PRPs and block ciphers

Recommended reading

KL, Sections 3.5.1 and 3.6.2

We have seen that PRFs yield CPA-secure encryption, and that they can be constructed from any PRG.

Can we push the notion of a PRF even further?

We introduce pseudorandom permutations. These are keyed permutations that are easy to compute and invert with the key, but that are indistinguishable from a uniformly random permutation without the key.

Def (PRPs): A function $F: \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^*$ is a pseudorandom permutation (PRP) if:

- For each $k \in \{0,1\}^*$, $F_k: \{0,1\}^{|k|} \to \{0,1\}^{|k|}$ is a permutation (bijective).
- There is a deterministic poly-time algorithm that given $k, x \in \{0,1\}^n$ outputs $F_k(x) := F(k,x)$
- There is a deterministic poly-time algo that given $k, y \in \{0,1\}^n$ outputs x such that $F_k(x) = y$.

→ Let $\mathcal{O}_{real,k}$ be an oracle that on input $x \in \{0,1\}^n$ outputs $F_k(x)$. Let $\mathcal{O}_{ideal}$ be an oracle that on input x behaves as follows: If x hasn't been queried before, sample $y \leftarrow \{0,1\}^n$ conditioned on y not having been output by the oracle before and output y. If x has been queried before, answer consistently with previous queries.

Then, for every PPT adv A there is a negl function $\epsilon(n)$ such that

$$\left| \Pr_{k \leftarrow \{0,1\}^n}\left(A^{\mathcal{O}_{real,k}}(1^n)=1\right) - \Pr\left(A^{\mathcal{O}_{ideal}}(1^n)=1\right)\right| \leq \epsilon(n).$$

Requiring only that a PRP F be indistinguishable from a random permutation is a purely aesthetic choice.

Thm: If F is a PRP then it is also a PRF.

Proof: Homework. Intuition: For a random function, the probability that two queries lead to an output collision is very small. Bound the probability of seeing a collision.

PRPs (also sometimes with extra useful properties) are commonly known in practice as "block ciphers". More on this in a bit.

It is natural to wonder how we can construct PRPs. It turns out that PRGs suffice! We will not see a proof. If you're interested, see KL, Section 7.6.

In practice... Constructions of concrete block ciphers do not follow the transformation from PRGs. Rather, they are ad hoc and heavily optimized for efficiency

Examples: DES, AES → based on a selection process by NIST

See KL, Section 6.2 for some discussion on this.

Block cipher modes of operation

We saw a CPA-secure scheme based on any PRF.
While this is great, this scheme has a drawback:
To encrypt an n-bit message we produce a $2n$-bit ciphertext. That's a large overhead when encrypting many messages.

Therefore, practitioners don't use this scheme, but rather use PRPs to encrypt traffic in other ways. These are called modes of operation.

The setting: We wish to encrypt several messages $m_1, m_2, \ldots, m_\ell$, with small overhead!

For much more on this, see

https://www.cs.ucdavis.edu/~rogaway/papers/modes.pdf

Some modes of operation

Let F be a PRP.

Electronic codebook (ECB) mode:

