Introduction to Cryptology

9.2 - Hybrid Encryption

Federico Pintore

Mathematical Institute, University of Oxford (UK)

UNIVERSITY OF OXFORD

Michaelmas term 2020

Hybrid Encryption

A combination of a public-key scheme and a symmetric-key encryption can be used to deal with arbitrary-length messages.

Hybrid Encryption

A combination of a public-key scheme and a symmetric-key encryption can be used to deal with arbitrary-length messages.

- The public-key primitive, called key-encapsulation mechanism (KEM), is used to obtain a shared key.
- The shared key is used with a symmetric-key encryption scheme, called data-encapsulation mechanism.

Hybrid Encryption

A combination of a public-key scheme and a symmetric-key encryption can be used to deal with arbitrary-length messages.

- The public-key primitive, called key-encapsulation mechanism (KEM), is used to obtain a shared key.
- The shared key is used with a symmetric-key encryption scheme, called data-encapsulation mechanism.

Symmetric-key encryption schemes are significantly faster (2 or 3 orders of magnitude) than public ones.

Key-encapsulation mechanisms (KEMs)

A key-encapsulation mechanism $(\mathrm{KeyGen}, \mathrm{Encaps}, \mathrm{Decaps})$ consists of three algorithms:

- $(\mathrm{PK}, \mathrm{SK}) \leftarrow \mathrm{KeyGen}(n)$: on input a security parameter n, it returns a pair of keys $(\mathrm{PK}, \mathrm{SK})$ - the public key PK and its matching secret key SK - each of length n.
- $(c, k) \leftarrow \mathrm{Encaps}(\mathrm{PK}, n)$: on input a public key PK and n, it outputs a ciphertext c and a key $k \in \{0,1\}^{\ell(n)}$.
- $k/\perp \leftarrow \mathrm{Decaps}(\mathrm{SK}, c)$: deterministic algorithm that takes a secret key SK and a ciphertext c, and returns a key k or $\perp$.

Correctness: for any $(\mathrm{PK}, \mathrm{SK})$ output by KeyGen on input n it holds

$$\Pr(\mathrm{Decaps}(\mathrm{SK}, c) = k | (c, k) \leftarrow \mathrm{Encaps}(\mathrm{PK}, n)) = 1$$

KEMs - Definition of Security

CPA Indistinguishability $\text{KEM}^{\text{cpa}}_{\mathcal{A},\Pi}(n)$

KEMs - Definition of Security

CPA Indistinguishability $\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A},\Pi}(n)$

Challenger Ch		Adversary $\mathcal{A}$
$(\mathrm{PK}, \mathrm{SK}) \leftarrow \mathrm{KeyGen}(n)$		
$(c, k) \leftarrow \mathrm{Encaps}(\mathrm{PK}, n)$		
$b \leftarrow \{0, 1\}$		
If $b = 0$, $\hat{k} := k$ else $\hat{k} \leftarrow \{0, 1\}^{\ell(n)}$	$\xrightarrow{(\mathrm{PK}, c, \hat{k})}$	
		Output their guess b'

KEMs - Definition of Security

CPA Indistinguishability $\text{KEM}^{\text{cpa}}_{\mathcal{A},\Pi}(n)$

Challenger Ch		Adversary $\mathcal{A}$
$(\text{PK}, \text{SK}) \leftarrow \text{KeyGen}(n)$		
$(c, k) \leftarrow \text{Encaps}(\text{PK}, n)$		
$b \leftarrow \{0, 1\}$		
If $b = 0$, $\hat{k} := k$ else $\hat{k} \leftarrow \{0, 1\}^{\ell(n)}$	$\xrightarrow{(\text{PK}, c, \hat{k})}$	
		Output their guess b'

$\mathcal{A}$ wins the game, i.e. $\text{KEM}^{\text{cpa}}_{\mathcal{A},\Pi}(n) = 1$, if $b' = b$.

KEMs - Definition of Security

CPA Indistinguishability $\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A},\Pi}(n)$

Challenger Ch		Adversary $\mathcal{A}$
$(\mathrm{PK}, \mathrm{SK}) \leftarrow \mathrm{KeyGen}(n)$		
$(c, k) \leftarrow \mathrm{Encaps}(\mathrm{PK}, n)$		
$b \leftarrow \{0, 1\}$		
If $b = 0$, $\hat{k} := k$ else $\hat{k} \leftarrow \{0, 1\}^{\ell(n)}$	$\xrightarrow{(\mathrm{PK}, c, \hat{k})}$	
		Output their guess b'

$\mathcal{A}$ wins the game, i.e. $\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A},\Pi}(n) = 1$, if $b' = b$.

