Cryptology complementary

✧

Symmetric modes of operation

Pierre Karpman
pierre.karpman@univ-grenoble-alpes.fr
https://www-ljk.imag.fr/membres/Pierre.Karpman/tea.html

2018–05–03

From two weeks ago

- A good *primitive* $\neq$ a good *cryptographic scheme*
  - Example: RSA (a good OWF w/ trapdoor) is not a good encryption scheme
  - $\rightsquigarrow$ need padding (e.g. OAEP)
  - Ditto for signatures (use e.g. PSS-R)
- This is true for asymmetric crypto (above)
- But also symmetric (today's topic)

Block cipher recalls

- Recall that a (binary) block cipher is a mapping $\mathcal{E} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ s.t. $\forall k \in \{0,1\}^\kappa$, $\mathcal{E}(k,\cdot)$ is a permutation
- A "good" block cipher is a family of permutations that "look random" and are independent of each other $\rightsquigarrow$ PRP-security
- Some implications for good BCs:
    - It is hard to find an unknown k given oracle access to $\mathcal{E}(k,\cdot)$
    - It is hard to find m given $c = \mathcal{E}(k,m)$ for an unknown k
    - It is hard to find $c = \mathcal{E}(k,m)$ for an unknown k given m
    - Etc.

Block ciphers are not enough

What block ciphers do:

- One-to-one encryption of fixed-size messages

What do we want:

- One-to-many encryption of variable-size messages
- Why?
  - Variable-size → kind of obvious?
  - One-to-many → necessary for "semantic security" → cannot tell if two ciphertexts are of the same message or not

Enter modes of operation

- A *mode of operation* transforms a block cipher into a *symmetric encryption scheme*
- $\approx \mathcal{E} \rightsquigarrow \mathsf{Enc}: \{0,1\}^\kappa \times \{0,1\}^r \times \{0,1\}^* \to \{0,1\}^*$
- For all $k \in \{0,1\}^\kappa$, $r \in \{0,1\}^r$, $\mathsf{Enc}(k, r, \cdot)$ is invertible
- $\{0,1\}^r$, $r \geq 0$ is used to make encryption non-deterministic
- A mode is "good" if it gives "good encryption schemes" when used with "good BCs"
- So what's a good encryption scheme?

IND-CPA for Symmetric encryption

IND-CPA for Enc: An adversary cannot distinguish $\mathsf{Enc}(k, m_0)$ from $\mathsf{Enc}(k, m_1)$ for an unknown key k and equal-length messages m_0, m_1 when given oracle access to an $\mathsf{Enc}(k, \cdot)$ oracle:

1. The Challenger chooses a key $k \xleftarrow{\$} \{0,1\}^{\kappa}$
2. The Adversary may repeatedly submit queries x_i to the Challenger
3. The Challenger answers a query with $\mathsf{Enc}(k, r_i, x_i)$
4. The Adversary now submits m_0, m_1 of equal length
5. The Challenger draws $b \xleftarrow{\$} \{0,1\}$, answers with $\mathsf{Enc}(k, r', m_b)$
6. The Adversary tries to guess b

- The choice of r_i, r' is defined by the mode (made explicit here, may be omitted)

IND-CPA comments

- A random adversary succeeds w/ prob. $1/2 \rightarrow$ the correct success measure is the *advantage* over this
  - Advantage (one possible definbition): $|\Pr[\text{Adversary answers } 1 : b = 0] - \Pr[\text{Adversary answers } 1 : b = 1]|$
- An adversary may always succeed w/ advantage 1 given enough ressources
  - Find the key spending time $t \leq 2^{\kappa}$ and a few oracle queries
- What matters is the "best possible" advantage in function of the attack complexity

First (non-) mode example: ECB

- ECB: just concatenate independent calls to $\mathcal{E}$

Electronic Code Book mode

$m_0 || m_1 || \ldots \mapsto \mathcal{E}(k, m_0) || \mathcal{E}(k, m_1) || \ldots$

- No security
  - Exercise: give a simple attack on ECB for the IND-CPA security notion w/ advantage 1, low complexity

Second (actual) mode example: CBC

- Cipher Block Chaining: Chain blocks together (duh)

Cipher Block Chaining mode

$r \times m_0||m_1||\ldots \mapsto c_0 := \mathcal{E}(k, m_0 \oplus r)||c1 := \mathcal{E}(k, m_1 \oplus c_0)||\ldots$

- Output block i (ciphtertext) added (XORed) w/ input block $i+1$ (plaintext)
- For first (m_0) block: use random IV r
- Okay security in theory $\rightsquigarrow$ okay security in practice *if used properly*

CBC IVs

CBC has bad IND-CPA security if the IVs are not random

- Consider an IND-CPA adversary who asks an oracle query $\mathrm{CBC}-\mathrm{ENC}(m)$, gets $r, c = \mathcal{E}(k, m \oplus r)$ (where $\mathcal{E}$ is the cipher used in $\mathrm{CBC}-\mathrm{ENC}$)
- Assume the adversary knows that for the next IV r', $\Pr[r' = x] = p$
- Sends two challenges $m_0 = m \oplus r \oplus x$, $m_1 = m_0 \oplus 1$
- Gets $c_b = \mathrm{CBC}-\mathrm{ENC}(m_b)$, $b \xleftarrow{\$} \{0, 1\}$
- If $c_b = c$, guess $b = 0$, else $b = 1$
    - Exercise: what is the adversary's advantage? (If $q := \Pr[r' = x \oplus 1] \leq (1 - p)$.)