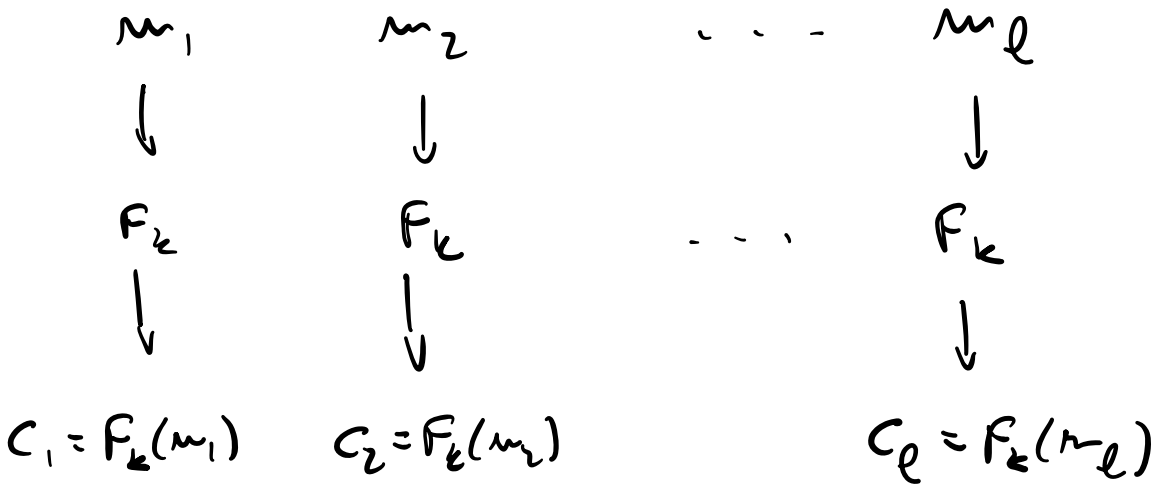

No overhead! However, offers poor security... For example, it is deterministic, so cannot be CPA-secure.

Cipher Block Chaining (CBC) Mode:

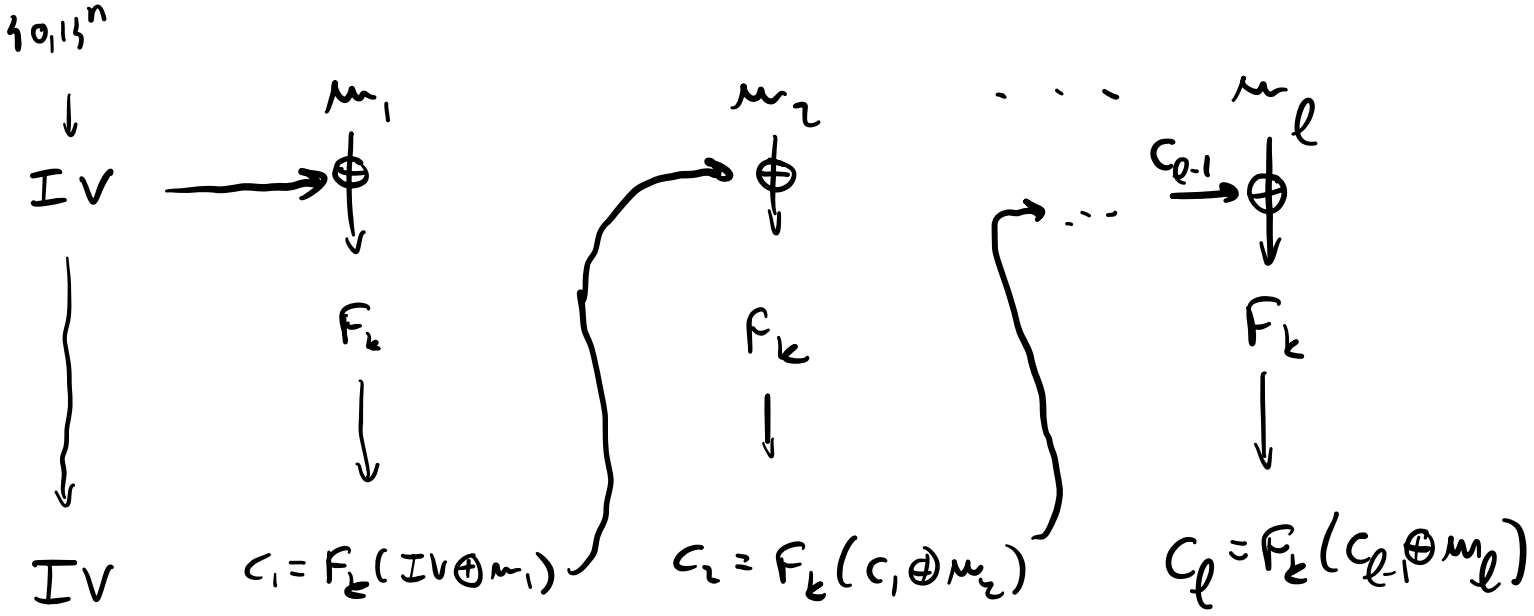

Only overhead is the random initialization vector.