Definition

A KEM Π *is CPA-secure if, for every PPT adversary $\mathcal{A}$, it holds*

$$\mathrm{Adv}^{\mathrm{cpa}}_{\mathcal{A},\Pi}(n) = \Pr(\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A},\Pi}(n) = 1) \leq 1/2 + \mathrm{negl}(n)\,.$$

Hybrid Encryption

A hybrid encryption scheme $(\text{KeyGen}^{hy}, \text{Enc}^{hy}, \text{Dec}^{hy})$ is a public-key encryption scheme obtained combining a KEM $\Pi = (\text{KeyGen}, \text{Encaps}, \text{Decaps})$ and a symmetric-key encryption scheme $E = (\text{KeyGen}', \text{Enc}, \text{Dec})$ as follows.

- $(\text{PK}, \text{SK}) \leftarrow \text{KeyGen}^{hy}(n)$: it runs KeyGen on input a security parameter n, and returns its output (PK, SK).
- $(c, c') \leftarrow \text{Enc}^{hy}(\text{PK}, m \in \{0,1\}^*)$: given a public key PK and a message m it
  - computes $(c, k) \leftarrow \text{Encaps}(\text{PK}, n)$;
  - computes $c' \leftarrow \text{Enc}(k, m)$;
  - outputs the ciphertext (c, c').
- $m \leftarrow \text{Dec}^{hy}(\text{SK}, (c, c'))$: on input a secret key SK and a ciphertext (c, c'), it
  - computes $k \leftarrow \text{Decaps}(\text{SK}, c)$;
  - outputs $m \leftarrow \text{Dec}(k, c')$.

Hybrid Encryption: Efficiency

Consider $\alpha = \text{cost}(\text{Encaps}(\cdot, n))$ and $\beta = \text{cost}(\text{Enc}(\cdot, 1 \text{ bit}))$ for a fixed security parameter n.

Hybrid Encryption: Efficiency

Consider $\alpha = \text{cost}(\text{Encaps}(\cdot, n))$ and $\beta = \text{cost}(\text{Enc}(\cdot, 1 \text{ bit}))$ for a fixed security parameter n.

To encrypt a message m, the cost per bit is:

$$\text{cost}(\text{Enc}^{hy}(\cdot, 1 \text{ bit})) = \frac{\alpha + \beta \cdot |m|}{|m|} = \frac{\alpha}{|m|} + \beta.$$

Hybrid Encryption: Efficiency

Consider $\alpha = \text{cost}(\text{Encaps}(\cdot, n))$ and $\beta = \text{cost}(\text{Enc}(\cdot, 1 \text{ bit}))$ for a fixed security parameter n.

To encrypt a message m, the cost per bit is:

$$\text{cost}(\text{Enc}^{hy}(\cdot, 1 \text{ bit})) = \frac{\alpha + \beta \cdot |m|}{|m|} = \frac{\alpha}{|m|} + \beta.$$

For a sufficiently long m, $\text{cost}(\text{Enc}^{hy}(\cdot, 1 \text{ bit}))$ approaches β, i.e.

$$\text{cost}(\text{Enc}^{hy}(\cdot, 1 \text{ bit})) \approx \text{cost}(\text{Enc}(\cdot, 1 \text{ bit})).$$

Security of the Hybrid Encryption Scheme

Theorem

Consider the hybrid encryption scheme E^{hy}. If

- $\Pi = (\mathrm{KeyGen}, \mathrm{Encaps}, \mathrm{Decaps})$ *is a CPA-secure key-encapsulation mechanism,*
- $E = (\mathrm{KeyGen}', \mathrm{Enc}, \mathrm{Dec})$ *is a symmetric-key encryption scheme which has indistinguishable encryptions in the presence of an eavesdropper,*

then E^{hy} is a CPA-secure public-key encryption scheme.

