Provable Security VI

Pseudorandom Permutations and Block Ciphers

Pseudorandom Permutations (PRPs)

Definition 1 (Pseudorandom permutations). Let $P : \{0,1\}^{\kappa(n)} \times \{0,1\}^n \to \{0,1\}^n$ be an efficient keyed permutation ($P_k(x) \stackrel{\text{def}}{=} P(k,x)$). We say P is a pseudorandom permutation if for all probabilistic polynomial-time distinguishers D, there exists a negligible function negl such that:

$$\left| \Pr_{k \xleftarrow{\$} \{0,1\}^{\kappa(n)}} [\mathsf{D}^{P_k(\cdot), P_k^{-1}(\cdot)}(1^n) = 1] - \Pr_{\pi \xleftarrow{\$} \Pi} [\mathsf{D}^{\pi(\cdot), \pi^{-1}(\cdot)}(1^n) = 1] \right| \leq \mathsf{negl}(n)$$

where $k \leftarrow \{0,1\}^{\kappa(n)}$ and $\pi(\cdot)$ is a permutation chosen uniformly at random from the set of all permutations on n-bit strings (denoted by Π).

Informally, the definition of PRP states that no efficient algorithm A can distinguish $\langle P_k(\cdot), P_k^{-1}(\cdot) \rangle$ from $\langle \pi(\cdot), \pi^{-1}(\cdot) \rangle$, where k is randomly chosen from $k \in \{0,1\}^{\kappa(n)}$ and $\pi\colon \{0,1\}^n \to \{0,1\}^n$ randomly is chosen from all permutations over $\{0,1\}^n$. In practice,

Feistel Networks

FEISTEL NETWORKS are a symmetric structure (named after the German cryptographer Horst Feistel) which is used in the construction of pseudorandom permutations (block ciphers in practice).

ONE-ROUND FEISTEL. Given any function $F\colon \{0,1\}^n \to \{0,1\}^n$, we can construct a permutation $D_F\colon \{0,1\}^{2n} \to \{0,1\}^{2n}$ as:

$$D_F(x,y) := (y,\ F(y) \oplus x),$$

where x and y are n-bit strings. Its inverse is given by

$$D_F^{-1}(z,w) := (F(z) \oplus w, z).$$

It is thus easy to verify that D_F is permutation and is efficiently computable (both in forward and backward directions) if F is efficient.

t-ROUND FEISTEL. Given t functions $F_1, \ldots, F_t\colon \{0,1\}^n \to \{0,1\}^n$, the t-round Feistel is simply the composition of the functions $D_{F_t} \circ \ldots \circ D_{F_1}$.

The Luby-Rackoff PRP from PRF

Theorem 1 *For functions* $F : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$, *we define the following permutation (based on 4-round Feistel Networks) given a key* $k = \langle k_1, \ldots, k_4 \rangle$ *and an input* x

$$P : \{0,1\}^{4\kappa} \times \{0,1\}^{2n} \to \{0,1\}^{2n}$$

$$P_k(x) := D_{F_{k_4}}(D_{F_{k_3}}(D_{F_{k_2}}(D_{F_{k_1}}(x)))).$$

If F *is a* (t, q, ε)*-secure pseudorandom function, then* P *is a* $(t - O(q), q, 4\varepsilon + \frac{2q^2}{2^n})$ *secure pseudorandom permutation.*

Note: here we use t *and* q *to denote the running time and query complexity respectively. Namely,* (t,q,ε)*-security refers to every probabilistic algorithm of running time* t *who is bounded to make up to* q *queries gains advantage no more than* ε *in distinguishing the PRF (PRP) from a random function (permutation).*

PROOF SKETCH. Define the following hybrids of functions,

$$P_k^0(x) := D_{F_{k_4}}(D_{F_{k_3}}(D_{F_{k_2}}(D_{F_{k_1}}(x))))$$

$$P_k^1(x) := D_{F_{k_4}}(D_{F_{k_3}}(D_{F_{k_2}}(D_{R_1}(x))))$$

$$P_k^2(x) := D_{F_{k_4}}(D_{F_{k_3}}(D_{R_2}(D_{R_1}(x))))$$

$$P_k^3(x) := D_{F_{k_4}}(D_{R_3}(D_{R_2}(D_{R_1}(x))))$$

$$P_k^4(x) := D_{R_4}(D_{R_3}(D_{R_2}(D_{R_1}(x))))$$

$$P_k^5(x) := \pi(x),$$

where R_1, R_2, R_3, R_4 are independent random functions from $\{\{0,1\}^{2n} \to \{0,1\}^{2n}\}$ and π is a random permutation from all permutations over $\{0,1\}^{2n}$. By triangle inequality, we have