Generic CBC collision attack

Even with random IVs, CBC has some drawbacks
An observation:

- In CBC, inputs to $\mathcal{E}$ are of the form $x \oplus y$ where x is a message block and y an IV or a ciphertext block
- If $x \oplus y = x' \oplus y'$, then $\mathcal{E}(k, x \oplus y) = \mathcal{E}(k, x' \oplus y')$

A consequence:

- If $c_i = \mathcal{E}(k, m_i \oplus c_{i-1}) = c'_j = \mathcal{E}(k, m'_j \oplus c'_{j-1})$, then $c_{i-1} \oplus c'_{j-1} = m_i \oplus m'_j$
- $\rightsquigarrow$ knowing identical ciphertext blocks reveals information about the message blocks
- $\Rightarrow$ breaks IND-CPA security
- Regardless of the security of $\mathcal{E}$!

CBC collisions: how likely?

How soon does a collision happen?

- Proposition: the distribution of the $(x \oplus y)$ is $\approx$ uniform
  - If y is an IV it has to be (close to) uniformly random, otherwise we have an attack (two slides ago)
  - If $y = \mathcal{E}(k, z)$ is a ciphertext block, ditto for y knowing z, otherwise we have an attack on $\mathcal{E}$
- $\Rightarrow$ A collision occurs w.h.p. after $\sqrt{\#\{0,1\}^n} = 2^{n/2}$ blocks are observed (with identical key k) $\leftarrow$ *The birthday bound*
- (Slightly more precisely, w/ prob. $\approx q^2/2^n, q \leq 2^{n/2}$ after q blocks)

Some CBC recap

A decent mode, but

- Must use random IVs
- Must change key *much* before encrypting $2^{n/2}$ blocks when using an n-bit block cipher
- And this *regardless of the key size* κ
- This is a common restriction for modes of operation (cf. next slide)

Another classical mode: CTR

Counter mode

$m_0||m_1||\ldots \mapsto \mathcal{E}(k, \mathtt{s}\texttt{++}) \oplus m_0||\mathcal{E}(k, \mathtt{s}\texttt{++}) \oplus m_1||\ldots$

- This uses a global state $\mathtt{s}$ for the *counter*, with C-like semantics for $\mathtt{s}\texttt{++}$
- Encrypts a public counter $\rightsquigarrow$ pseudo-random keystream $\rightsquigarrow$ (perfect) one-time-pad approximation (i.e. a *stream cipher*)
- Like CBC, must change key *much* before encrypting $2^{n/2}$ blocks when using an n-bit block cipher
  - Question: why?

How to go further: the tweakable option

- A (binary) *tweakable* block cipher is a mapping $\tilde{\mathcal{E}} : \{0,1\}^\kappa \times \{0,1\}^\theta \times \{0,1\}^n \to \{0,1\}^n$ s.t. $\forall k \in \{0,1\}^\kappa, t \in \{0,1\}^\theta$, $\tilde{\mathcal{E}}(k,t,\cdot)$ is a permutation
- The *tweak* t is "like a key", but known & may be chosen by the adversary
- A necessary condition for $\tilde{\mathcal{E}}$ to be a good TBC is for $\tilde{\mathcal{E}}(\cdot,t,\cdot)$ to be a good BC for all t.
  - But an adversary may further try to exloit relations between $\tilde{\mathcal{E}}$ for $\neq$ tweaks

TBC constructions

How to build a TBC?

- From scratch, like any block cipher (see for instance Jean et al., 2014)
- From an existing block cipher treated as a black box (see for instance Liskov et al., 2002)
- Still a quite active research topic

A simple (not ideal) example:

- $\tilde{\mathcal{E}}(k, t, \cdot) := \mathcal{E}(k \oplus t, \cdot)$
- (Relies on the analysis of $\mathcal{E}$ in a *XOR-Related-key setting*)

TBC: why?

- Many modes (like CBC) fail when encrypting too many blocks with the same permutation
- $\rightsquigarrow$ Change permutation as often as possible
- Change key at every block?
  - Not so clean to define, possible efficiency issues
- $\rightsquigarrow$ Add a tweak, change tweak at every block
  - Clean, possibly more efficient, but a more "complex" primitive

A simple mode for TBCs: TIE

- "Like ECB", but with distinct tweaks for every call to $\tilde{\mathcal{E}}$

Tweak Incrementation Encryption

$m_0||m_1||\ldots \mapsto c_0 := \tilde{\mathcal{E}}(k, \mathtt{s}{+}{+}, m_0)||c1 := \tilde{\mathcal{E}}(k, \mathtt{s}{+}{+}, m_1)||\ldots$

- Again uses a global state $\mathtt{s}$, this time for the tweak
- Security directly reduces to the one of $\tilde{\mathcal{E}}$ as long as tweaks don't repeat
  - Intuitively if $\tilde{\mathcal{E}}(k, t, \cdot)$ and $\tilde{\mathcal{E}}(k, t' \neq t, \cdot)$ are independent random permutations, $\tilde{\mathcal{E}}(k, t, x)$ and $\tilde{\mathcal{E}}(k, t', x')$ are independent random values for any x, x'

To go even further

- TBCs are great to define *authenticated encryption* (AE) modes, like TAE
- *Authentication*: "Only someone knowing the key k knows how to create and verify 'valid' messages"
- (Beyond the scope of this course)

About the exam

- One hour out of the three
- Probably $\approx$ two independent exercises
- Mostly on symmetric notions