Secret-Key Encryption

600.442: Modern Cryptography

Fall 2016

The Setting

- Alice and Bob share a secret $s \in \{0,1\}^n$
- Alice wants to send a private message m to Bob
- Goals:
  - **Correctness:** Alice can compute an encoding c of m using s. Bob can decode m from c correctly using s
  - **Security:** No eavesdropper can distinguish between encodings of m and m'

Definition

- **Syntax:**
  - $\mathsf{Gen}(1^n) \rightarrow s$
  - $\mathsf{Enc}(s, m) \rightarrow c$
  - $\mathsf{Dec}(s, c) \rightarrow m'$ or $\perp$

  All algorithms are polynomial time

- **Correctness:** For every m, $\mathsf{Dec}(s, \mathsf{Enc}(s, m)) = m$, where $s \xleftarrow{\$} \mathsf{Gen}(1^n)$

- **Security:** ?

Security

Definition (Indistinguishability Security)

A secret-key encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is secure if for all n.u. PPT adversaries $\mathcal{A}$, there exists a negligible function $\mu(\cdot)$ s.t.:

$$\Pr\left[\begin{array}{r} s \xleftarrow{\$} \mathsf{Gen}(1^n), \\ (m_0, m_1) \leftarrow \mathcal{A}(1^n), \\ b \xleftarrow{\$} \{0,1\} \end{array} : \mathcal{A}\left(\mathsf{Enc}\left(m_b\right)\right) = b\right] \leqslant \frac{1}{2} + \mu(n)$$

1. <u>Think:</u> Computational Indistinguishability style definition?
2. <u>Think</u> How does this definition guarantee that m is "hidden?"

One-Time Pads

- $\mathsf{Gen}(1^n) := s \xleftarrow{\$} \{0,1\}^n$
- $\mathsf{Enc}(s, m) := m \oplus s$
- Security:

$$\mathsf{Enc}\left(s \xleftarrow{\$} \{0,1\}^n, m_0\right) \equiv \mathsf{Enc}\left(s \xleftarrow{\$} \{0,1\}^n, m_1\right)$$

1. Think: Can we use the pad s to encrypt two messages?
2. Think: How to encrypt messages longer than n bits?

Encryption using PRGs

- $\mathsf{Gen}(1^n) := s \xleftarrow{\$} \{0,1\}^n$
- $\mathsf{Enc}(s, m) := m \oplus PRG(s)$
- Security:

$$\mathsf{Enc}\left(s \xleftarrow{\$} \{0,1\}^n, m_0\right) \approx \mathsf{Enc}\left(s \xleftarrow{\$} \{0,1\}^n, m_1\right)$$

Note: m can be polynomially long if we use poly-stretch PRG

- Think: Proof?
- Think: How to encrypt more than one message?

Multi-message Secure Encryption

Definition (Multi-message Secure Encryption)

A secret-key encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is multi-message secure if for all n.u. PPT adversaries $\mathcal{A}$, for all polynomials $q(\cdot)$, there exists a negligible function $\mu(\cdot)$ s.t.:

$$\Pr\left[\begin{array}{r} s \xleftarrow{\$} \mathsf{Gen}(1^n), \\ \left\{\left(m_0^i, m_1^i\right)\right\}_{i=1}^{q(n)} \leftarrow \mathcal{A}(1^n), \\ b \xleftarrow{\$} \{0,1\} \end{array} : \mathcal{A}\left(\left\{\mathsf{Enc}\left(m_b^i\right)\right\}_{i=1}^{q(n)}\right) = b\right] \leqslant \frac{1}{2} + \mu(n)$$

1. Think: Computational Indistinguishability style definition
2. Think Security against *adaptive* adversaries?

Necessity of Randomized Encryption

Theorem (Randomized Encryption)

A multi-message secure encryption scheme cannot be deterministic and stateless.

Think: Proof?

Encryption using PRFs

Let $\{f_s : \{0,1\}^n \rightarrow \{0,1\}^n\}$ be a family of PRFs

- $\mathsf{Gen}(1^n)$: $s \xleftarrow{\$} \{0,1\}^n$
- $\mathsf{Enc}(s, m)$: Pick $r \xleftarrow{\$} \{0,1\}^n$. Output $(r, m \oplus f_s(r))$
- $\mathsf{Dec}(s, (r, c))$: Output $c \oplus f_s(r)$

Theorem (Encryption from PRF)

$(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ *is a multi-message secure encryption scheme*

- Think: Proof?

Proof of Security

Proof via hybrids:

- H_1: Real experiment with $m_0^1, \ldots, m_0^{q(n)}$ (i.e., $b = 0$)
- H_2: Replace f_s with random function $f \xleftarrow{\$} \mathcal{F}_n$
- H_3: Switch to one-time pad encryption
- H_4: Switch to encryption of $m_1^1, \ldots, m_1^{q(n)}$
- H_5: Use random function $f \xleftarrow{\$} \mathcal{F}_n$ to encrypt
- H_6: Encrypt using f_s. Same as real experiment with $m_0^1, \ldots, m_0^{q(n)}$ (i.e., $b = 1$)

Think: Non-adaptive vs adaptive queries

Semantic Security

Definition (Semantic Security)

A secret-key encryption scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ is semantically secure if there exists a PPT simulator algorithm $\mathcal{S}$ s.t. the following two experiments generate computationally indistinguishable outputs:

$$\left\{ \begin{array}{r} (m, z) \leftarrow M(1^n), \\ s \leftarrow \mathsf{Gen}(1^n), \\ \text{Output } (\mathsf{Enc}(s, m), z) \end{array} \right\} \approx \left\{ \begin{array}{l} (m, z) \leftarrow M(1^n), \\ \text{Output } S(1^n, z) \end{array} \right\}$$

where M is a machine that randomly samples a message from the message space and arbitrary auxiliary information.

- Indistinguishability security $\Leftrightarrow$ Semantic security
- Think: Proof?

Food for Thought

Secret-key Encryption in practice:

- Block ciphers with fixed input length (e.g., AES)
- Encryption modes to encrypt arbitrarily long messages (e.g., CBC)
- Stream ciphers for stateful encryption
- Cryptanalysis (e.g., Differential Cryptanalysis)