A parallelizable MAC (PMAC) — general idea:

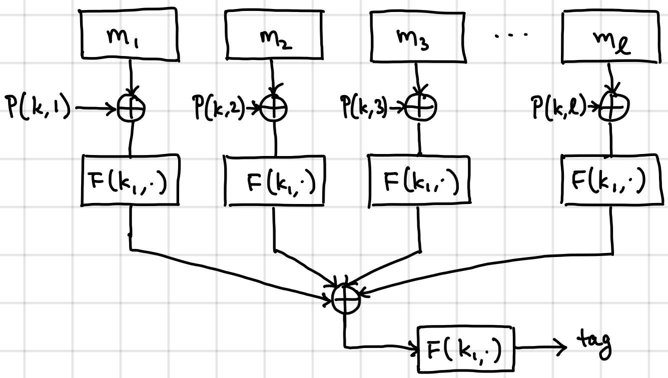

$P(k,\cdot)$ are important — otherwise, adversary can permute the blocks
(derived as $F(k_1, 0^n)$ — so key is just k_1)

↳ "mask" term is of the form $\gamma_i \cdot k$ where multiplication is done over $GF(2^n)$ where n is the block size (constants γ_i carefully chosen for efficient evaluation)

Can use similar ideas as CMAC (randomized prefix-free encoding) to support messages that is not constant multiple of block size

Parallel structure of PMAC makes it easily updateable (assuming F is a PRP)

↳ suppose we change block i from $m[i]$ to $m'[i]$:

compute $F^{-1}(k_1, \text{tag}) \oplus \underbrace{F(k_1, m[i] \oplus P(k,i))}_{\text{old value}} \oplus \underbrace{F(k_1, m'[i] \oplus P(k,i))}_{\text{new value}}$

} PMAC is "incremental": can make local updates without full recomputation

In terms of performance:

- On sequential machine, PMAC comparable to ECBC, NMAC, CMAC
- On parallel machine, PMAC much better

} Best MAC we've seen so far, but not used... Reason: patents :([not patented anymore!]

Summary: Many techniques to build a large-domain PRF from a small-domain one (domain extension for PRF)

↳ Each method (ECBC, CMAC, PMAC) gives a MAC on variable-length messages

↳ Many of these designs (or their variants) are standardized

How do we combine confidentiality and integrity?

↳ Systems with both guarantees are called authenticated encryption schemes – gold standard for symmetric encryption

Two natural options:

1. Encrypt-then MAC (TLS 1.2+, IPsec) ← guaranteed to be secure if we instantiate using CPA-secure encryption and a secure MAC
2. MAC-then-encrypt (SSL 3.0/TLS 1.0, 802.11i) ← as we will see, not always secure

Definition. An encryption scheme $\Pi_{SE} = (\text{Encrypt}, \text{Decrypt})$ is an authenticated encryption scheme if it satisfies the following two properties:

- CPA security [confidentiality]
- ciphertext integrity [integrity]

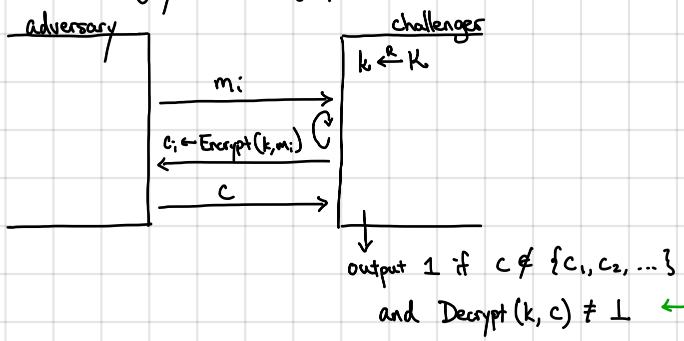

special symbol $\bot$ to denote invalid ciphertext

Define $\text{CIAdv}[A, \Pi_{SE}]$ to be the probability that output of above experiment is 1. The scheme Π_{SE} satisfies ciphertext integrity if for all efficient adversaries A,

$$\text{CIAdv}[A, \Pi_{SE}] = \text{negl}(\lambda)$$

↑ security parameter determines key length

Ciphertext integrity says adversary cannot come up with a new ciphertext: only ciphertexts it can generate are those that are already valid. Why do we want this property?