Security of the Hybrid Encryption Scheme

Proof

Let $\mathcal{A}^{hy}$ be an adversary playing the $\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A}^{hy},S^{hy}}(n)$ game. The goal is proving that:

$$\Pr(\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A}^{hy},S^{hy}}(n) = 1) \leq \frac{1}{2} + \mathrm{negl}(n)\,.$$

Security of the Hybrid Encryption Scheme

Proof

Let $\mathcal{A}^{hy}$ be an adversary playing the $\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A}^{hy},\mathcal{S}^{hy}}(n)$ game. The goal is proving that:

$$\Pr(\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A}^{hy},\mathcal{S}^{hy}}(n) = 1) \leq \frac{1}{2} + \mathrm{negl}(n)\,.$$

From the union formula and the definition of conditional probability we deduce:

$$\Pr(\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A}^{hy},\mathcal{S}^{hy}}(n) = 1) = \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0 | m = m_0) + \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1 | m = m_1).$$

Security of the Hybrid Encryption Scheme

Using $\mathcal{A}^{hy}$ as a subroutine, we construct an adversary $\mathcal{A}_1$ against the CPA-security of Π.

- $\mathcal{A}_1$ receives $(\mathrm{PK}, c, \hat{k})$ from Ch and sends PK to $\mathcal{A}^{hy}$;
- upon reception of (m_0, m_1) from $\mathcal{A}^{hy}$, it obtains c' running Enc on input $\hat{k}$ and m_0, and sends (c, c') to $\mathcal{A}^{hy}$;
- $\mathcal{A}_1$ outputs the bit b' received from $\mathcal{A}^{hy}$.

Security of the Hybrid Encryption Scheme

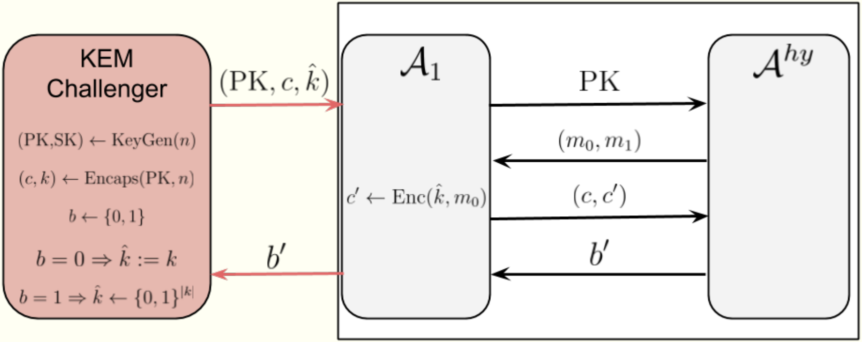

$$\Pr(\mathcal{A}_1 \text{ outputs } 0 | b = 0) = \Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k, m = m_0)$$
$$\Pr(\mathcal{A}_1 \text{ outputs } 1 | b = 1) = \Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k', m = m_0)$$

Security of the Hybrid Encryption Scheme

Since the key-encapsulation scheme Π is CPA-secure, we have:

$$\begin{aligned}
\Pr(\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A}_1,\Pi}(n) = 1) &= \frac{1}{2}\Pr(\mathcal{A}_1 \text{ outputs } 0 | b = 0) + \\
&+ \frac{1}{2}\Pr(\mathcal{A}_1 \text{ outputs } 1 | b = 1) = \\
&= \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k, m = m_0) + \\
&+ \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k', m = m_0) \leq \\
&\leq \frac{1}{2} + \mathrm{negl}_1(n)
\end{aligned}$$

Security of the Hybrid Encryption Scheme

Using $\mathcal{A}^{hy}$ as a subroutine, we construct an adversary $\mathcal{A}_2$ against the CPA-security of Π.

- $\mathcal{A}_2$ receives $(\mathrm{PK}, c, \hat{k})$ from Ch and sends PK to $\mathcal{A}^{hy}$;
- upon reception of (m_0, m_1) from $\mathcal{A}^{hy}$, it obtains c' running Enc on input $\hat{k}$ and m_1, and sends (c, c') to $\mathcal{A}^{hy}$;
- $\mathcal{A}_2$ outputs $1 - b'$, where b' is the bit received from $\mathcal{A}^{hy}$.

