Security Protocols IX: The Computational Soundness of Formal Encryption

Christoph Sprenger and Sebastian Mödersheim
Department of Computer Science
ETH Zurich

FMSEC Module 10, v.2
November 23, 2009

Outline

Two Views of Cryptography

	Formal View	Computational View
	$\{\!\{K_1\}\!\}_{K_2} \in \mathcal{DY}(\{\{\!\{\!\{K_1\}\!\}_{K_2}\}\!\}_{K_3}, K_3\})$	$Pr[k, k' \xleftarrow{R} \mathcal{K}(\eta) : A^{\mathcal{E}_k(\cdot),\mathcal{E}_{k'}(\cdot)}(\eta) = 1] -$ $Pr[k \xleftarrow{R} \mathcal{K}(\eta) : A^{\mathcal{E}_k(\perp),\mathcal{E}_k(\perp)}(\eta) = 1]$
Messages	Symbolic terms.	Strings of bits.
Adversary	Dolev-Yao deduction.	Probabilistic polynomial-time Turing machine.
Assumes	Perfect cryptography.	Hard problems (e.g., factoring, DL).
Analysis	Tool-supported analysis.	Manual reduction proofs.
	Two separate research communities!	

Abadí and Rogaway started to bridge the gap and relate the two views:

$$M \cong N \qquad \text{implies} \qquad [\![M]\!]_\Pi \approx [\![N]\!]_\Pi$$

$\Longrightarrow$ "Statements of the formal model also hold in the real world"

Motivation

Computational Soundness

- This lecture discusses the approach by Abadí and Rogaway (early 2000's) covering passive adversaries.
- This work has paved the way for a very active area of research linking the formal and cryptographic views: The computational soundness of formal protocol models.

Benefits

- Conceptual clarification of the relation between these two views, including their assumptions.
- Enables combination of advantages: we may use symbolic tools (efficiency) to obtain cryptographic guarantees (precision).

Outline

Messages as Terms

Definition (Terms)

$M, N ::=$	terms	
K	key	(for $K \in Keys$)
$\mid b$	bit	(for $b \in Bool = \{0, 1\}$)
$\mid (M, N)$	pair	
$\mid \{\!\mid M \mid\!\}_K$	encryption	(for $K \in Keys$)

- This work only considers symmetric encryption and atomic keys.
- Interpreted in free algebra, e.g.,

$$\{\!\mid M \mid\!\}_K = \{\!\mid M' \mid\!\}_{K'} \text{ iff } M = M' \text{ and } K = K'$$

Key cycles

Definition (Cyclic messages)

- We say that K encrypts K' in M if there is a subterm $\{\!|N|\!\}_K$ of M such that N contains an occurrence of K'.
- We say that M is cyclic if the relation "encrypts in M" on keys is cyclic.

Examples

- K encrypts both K_1 and K_2 in $\{\!|\{\!|K_2|\!\}_{K_1}|\!\}_K$.
- $(1, \{\!|K|\!\}_K)$ is cyclic.
- $(\{\!|K_1|\!\}_{K_2}, \{\!|K_2|\!\}_{K_1})$ is cyclic.
- $(\{\!|K_1|\!\}_{K_2}, \{\!|0|\!\}_{K_1})$ is acyclic.

The soundness theorem is restricted to acyclic messages.

Message Deduction by Adversary

Definition (Adversary deduction)

$$\frac{M \in H}{M \in \mathcal{DY}(H)}\ \text{Inject} \qquad \frac{b \in Bool}{b \in \mathcal{DY}(H)}\ \text{Bool}$$

$$\frac{M_1 \in \mathcal{DY}(H) \quad M_2 \in \mathcal{DY}(H)}{(M_1, M_2) \in \mathcal{DY}(H)}\ \text{Pair} \qquad \frac{(M_1, M_2) \in \mathcal{DY}(H)}{M_i \in \mathcal{DY}(H)}\ \text{Proj}_i$$

$$\frac{K \in \mathcal{DY}(H) \quad M \in \mathcal{DY}(H)}{\{\!|M|\!\}_K \in \mathcal{DY}(H)}\ \text{Enc} \qquad \frac{\{\!|M|\!\}_K \in \mathcal{DY}(H) \quad K \in \mathcal{DY}(H)}{M \in \mathcal{DY}(H)}\ \text{Dec}$$

