hybrid argument" for more systematic sequences of games—when the change from Game i to Game $i+1$ follows some simple rule.)

Proof of Theorem 2.1. We define the following sequence of games. In particular, Games 1 and 2 will differ by replacing F_k with a uniformly random function $H \sim \{f : \{0,1\}^n \to \{0,1\}^n\}$. So, we think of "moving from Game 1 to Game 2 as applying PRF security."

- Game 1 is the many-message semantic security game against our scheme. In particular, the ciphertexts $c_1, \ldots, c_\ell$ sent to the adversary have the form $c_i := (r_i, F_k(r_i) \oplus m_{i,b})$.
- Game 2 is the same as Game 1 except the ciphertext c_i is replaced by $c'_i := (r_i, H(r_i) \oplus m_{i,b})$ where $H \sim \{f : \{0,1\}^n \to \{0,1\}^n\}$ is a uniformly random function.

Claim 2.2. *If the PRF is secure, then for any PPT adversary $\mathcal{A}$ there exists negligible $\varepsilon(n)$ such that*

$$\Pr[\mathcal{A} \textit{ wins Game } 1] - \Pr[\mathcal{A} \textit{ wins Game } 2] \leq \varepsilon(n)$$

Proof. Suppose that $\mathcal{A}$ is an adversary such that

$$\Pr[\mathcal{A} \text{ wins Game } 1] - \Pr[\mathcal{A} \text{ wins Game } 2] = \delta(n)$$

for some non-negligible $\delta(n)$. Then, we construct an adversary $\mathcal{A}'$ in the PRF game as follows.

$\mathcal{A}'$ has oracle access to some oracle $\mathcal{O}$ which is either a random oracle $H \sim \{f : \{0,1\}^n \to \{0,1\}^n\}$ or F_k for $k \sim \{0,1\}^n$. It simply runs $\mathcal{A}$ on input 1^n to receive plaintexts $m_{1,0}, \ldots, m_{\ell,0}, m_{1,1}, \ldots, m_{\ell,1} \in \{0,1\}^n$. Then it flips a coin $b \sim \{0,1\}$, and for $i = 1, \ldots, \ell$, it samples $r_i \sim \{0,1\}^n$ and queries its oracle to compute $y_i := \mathcal{O}(r_i)$.

Finally, $\mathcal{A}'$ sends the ciphertexts $(r_1, y_1 \oplus m_{1,b}), \ldots, (r_\ell, y_\ell \oplus m_{\ell,b})$ to $\mathcal{A}$, receiving as output some bit b'. (Here, $\mathcal{A}'$ is "behaving like the challenger in the semantic security game.") If $b = b'$, $\mathcal{A}'$ outputs 1 (i.e., it guesses that $\mathcal{O} = F_k$). Otherwise, it outputs 0.

Clearly $\mathcal{A}'$ is efficient. Notice that when $\mathcal{O} = F_k$, the view of $\mathcal{A}$ is distributed identically to its view in Game 1,[1] and when $\mathcal{O} = H$, its view is identical to its view in Game 2. It follows that

$$\Pr_{k \sim \{0,1\}^n}[(\mathcal{A}')^{F_k}(1^n) = 1] - \Pr_{H \sim \{f:\{0,1\}^n \to \{0,1\}^n\}}[(\mathcal{A}')^{H}(1^n) = 1] = \Pr[\mathcal{A} \text{ wins Game } 1] - \Pr[\mathcal{A} \text{ wins Game } 2] = \delta(n) \ ,$$

which is non-negligible. This contradicts the assumption that F_k is a PRF, so $\delta(n)$ must be negligible, as needed. □

We could now attempt to reason directly about Game 2. Intuitively, we would like to say that the ciphertexts in Game 2 are just uniformly random strings that are independent of b, so that no adversary can win Game 2 with probability better than 1/2. But, this is not *exactly* true. It would be true if we conditioned on $r_i \neq r_j$ for all $i \neq j$. But, if it happens to be the case that $r_i = r_j$ for some $i \neq j$, then of course $H(r_i) = H(r_j)$, which means that the ciphertexts could depend on b. (Another way of saying this is that if we happen to have $r_i = r_j$ for $i \neq j$, then applying our encryption scheme will result in reuse of a one-time pad. So, somewhere our argument *must* use the fact that this is very unlikely to happen.) Of course, this does not happen very often, but we must account for it. The nicest way to account for it is to simply define a third game in which this is no longer an issue.

[1] This concept of "the view of $\mathcal{A}$" means everything that $\mathcal{A}$ sees in the game. So, here this means the joint distribution of the plaintexts $m_{i,b}$, the ciphertexts c_i, and the bit b'—all of which are random variables.

- In Game 3, the ciphertext c_i is replaced by $(r_i, y_i \oplus m_{b,i})$, where $y_i \sim \{0,1\}^n$ is sampled uniformly at random and independent of everything else.

Claim 2.3. *For any adversary $\mathcal{A}$ (even computationally unbounded adversaries),*

$$\Pr[\mathcal{A} \textit{ wins Game } 2] - \Pr[\mathcal{A} \textit{ wins Game } 3] \leq \ell^2/2^n \ ,$$

where ℓ is the length of the lists of plaintexts chosen by the adversary in the game. In particular, if $\ell \leq \text{poly}(n)$, then this difference is negligible.

Proof. Notice that Game 2 and Game 3 are identical unless there exists an $i \neq j$ such that $r_i = r_j$. For every pair $i \neq j$, $\Pr[r_i = r_j] = 2^{-n}$. There are fewer than q^2 such pairs, and the result follows by union bound. □

Finally, we pedantically observe that no adversary can win Game 3 with probability larger than 1/2.

Claim 2.4. *For any adversary $\mathcal{A}$ (even computationally unbounded adversaries),*

$$\Pr[\mathcal{A} \textit{ wins Game } 3] \leq 1/2 \ .$$

Proof. The ciphertexts $(r_i, y_i \oplus m_{b,i})$ are uniformly random and independent of b. So, the view of the adversary in this game is independent of b, and the result follows. □

Combining the three claims together, we see that for every PPT $\mathcal{A}$ there exists a negligible $\varepsilon^*(n)$ such that

$$\Pr[\mathcal{A} \text{ wins Game } 1] \leq 1/2 + \varepsilon^*(n) \ ,$$

which is exactly what we wanted to prove. (The final value of ε^* that we get is the $\varepsilon^*(n) := \varepsilon(n) + \ell^2/2^n$, where $\varepsilon(n)$ is the value of $\varepsilon(n)$ that we got from Claim 2.2 *plus* $\ell^2/2^n$. Of course, since ℓ is polynomially bounded (because $\mathcal{A}$ is PPT, and $\mathcal{A}$ outputs ℓ strings), $\ell^2/2^n$ is negligible. And, since the sum of two negligible functions is negligible, we conclude that $\varepsilon^*(n)$ is itself negligible.) □

References

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4), 1986.