Security of the Hybrid Encryption Scheme

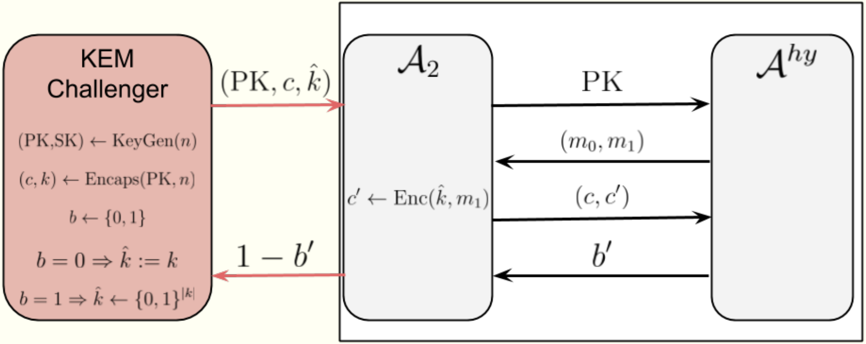

$$\Pr(\mathcal{A}_2 \text{ outputs } 0 | b = 0) = \Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k, m = m_1)$$
$$\Pr(\mathcal{A}_2 \text{ outputs } 1 | b = 1) = \Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k', m = m_1)$$

Security of the Hybrid Encryption Scheme

Since the key-encapsulation scheme Π is CPA-secure, we have:

$$\begin{aligned}
\Pr(\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A}_2,\Pi}(n) = 1) &= \frac{1}{2}\Pr(\mathcal{A}_2 \text{ outputs } 0 | b = 0) + \\
&+ \frac{1}{2}\Pr(\mathcal{A}_2 \text{ outputs } 1 | b = 1) = \\
&= \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k, m = m_1) + \\
&+ \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k', m = m_1) \leq \\
&\leq \frac{1}{2} + \mathrm{negl}_2(n)
\end{aligned}$$

Security of the Hybrid Encryption Scheme

Using $\mathcal{A}^{hy}$ as a subroutine, we construct an adversary $\mathcal{A}'$ against the indistinguishability of E.

- $\mathcal{A}'$ runs KeyGen, obtaining $(\mathrm{PK}, \mathrm{SK})$. They compute $(c, k) \leftarrow \mathrm{Encaps}(\mathrm{PK}, n)$ and send PK to $\mathcal{A}^{hy}$.
- Upon reception of (m_0, m_1) from $\mathcal{A}^{hy}$, $\mathcal{A}'$ sends them to the challenger, receiving a ciphertext c';
- $\mathcal{A}'$ sends (c, c') to $\mathcal{A}^{hy}$.
- $\mathcal{A}'$ outputs the bit b' received from $\mathcal{A}^{hy}$.

Security of the Hybrid Encryption Scheme

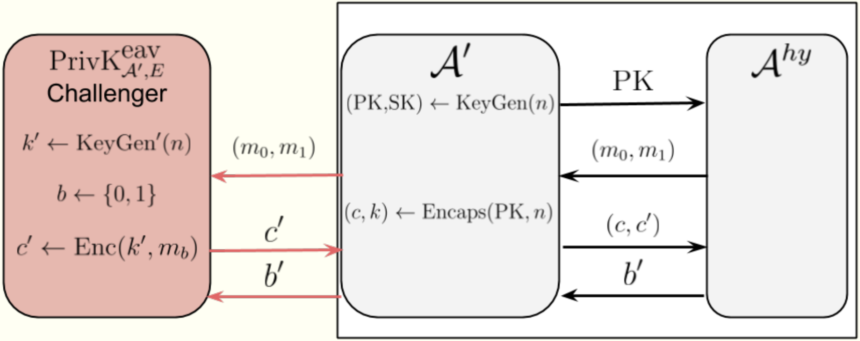

$$\Pr(\mathcal{A}' \text{ outputs } 0 | b = 0) = \Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k', m = m_0)$$
$$\Pr(\mathcal{A}' \text{ outputs } 1 | b = 1) = \Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k', m = m_1)$$

Security of the Hybrid Encryption Scheme

The symmetric-key encryption scheme E has indistinguishable encryptions in the presence of an eavesdropper. Therefore:

$$\begin{aligned}
\Pr(\mathrm{PrivK}^{\mathrm{eav}}_{\mathcal{A}',E}(n) = 1) &= \frac{1}{2}\Pr(\mathcal{A}' \text{ outputs } 0 | b = 0) + \\
&+ \frac{1}{2}\Pr(\mathcal{A}' \text{ outputs } 1 | b = 1) = \\
&= \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k', m = m_0) + \\
&+ \frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k', m = m_1) \leq \\
&\leq \frac{1}{2} + \mathrm{negl}'(n)
\end{aligned}$$

Security of the Hybrid Encryption Scheme

$\text{negl}_1(n) + \text{negl}_2(n) + \text{negl}'(n)$ is a negligible function $\text{negl}(n)$.