Example

For $H = \{K_3, \{\!|0, \{\!|K_1|\!\}_{K_2}|\!\}_{K_3}\}$ we have:

$$(\{\!|K_1|\!\}_{K_2}, \{\!|1|\!\}_{K_3}) \in \mathcal{DY}(H)$$

Patterns

Definition (Patterns)

$P, Q ::=$	patterns	
K	key	(for $K \in Keys$)
$\mid b$	bit	(for $b \in Bool$)
$\mid (M, N)$	pair	
$\mid \{\!\mid M \mid\!\}_K$	encryption	(for $K \in Keys$)
$\mid \square$	undecryptable	

- Intuitively, a pattern is a term that may have some parts that the intruder cannot decrypt.
- These undecryptable parts are marked by $\square$.
- We want to define a function $p(M, T)$ that, given a message M and a set of known keys T, replaces the undecryptable parts of M by $\square$.

Patterns from Terms

Definition (Pattern of a term M)

With respect to a set of keys T:

$$
\begin{array}{rcll}
p(K,T) & = & K & \text{for } K \in \mathit{Keys} \\
p(b,T) & = & b & \text{for } b \in \mathit{Bool} \\
p((M,N),T) & = & (p(M,T), p(N,T)) & \\
p(\{\!|M|\!\}_K, T) & = & \begin{cases} \{\!|p(M,T)|\!\}_K & \text{if } K \in T \\ \square & \text{otherwise} \end{cases} &
\end{array}
$$

With respect to the keys that are derivable from the message M:

$$\mathit{pattern}(M) = p(M, \mathit{Keys} \cap \mathcal{DY}(\{M\}))$$

Example

$$
\begin{array}{rcl}
\mathit{pattern}((\{\!|\{\!|K_1|\!\}_{K_2}, 0|\!\}_{K_3}, K_3)) & = & p((\{\!|\{\!|K_1|\!\}_{K_2}, 0|\!\}_{K_3}, K_3), \{K_3\}) \\
& = & (\{\!|\square, 0|\!\}_{K_3}, K_3)
\end{array}
$$

Equivalence of Messages

Idea: The adversary cannot distinguish messages with the same pattern.

Definition (Equivalence)

Two terms are equivalent if they yield the same pattern:

$$M \equiv N \quad \text{if and only if} \quad \mathit{pattern}(M) = \mathit{pattern}(N)$$

Examples

- $(\{\!|\{\!|K_1|\!\}_{K_2}|\!\}_{K_3}, K_3) \equiv (\{\!|\{\!|0|\!\}_{K_1}|\!\}_{K_3}, K_3)$
  Both terms yield the same pattern: $(\{\!|\Box|\!\}_{K_3}, K_3)$.
- $(\{\!|0|\!\}_K, K) \not\equiv (\{\!|0|\!\}_{K'}, K')$
  These terms differ only by a renaming of keys.

Definition (Equivalence up to renaming)

$$M \cong N \quad \text{if and only if} \quad \text{if there is a bijection } \sigma \text{ on } \mathit{Keys} \text{ such that } M \equiv N\sigma.$$

Exploring the Definition of Equivalence

Examples (basic)

- $0 \cong 0, 0 \not\cong 1$
- $(K, \{\!|0|\!\}_K) \not\cong (K, \{\!|1|\!\}_K)$ Decryptable messages
- $\{\!|0|\!\}_K \cong \{\!|1|\!\}_K$ Undecryptable messages
- $\{\!|0|\!\}_K \cong \{\!|1|\!\}_{K'}$ and even $\{\!|0|\!\}_K \equiv \{\!|1|\!\}_{K'}$

Examples (Properties of encryption)

- $\{\!|((1,1),(1,1)),(1,1))|\!\}_K \cong \{\!|0|\!\}_K$
  Cannot deduce the size of a plaintext from ciphertext.
- $(\{\!|0|\!\}_K, \{\!|0|\!\}_K) \cong (\{\!|0|\!\}_K, \{\!|1|\!\}_K)$
  Cannot distinguish two encryptions of a plaintext under given key.
- $(\{\!|0|\!\}_K, \{\!|1|\!\}_K) \cong (\{\!|0|\!\}_K, \{\!|1|\!\}_{K'})$
  Cannot detect whether two ciphertexts use the same key.