Consider the following active attack scenario:

- Each user shares a key with a mail server
- To send mail, user encrypts contents and send to mail server
- Mail server decrypts the email, re-encrypts it under recipient's key and delivers email

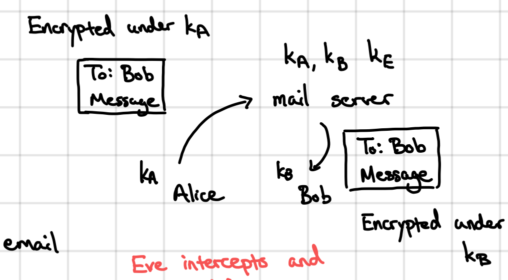

If Eve is able to tamper with the encrypted message, then she is able to learn the encrypted contents (even if the scheme is CPA-secure)

↳ More broadly, an adversary can tamper and inject ciphertexts into a system and observe the user's behavior to learn information about the decrypted values – against active attackers, we need stronger notion of security

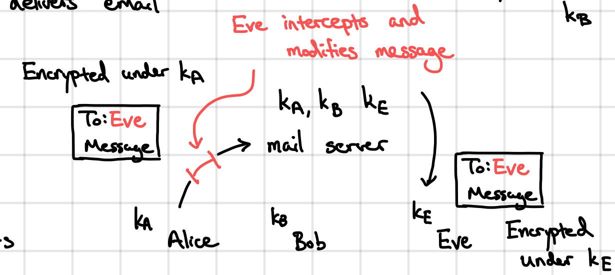

Definition. An encryption scheme $\Pi_{SE} = (\text{Encrypt}, \text{Decrypt})$ is secure against chosen-ciphertext attacks (CCA-secure) if for all efficient adversaries $\mathcal{A}$, $\text{CCAAdv}[\mathcal{A}, \Pi_{SE}] = \text{negl.}$ where we define $\text{CCAAdv}[\mathcal{A}, \Pi_{SE}]$ as follows:

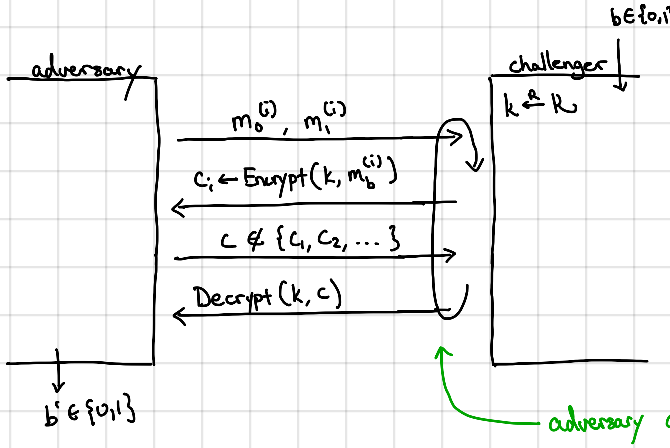

adversary can make arbitrary encryption and decryption queries, but cannot decrypt any ciphertexts it received from the challenger (otherwise, adversary can trivially break security)
↳ called an "admissibility" criterion

$$\text{CCAAdv}[\mathcal{A}, \Pi_{SE}] = \left| \Pr[b' = 1 \mid b = 0] - \Pr[b' = 1 \mid b = 1] \right|$$

CCA-security captures above attack scenario where adversary can tamper with ciphertexts
- ↳ Rules out possibility of transforming encryption of $x \| z$ to encryption of $y \| z$
- ↳ Necessary for security against <u>active</u> adversaries [CPA-security is for security against <u>passive</u> adversaries]
- ↳ We will see an example of a real CCA attack in HW1

Theorem. If an encryption scheme Π_{SE} provide authenticated encryption, then it is CCA-secure.

Proof (Idea). Consider an adversary $\mathcal{A}$ in the CCA-security game. Since Π_{SE} provides ciphertext integrity, the challenger's response to the adversary's decryption query will be $\perp$ with all but negligible probability. This means we can implement the decryption oracle with the "output $\perp$" function. But then this is equivalent to the CPA-security game.
[Formalize using a hybrid argument]

simple counter-example: concatenate unused bits to end of ciphertext in a CCA-secure scheme (stripped away during decryption)

Note: Converse of the above is not true since CCA-security $\not\Rightarrow$ ciphertext integrity.
- ↳ However, CCA-security + plaintext integrity $\Rightarrow$ authenticated encryption

Take-away: Authenticated encryption captures meaningful confidentiality + integrity properties; provides <u>active</u> security