$$\begin{aligned} & |\Pr[\mathsf{D}^{P_k^0(\cdot),P_k^{0\,-1}(\cdot)}(1^n)=1] \;-\; \Pr[\mathsf{D}^{P_k^5(\cdot),P_k^{5\,-1}}(1^n)=1]| \\ \leq \quad & \sum_{i=0}^{4} |\Pr[\mathsf{D}^{P_k^i(\cdot),P_k^{i\,-1}(\cdot)}(1^n)=1] \;-\; \Pr[\mathsf{D}^{P_k^{i+1}(\cdot),P_k^{i+1\,-1}}(1^n)=1]| \end{aligned}$$

thus it suffices to bound $|\Pr[\mathsf{D}^{P_k^i(\cdot),P_k^{i\,-1}(\cdot)}(1^n)=1] \;-\; \Pr[\mathsf{D}^{P_k^{i+1}(\cdot),P_k^{i+1\,-1}}(1^n)=1]|$ for every $0 \leq i \leq 4$.

Some lemmas for Theorem 1

Lemma 2. *For every* A *making up to* q *queries,*

$$\left| \Pr[\mathsf{D}^{P_k^4(\cdot),P_k^{4-1}(\cdot)}(1^n)=1] - \Pr[\mathsf{D}^{P_k^5(\cdot),P_k^{5-1}}(1^n)=1] \right| \leq \frac{2q^2}{2^n}$$

Note that the bound is information-theoretic, i.e., D *is only bounded by query complexity (no restrictions on running time).*

The proof is a bit involved and isn't a typical reduction type of proof, so we're not introducing it in this course (or giving any exercises on this).

Lemma 3. *If* F *is a* (t,q,ε)*-secure pseudorandom function, then for every* D *of running time* $t-O(q)$ *and making up to* q *queries, and for every* $0 \leq i \leq 3$

$$\left| \Pr[\mathsf{D}^{P_k^i(\cdot),P_k^{i-1}(\cdot)}(1^n)=1] - \Pr[\mathsf{D}^{P_k^{i+1}(\cdot),P_k^{i+1-1}}(1^n)=1] \right| \leq \varepsilon$$

The proof is quite like the one we did in lecture 5.

Constructions of Block Ciphers

- In theory:

  one-way functions → pseudorandom generators

→ pseudorandom functions →pseudorandom permutations

- In practice: from scratch
- Advanced Encryption Standard (AES)
- Data Encryption Standard (DES)

Chosen Ciphertext Attacks (CCA) security

THE CCA INDISTINGUISHABILITY EXPERIMENT $\mathsf{PrivK}^{cca}_{\mathsf{A},\Pi}(n)$.

1. A random key k is generated: $k \leftarrow \mathsf{Gen}(1^n)$.
2. The adversary A is given input 1^n and oracle access to $\mathsf{Enc}_k(\cdot)$ and $\mathsf{Dec}_k(\cdot)$, and outputs a pair of messages m_0, m_1 of the same length.
3. A random bit $b \leftarrow \{0,1\}$ is chosen, and then a challenge ciphertext $c \leftarrow \mathsf{Enc}_k(m_b)$ and $c \leftarrow \mathsf{Dec}_k(m_b)$ are computed and given to A.
4. The adversary A continues to have oracle access to $\mathsf{Enc}_k(\cdot)$ and $\mathsf{Dec}_k(\cdot)$, but is not allowed to query the latter on the challenge ciphertext c itself. Finally, it outputs a bit b'.
5. The output of the experiment is defined to be 1 iff $b' = b$ (indicating that A succeeded), and 0 otherwise.

Definition 2 (CCA security)**.** A private-key encryption scheme $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ has indistinguishable encryptions under a chosen-ciphertext attack (or is CCA- secure) if for all probabilistic polynomial-time adversaries A there exists a negligible function negl such that

$$\Pr[\mathsf{PrivK}^{cca}_{\mathsf{A},\Pi}(n) = 1] \leq \frac{1}{2} + \mathsf{negl}(n)$$

where the probability is taken over the random coins used by A, as well as the random coins used in the experiment.