Summing all the above inequalities we obtain:

$$
\begin{aligned}
&\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k, m = m_0) + \\
&\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k', m = m_0) + \\
&\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k, m = m_1) + \\
&\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k', m = m_1) + \\
&\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0 | \hat{k} = k', m = m_0) + \\
&\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1 | \hat{k} = k', m = m_1) \\
&\leq \frac{3}{2} + \text{negl}(n).
\end{aligned}
$$

Security of the Hybrid Encryption Scheme

Furthermore, we have:

$$\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1|\hat{k}=k', m=m_0)+$$
$$\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0|\hat{k}=k', m=m_0)=\frac{1}{2}$$

and

$$\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0|\hat{k}=k', m=m_1)+$$
$$\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1|\hat{k}=k', m=m_1)=\frac{1}{2}.$$

Security of the Hybrid Encryption Scheme

Hence, it remains

$$\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 0|\hat{k} = k, m = m_0) +$$
$$\frac{1}{2}\Pr(\mathcal{A}^{hy} \text{ outputs } 1|\hat{k} = k, m = m_1) =$$
$$\Pr(\mathrm{PubK}^{\mathrm{eav}}_{\mathcal{A}^{hy},S^{hy}}) \leq \frac{1}{2} + \mathrm{negl}(n)\,,$$

which concludes the proof.

□

Security of the Hybrid Encryption Scheme

The definition of CCA-security of a KEM relies on an game, similar to $\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A},\Pi}(n)$, where $\mathcal{A}$ is also given access to a decapsulation oracle $\mathrm{Decaps}(\mathrm{SK}, \cdot)$.

Security of the Hybrid Encryption Scheme

The definition of CCA-security of a KEM relies on an game, similar to $\mathrm{KEM}^{\mathrm{cpa}}_{\mathcal{A},\Pi}(n)$, where $\mathcal{A}$ is also given access to a decapsulation oracle $\mathrm{Decaps}(\mathrm{SK}, \cdot)$.

Theorem

If Π is a CCA-secure key-encapsulation mechanism and E is a CCA-secure symmetric-key encryption scheme, the corresponding hybrid encryption scheme E^{hy} is a CCA-secure public-key encryption scheme.

Further Reading I

- Mihir Bellare, Alexandra Boldyreva, and Silvio Micali.
  Public-Key Encryption in a Multi-user Setting: Security Proofs and Improvements.
  In Bart Preneel, editor, Advances in Cryptology — EUROCRYPT 2000, volume 1807 of Lecture Notes in Computer Science, pages 259–274. Springer Berlin Heidelberg, 2000.

- Dan Boneh.
  Simplified OAEP for the RSA and Rabin Functions.
  In Joe Kilian, editor, Advances in Cryptology — CRYPTO 2001, volume 2139 of Lecture Notes in Computer Science, pages 275–291. Springer Berlin Heidelberg, 2001.

Further Reading II

Ronald Cramer and Victor Shoup.
Design and analysis of practical public-key encryption schemes secure against adaptive chosen ciphertext attack.
SIAM Journal on Computing, 33(1):167–226, 2003.

Whitfield Diffie and Martin E Hellman.
New directions in cryptography.
Information Theory, IEEE Transactions on, 22(6):644–654, 1976.

Itai Dinur, Orr Dunkelman, Nathan Keller, and Adi Shamir.
New attacks on Feistel Structures with Improved Memory Complexities.
In Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I, pages 433–454, 2015.

Further Reading III

Naofumi Homma, Atsushi Miyamoto, Takafumi Aoki, Akashi Satoh, and Adi Shamir.
Collision-Based Power Analysis of Modular Exponentiation Using Chosen-Message Pairs.
In Cryptographic Hardware and Embedded Systems - CHES 2008, 10th International Workshop, Washington, D.C., USA, August 10-13, 2008. Proceedings, pages 15–29, 2008.