Outline

Preliminaries: Indistinguishability & Co.

Definitions

- A probability ensemble is a collection of distributions on strings, $D = \{D_\eta\}$, one for each $\eta \in \mathbb{N}$.
- We write $x \xleftarrow{R} D_\eta$ to indicate that x is sampled from D_η.
- We write $Pr[x \xleftarrow{R} D_\eta : E]$ for the probability of event E when x is sampled from D_η.
- A function $\varepsilon : \mathbb{N} \to \mathbb{R}$ is negligible if for all $c > 0$ there exists a N_c such that $\varepsilon(\eta) \leq \eta^{-c}$ for all $\eta \geq N_c$.
- Two ensembles $D = \{D_\eta\}$ and $D' = \{D'_\eta\}$ are (computationally) indistinguishable, written $D \approx D'$, if for every probabilistic polynomial time adversary, the following function is negligible:

$$\varepsilon(\eta) = \left| Pr[x \xleftarrow{R} D_\eta : A(\eta, x) = 1] - Pr[x \xleftarrow{R} D'_\eta : A(\eta, x) = 1] \right|$$

Here, η is called the security parameter.

Encryption Scheme (1)

Domains

- Plaintext, Ciphertext, Key $\subseteq$ String $= \{0,1\}^*$,
- Parameter $= 1^*$.

Definition (Encryption scheme)

An encryption scheme is a triple $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ where

- $\mathcal{K}$ is a probabilistic key generation algorithm with
  input: a security parameter $1^\eta \in$ Parameter, and
  output: a key $\mathcal{K}(1^\eta) \in$ Key,
- $\mathcal{E}$ is a probabilistic encryption algorithm with
  input: a key $k \in$ Key and a message $m \in$ String, and
  output: a ciphertext $\mathcal{E}_k(m) \in$ Ciphertext.
- $\mathcal{D}$ is a deterministic decryption algorithm with
  input: a key $k \in$ Key and a message $c \in$ String, and
  output: a cleartext $\mathcal{D}_k(c) \in$ Plaintext.

Each algorithm runs in time polynomial in the size of its input.

Encryption Schemes (2)

Assumptions

- $\perp \in$ Plaintext,
- For $k \in \mathcal{K}(1^\eta)$ and $\eta \in$ Parameter:

$$\mathcal{D}_k(\mathcal{E}_k(m)) = \begin{cases} m & \text{if } m \in \text{Plaintext} \\ \perp & \text{otherwise} \end{cases}$$

- $|\mathcal{E}_k(x)|$ depends only on η and $|x|$ when $k \in \mathcal{K}(1^\eta)$.

Type-0 Security (1)

Repetition concealing (r) Given ciphertexts c and c', are their underlying plaintexts equal?

Formal view: $(\{\!|0|\!\}_K, \{\!|0|\!\}_K) \cong (\{\!|0|\!\}_K, \{\!|1|\!\}_K)$
Implementation: requires randomized encryption.

Which-key concealing (k) Given ciphertexts c and c', are they encrypted with the same key?

Formal view: $(\{\!|0|\!\}_K, \{\!|1|\!\}_K) \cong (\{\!|0|\!\}_K, \{\!|1|\!\}_{K'})$

Message-length concealing (l) Does the length of ciphertext reveal the length of plaintext?

Formal view: $\{\!|(((1,1),(1,1)),(1,1))|\!\}_K \cong \{\!|0|\!\}_K$
Implementation: requires maximal message size and padding.

Type-0 security: the encryption scheme does not reveal any of these informations. (There is a type-rkl security for $r, k, l \in Bool$.)

Type-0 Security (2)

Definition (Type-0 security)

An encryption scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is type-0 secure if for every probabilistic polynomial-time adversary A the advantage of A,

$$\begin{aligned} Adv^0_{\Pi[\eta]}(A) \quad = \quad \Big| & Pr[k, k' \stackrel{R}{\leftarrow} \mathcal{K}(\eta) : A^{\mathcal{E}_k(\,\cdot\,), \mathcal{E}_{k'}(\,\cdot\,)}(\eta) = 1] - \\ & Pr[k \stackrel{R}{\leftarrow} \mathcal{K}(\eta) : \quad A^{\mathcal{E}_k(\perp), \mathcal{E}_k(\perp)}(\eta) = 1] \Big| \end{aligned}$$

as a function of η is negligible.

