EE 720: An Introduction to Number Theory and Cryptography (Spring 2018)

Lecture 8 — February 2, 2018

Instructor: Saravanan Vijayakumaran *Scribe: Saravanan Vijayakumaran*

1 Lecture Plan

- See examples of stream ciphers used in practice.
- Define CPA-security

2 Recap

A Secure Fixed-Length Encryption Scheme

- Let G be a pseudorandom generator with expansion factor l. Define a private-key encryption scheme for messages of length l as follows:
  - `Gen`: On input 1^n, choose k uniformly from $\{0,1\}^n$.
  - `Enc`: Given $k \in \{0,1\}^n$ and message $m \in \{0,1\}^{l(n)}$, output the ciphertext

    $$c := G(k) \oplus m.$$

  - `Dec`: Given $k \in \{0,1\}^n$ and ciphertext $c \in \{0,1\}^{l(n)}$, output the message

    $$m := G(k) \oplus c.$$

3 Stream Ciphers

- Stream ciphers are practical systems which behave like pseudorandom generators. However, there are no proofs available that a particular stream cipher is in fact a pseudorandom generator.
- Stream ciphers can be designed for either efficient hardware implementation or efficient software implementation.
- Hardware-oriented stream ciphers are based on feedback shift registers (FSRs).
- Linear feedback shift registers (LFSRs) are FSRs where the feedback function is linear.
- Example: Consider a four-bit shift register where the feedback is the XOR of all the four bits. If we initialize the state to 1100, then we get a cycle of period 5. The states are 1100, 1000, 0001, 0011, 0110.

- The output depends on the state of the LFSR. Once a state repeats, the output repeats. If an LFSR has n bits, then the period of the output sequence can be at most $2^n - 1$.
- Each LFSR can be associated with a feedback polynomial. If the feedback polynomial is primitive, then the period is maximal. A polynomial of degree n over GF(2) is primitive if it is irreducible and the smallest value of m for which the polynomial divides $X^m + 1$ is $m = 2^n - 1$. Example: $1 + X^3 + X^4$.

3.1 A5/1

- Used to provide voice encryption in the GSM cellular system.
- Developed in 1987. Reverse engineered in 1999.
- Uses three LFSRs of lengths 19, 22, and 23.
- More details at `https://en.wikipedia.org/wiki/A5/1`.

3.2 RC4

- A software-oriented cipher designed by Ron Rivest of RSA Security in 1987. Reverse engineered and leaked in 1994.
- Has an internal state of 256 bytes initialized to $S[i] = i$ for $i = 0, 1, \ldots, 255$.
- More details on pages 92–93 in Chapter 5 of Serious Cryptography.
- It took 20 years for cryptanalysts to find flaws. Used in WEP (the first generation Wi-fi security protocol) and TLS (the protocol underlying HTTPS).
- **No longer recommended for use**.

3.3 Salsa20

- A software-oriented stream cipher announced by Daniel J. Bernstein in 2005. Part of eSTREAM software portfolio.
- It is counter-based stream cipher which generated 512-bit blocks of keystream at a time.
- More details on pages 95–97 in Chapter 5 of Serious Cryptography.

4 Chosen-Plaintext Attacks and CPA-Security

- Consider a scenario where the honest parties share a key k and the attacker can influence these parties to encrypt messages $m_1, m_2, \ldots$ using k. At some later point, the attacker observes the encryption of a message m (using the same key k). He even knows m is one of the messages $m_1, m_2, \ldots$. Security against chosen-plaintext attacks means that the attacker cannot tell which message was encrypted with probability significantly better than random guessing.

- Real-world chosen-plaintext attacks: WWII British mine locations, Battle of Midway
- Formally, chosen-plaintext attacks are modeled by giving the adversary $\mathcal{A}$ access to an *encryption oracle*. It can be considered a black box which encrypts messages of $\mathcal{A}$'s choosing using a key k which is unknown to $\mathcal{A}$.
- Consider the following experiment $\texttt{PrivK}^{\texttt{cpa}}_{\mathcal{A},\Pi}(n)$:
  1. A key k is generated by running $\texttt{Gen}(1^n)$.
  2. The adversary $\mathcal{A}$ is given 1^n and oracle access to $\texttt{Enc}_k(\cdot)$, and outputs a pair of messages $m_0, m_1 \in \mathcal{M}$ with $|m_0| = |m_1|$.
  3. A uniform bit $b \in \{0, 1\}$ is chosen. Ciphertext $c \leftarrow \texttt{Enc}_k(m_b)$ is computed and given to $\mathcal{A}$.
  4. The adversary $\mathcal{A}$ continues to have oracle access to $\texttt{Enc}_k(\cdot)$, and outputs a bit b'.
  5. The output of the experiment is defined to be 1 if $b' = b$, and 0 otherwise. If output is 1, we say that $\mathcal{A}$ succeeds.

Definition. *A private-key encryption scheme $\Pi = (\texttt{Gen}, \texttt{Enc}, \texttt{Dec})$ has **indistinguishable encryptions under a plaintext attack**, or is **CPA-secure**, if for all probabilistic polynomial-time adversaries $\mathcal{A}$ there is a negligible function $\texttt{negl}$ such that, for all n,*

$$\Pr\left[\texttt{PrivK}^{\texttt{cpa}}_{\mathcal{A},\Pi}(n) = 1\right] \le \frac{1}{2} + \texttt{negl}(n).$$

5 References and Additional Reading

- Chapter 5 of *Serious Cryptography* by J.-P. Aumasson.
- Section 3.4 from Katz/Lindell