F is a PRP $\Rightarrow$ CBC mode is CPA-secure.

Con: Cannot be parallelized...

Counter (CTR) Mode:

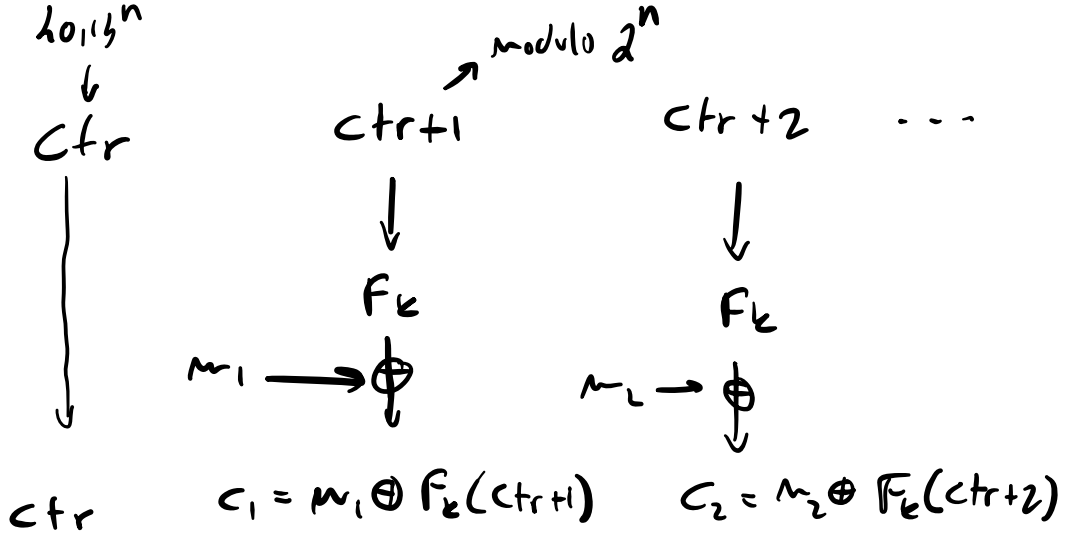

CTR mode is CPA-secure if F is a PRF.

Can be parallelized!

A note on chosen-ciphertext attacks

Recommended reading: KL, Section 3.7.1

We have seen CPA-security, where the adversary has adaptive access to an encryption oracle. This is meant to model chosen-plaintext attacks.

But what if the adversary can learn **decryptions** of ciphertexts of their choice? These are called **chosen-ciphertext attacks**.

A real-world instance (in a different context): Bleichenbacher's attack on PKCS#1 v1.5.

https://archiv.infsec.ethz.ch/education/fs08/secsem/bleichenbacher98.pdf

How can we define "CCA-secure" encryption? Just like CPA-security, except we allow adversaries access also to a decryption oracle that on input c outputs $\mathrm{Dec}(k,c)$.

Recall our CPA-secure encryption scheme

$k \leftarrow \{0,1\}^n$

$Enc(k, M) = (r, F_k(r) \oplus M)$ for $r \leftarrow \{0,1\}^n$, F a PRF

This isn't CCA-secure!

Consider the following attack:

→ A chooses $m_0 = 0^n$, $m_1 = 1^n$ and learns

$$c = (r, y), \quad y = F_k(r) \oplus m_b.$$

→ A computes $c' = (r, y \oplus (1, 0, \dots, 0))$

c' to the decryption oracle to get back m'.

Note that c' decrypts to $m_b \oplus (1, 0, \dots, 0)$.

If $b = 0$, then $m' = 10^{n-1}$.

If $b = 1$, then $m' = 01^{n-1}$.

$\Rightarrow$ A guesses b with probability 1.