- The adversary $A^{f,g}(\eta)$ has access to two oracles f and g.
- First probability: Choose two keys k and k' by independently running the key generator. Then run the adversary A giving him access to the oracles $f = \mathcal{E}_k(\,\cdot\,)$ and $g = \mathcal{E}_{k'}(\,\cdot\,)$.
- Second probability: Choose a key k by running the key generator. Then run the adversary A with oracles $f = g = \lambda m.\ \mathcal{E}_k(\perp)$. These oracles return samples from $\mathcal{E}_k(\perp)$, ignoring the input m given.

A Computational Semantics of Terms

Definition (Computational semantics of terms)

The distribution $[\![M]\!]_{\Pi[\eta]}$ on strings is induced by the following probabilistic algorithm, which defines how a bit string y is drawn from $[\![M]\!]_{\Pi[\eta]}$:

$$\tau \xleftarrow{R} \mathit{Initialize}(\eta, M);$$
$$y \xleftarrow{R} \mathit{Convert}(\tau, M)$$

where

$$\begin{array}{l} \mathit{Init}(\eta, M): \\ \quad \text{for each } K \in \mathit{Keys}(M) \text{ do } \tau(K) \xleftarrow{R} \mathcal{K}(\eta) \text{ od;} \\ \quad \text{return } \tau \end{array}$$

and

$$\begin{array}{rcl} \mathit{Convert}(\tau, K) & = & (\tau(K), \text{"}\mathit{key}\text{"}) \\ \mathit{Convert}(\tau, b) & = & (b, \text{"}\mathit{bool}\text{"}) \\ \mathit{Convert}(\tau, (M_1, M_2)) & = & (\mathit{Convert}(\tau, M_1), \mathit{Convert}(\tau, M_2), \text{"}\mathit{pair}\text{"}) \\ \mathit{Convert}(\tau, \{\!|M|\!\}_K) & = & (\mathcal{E}_{\tau(K)}(\mathit{Convert}(\tau, M)), \text{"}\mathit{ciphertext}\text{"}) \end{array}$$

Outline

The Soundness Theorem

Theorem (Computational Soundness)

Let M and N be acyclic terms and let Π be a type-0 secure encryption scheme. Then

$$M \cong N \quad \textit{implies} \quad [\![M]\!]_\Pi \approx [\![N]\!]_\Pi.$$

The rest of this module is devoted to the proof of this theorem.

Overview of the Proof

$$
\begin{array}{ccc}
M & \cong & N \\
\| & & \| \\
M_m & & N_n \\
| & & | \\
\vdots & & \vdots \\
| & & | \\
M_0 & = & N_0
\end{array}
\qquad\qquad
\begin{array}{ccc}
[\![M]\!]_\Pi & \not\approx & [\![N]\!]_\Pi \\
\| & & \| \\
[\![M_m]\!]_\Pi & & [\![N_n]\!]_\Pi \\
| & & | \\
\vdots & & \vdots \\
| & & | \\
[\![M_0]\!]_\Pi & = & [\![N_0]\!]_\Pi
\end{array}
$$

- We will construct a chain of hybrid patterns connecting M and N such that $M_0 = \mathit{pattern}(M) = \mathit{pattern}(N) = N_0$.
- Assume $[\![M]\!]_\Pi \not\approx [\![N]\!]_\Pi$.
- Show that there is a "large gap" in the chain: $[\![M_{i-1}]\!]_\Pi \not\approx [\![M_i]\!]_\Pi$ for some i or $[\![N_{j-1}]\!]_\Pi \not\approx [\![N_j]\!]_\Pi$ for some j.
- From the adversary A distinguishing the two ensembles at the gap, construct an adversary A_0 violating the type-0 security of the encryption scheme Π, contradicting the assumption of the Theorem.