Encrypt-then-MAC: Let (Encrypt, Verify) be a CPA-secure encryption scheme and (Sign, Verify) be a secure MAC. We define Encrypt-then-MAC to be the following scheme:

$\text{Encrypt}'((k_E, k_M), m)$: $c \leftarrow \text{Encrypt}(k_E, m)$
$t \leftarrow \text{Sign}(k_M, c)$
output (c, t)

(k_E, k_M: independent keys)

$\text{Decrypt}'((k_E, k_M), (c, t))$: if $\text{Verify}(k_M, c, t) = 0$, output $\perp$
else, output $\text{Decrypt}(k_E, c)$

Theorem. If (Encrypt, Decrypt) is CPA-secure and (Sign, Verify) is a secure MAC, then (Encrypt', Verify') is an authenticated encryption scheme

Proof. (Sketch). CPA-security follows by CPA-security of (Encrypt, Decrypt). Specifically, the MAC is computed on ciphertexts and not the messages. MAC key is independent of encryption key so cannot compromise CPA-security.

Ciphertext integrity follows directly from MAC security (i.e., any valid ciphertext must contain a new tag on some ciphertext that was not given to the adversary by the challenger)

Important notes:
- Encryption + MAC keys must be independent. Above proof required this (in the formal reduction, need to be able to simulate ciphertexts / MACs — only possible if reduction can choose its own key).
  - ↳ Can also give explicit constructions that are completely broken if same key is used (i.e., both properties fail to hold)
  - ↳ In general, never reuse cryptographic keys in different schemes; instead, sample fresh, independent keys!
- MAC needs to be computed over the entire ciphertext
  - Early version of ISO 19772 for AE did not MAC IV (CBC used for CPA-secure encryption)
  - RNCryptor in Apple iOS (for data encryption) also problematic (HMAC not applied to encryption IV)

  ] means first block (i.e., "header") is malleable

MAC-then-Encrypt: Let (Encrypt, Verify) be a CPA-secure encryption scheme and (Sign, Verify) be a secure MAC. We define MAC-then-Encrypt to be the following scheme:

Encrypt'$((k_E, k_M), m)$: $t \leftarrow \text{Sign}(k_M, m)$
$c \leftarrow \text{Encrypt}(k_E, (m,t))$
output c

Decrypt'$((k_E, k_M), (c,t))$: compute $(m,t) \leftarrow \text{Decrypt}(k_E, c)$
if $\text{Verify}(k_M, m, t) = 1$, output m, else, output $\perp$

Not generally secure! SSL 3.0 (precursor to TLS) used randomized CBC + secure MAC
↳ Simple CCA attack on scheme (by exploiting padding in CBC encryption)
[POODLE attack on SSL 3.0 can decrypt all encrypted traffic using a CCA attack]
Padding is a common source of problems with MAC-then-Encrypt systems [see HW2 for an example]

In the past, libraries provided separate encryption + MAC interfaces — common source of errors
↳ Good library design for crypto should minimize ways for users to make errors, not provide more flexibility

Today, there are standard block cipher modes of operation that provide authenticated encryption
- One of the most widely used is GCM (Galois counter mode) — standardized by NIST in 2007

GCM mode: follows encrypt-then-MAC paradigm
- CPA-secure encryption is nonce-based counter mode
- MAC is a Carter-Wegman MAC
  - ↳ "encrypted one-time MAC"

} Most commonly used in conjuction with AES (AES-GCM provides authenticated encryption)

GCM encryption: encrypt message with AES in counter mode

compute Carter-Wegman MAC on resulting message using GHASH as the underlying hash function

and the block cipher as underlying PRF

(Galois Hash — GHASH)

(key derived from PRF evaluation at 0^n)

↳ GHASH operates on blocks of 128-bits

operations can be expressed as operations over $GF(2^{128})$ — Galois field with 2^{128} elements

implemented in hardware — very fast!

Typically, use AES-GCM for authenticated encryption

Oftentimes, only part of the payload needs to be hidden, but still needs to be authenticated

↳ e.g., sending packets over a network: desire confidentiality for packet body, but only integrity for packet headers (otherwise, cannot route!)

AEAD: authenticated encryption with associated data

↳ augment encryption scheme with additional plaintext input; resulting ciphertext ensures integrity for associated data, but not confidentiality (will not define formally here but follows straightforwardly from AE definitions)

↳ can construct directly via "encrypt-then-MAC": namely, encrypt payload and MAC the ciphertext + associated data

↳ AES-GCM is an AEAD scheme