Key Renaming (1)

Definition (Recoverable and hidden keys)

Let $Keys(M)$ be the set of keys occurring in M.

$$\begin{aligned} rec(M) &= Keys(M) \cap \mathcal{DY}(\{M\}) \\ hid(M) &= Keys(M) - rec(M) \end{aligned}$$

Lemma (Key renaming)

It is possible to rename the keys in M and N, obtaining M' and N', such that:

- $pattern(M') = pattern(N')$,
- $rec(M') = rec(N') = \{J_1, \ldots, J_\mu\}$,
- $hid(M') = \{K_1, \ldots, K_m\}$,
- $hid(N') = \{K_1, \ldots, K_n\}$, *and*
- *if K_i encrypts K_j in M' or N' then $i \geq j$.*

Key Renaming (2)

Example (Key renaming)

Original terms: $M \cong N$

$$M = (\ \{\!|1,1|\!\}_{K_4},\quad K_6,\quad \{\!|K_5|\!\}_{K_2},\quad \{\!|K_3,K_5|\!\}_{K_4},\quad \{\!|0,K_3|\!\}_{K_5},\quad \{\!|K_1|\!\}_{K_6}\ )$$
$$N = (\ \{\!|K_3|\!\}_{K_5},\quad K_1,\quad \{\!|K_4,1|\!\}_{K_2},\quad \{\!|0|\!\}_{K_4},\quad \{\!|K_4|\!\}_{K_5},\quad \{\!|K_3|\!\}_{K_1}\ )$$

Renamed terms: $M' \equiv N'$

$$M' = (\ \{\!|1,1|\!\}_{K_4},\quad J_1,\quad \{\!|K_2|\!\}_{K_3},\quad \{\!|K_1,K_2|\!\}_{K_4},\quad \{\!|0,K_1|\!\}_{K_2},\quad \{\!|J_2|\!\}_{J_1}\ )$$
$$N' = (\ \{\!|J_2|\!\}_{K_3},\quad J_1,\quad \{\!|K_1,1|\!\}_{K_2},\quad \{\!|0|\!\}_{K_1},\quad \{\!|K_1|\!\}_{K_3},\quad \{\!|J_2|\!\}_{J_1}\ )$$

The resulting terms satisfy the conclusion of Key renaming lemma:

- $pattern(M') = pattern(N') = (\Box, J_1, \Box, \Box, \Box, \{\!|J_2|\!\}_{J_1})$
- $rec(M') = rec(N') = \{J_1, J_2\}$
- $hid(M') = \{K_1, K_2, K_3, K_4\}$
- $hid(N') = \{K_1, K_2, K_3\}$
- $encrypts\_in(M') = \{(K_4, K_2), (K_4, K_1), (K_3, K_2), (K_2, K_1)\}$
- $encrypts\_in(N') = \{(K_3, K_1), (K_2, K_1)\}$

The Hybrid Patterns (1)

We introduce pattern $M_0, M_1, \ldots, M_m$ and $N_0, N_1, \ldots, N_n$ so that these patterns form a chain from M to N.

Definition (Hybrid patterns)

For $i \in \{0, \ldots, m\}$ and $j \in \{0, \ldots, n\}$:

$$\begin{aligned} M_i &= p(M', rec(M') \cup \{K_1, \ldots, K_i\}) \\ N_j &= p(N', rec(N') \cup \{K_1, \ldots, K_j\}) \end{aligned}$$

- $M_m = M'$ and $N_n = N'$
- $M_0 = pattern(M') = pattern(N') = N_0$
- Recall: $M_0 = N_0$ as a result of renaming.
- Intuitively, M_i and N_i are the patterns that the adversary sees if he has a priori knowledge of the otherwise hidden keys $K_1, \ldots, K_i$.
- The ordering of the keys guarantees that this knowledge does not permit the discovery of other hidden keys.

The Hybrid Patterns (2)

Example

$$
\begin{array}{lcccccc}
M' & & & & & & \\
\| & & & & & & \\
M_4 = (& \{\!|1,1|\!\}_{K_4}, & J_1, & \{\!|K_2|\!\}_{K_3}, & \{\!|K_1,K_2|\!\}_{K_4}, & \{\!|0,K_1|\!\}_{K_2}, & \{\!|J_2|\!\}_{J_1} \;) \\
M_3 = (& \square, & J_1, & \{\!|K_2|\!\}_{K_3}, & \square, & \{\!|0,K_1|\!\}_{K_2}, & \{\!|J_2|\!\}_{J_1} \;) \\
M_2 = (& \square, & J_1, & \square, & \square, & \{\!|0,K_1|\!\}_{K_2}, & \{\!|J_2|\!\}_{J_1} \;) \\
M_1 = (& \square, & J_1, & \square, & \square, & \square, & \{\!|J_2|\!\}_{J_1} \;) \\
M_0 = (& \square, & J_1, & \square, & \square, & \square, & \{\!|J_2|\!\}_{J_1} \;) \\
\| & & & & & & \\
N_0 = (& \square, & J_1, & \square, & \square, & \square, & \{\!|J_2|\!\}_{J_1} \;) \\
N_1 = (& \square, & J_1, & \square, & \{\!|0|\!\}_{K_1}, & \square, & \{\!|J_2|\!\}_{J_1} \;) \\
N_2 = (& \square, & J_1, & \{\!|K_1,1|\!\}_{K_2}, & \{\!|0|\!\}_{K_1}, & \square, & \{\!|J_2|\!\}_{J_1} \;) \\
N_3 = (& \{\!|J_2|\!\}_{K_3}, & J_1, & \{\!|K_1,1|\!\}_{K_2}, & \{\!|0|\!\}_{K_1}, & \{\!|K_1|\!\}_{K_3}, & \{\!|J_2|\!\}_{J_1} \;) \\
\| & & & & & & \\
N' & & & & & &
\end{array}
$$

Defining Ensembles for the Hybrid Patterns

Definition (Computational semantics of patterns)

The distribution $[\![M]\!]_{\Pi[\eta]}$ on strings is induced by the following probabilistic algorithm, which defines how a bit string y is drawn from $[\![M]\!]_{\Pi[\eta]}$:

$$\tau \xleftarrow{R} \mathit{Init}(\eta, M);$$
$$y \xleftarrow{R} \mathit{Convert}(\tau, M)$$

where for some new, fixed key K_0 not occurring elsewhere:

$$\begin{aligned} &\mathit{Init}(\tau, M): \\ &\quad \text{for each } K \in \mathit{Keys}(M) \cup \{K_0\} \text{ do } \tau(K) \xleftarrow{R} \mathcal{K}(\eta) \text{ od;} \\ &\quad \text{return } \tau \end{aligned}$$

and

$$\begin{aligned} \mathit{Convert}(\tau, K) &= (\tau(K), \text{"}\mathit{key}\text{"}) \\ \mathit{Convert}(\tau, i) &= (i, \text{"}\mathit{bool}\text{"}) \\ \mathit{Convert}(\tau, (M_1, M_2)) &= (\mathit{Convert}(\tau, M_1), \mathit{Convert}(\tau, M_2), \text{"}\mathit{pair}\text{"}) \\ \mathit{Convert}(\tau, \{\!|M|\!\}_K) &= (\mathcal{E}_{\tau(K)}(\mathit{Convert}(\tau, M)), \text{"}\mathit{ciphertext}\text{"}) \\ \mathit{Convert}(\tau, \Box) &= (\mathcal{E}_{\tau(K_0)}(\perp), \text{"}\mathit{ciphertext}\text{"}) \end{aligned}$$

Finding a Large Gap (1)

- Recall: We assume $[\![M]\!]_\Pi \not\approx [\![N]\!]_\Pi$ to derive a contradiction with type-0 security of Π.
- Clearly $[\![M]\!]_\Pi = [\![M']\!]_\Pi$, since M and M' differ only by key renaming. Similarly, we have $[\![N]\!]_\Pi = [\![N']\!]_\Pi$.
- Therefore, we have $[\![M']\!]_\Pi \not\approx [\![N']\!]_\Pi$, i.e., there is a probabilistic polynomial-time adversary A such that

$$\begin{aligned}\lambda(\eta) \quad = \quad & \Big| Pr[x \xleftarrow{R} [\![M']\!]_{\Pi[\eta]} : A(\eta, x) = 1] - \\ & \quad Pr[x \xleftarrow{R} [\![N']\!]_{\Pi[\eta]} : A(\eta, x) = 1] \Big|\end{aligned}$$

  is not negligible.
- This means that there is a constant $c \in \mathbb{N}$ and an infinite set $\mathcal{F} \subseteq \mathbb{N}$ such that $\lambda(\eta) > \eta^{-c}$ for all $\eta \in \mathcal{F}$.

Finding a Large Gap (2)

Definition (Functions p_i and q_j)

$$\begin{aligned} p_i(\eta) &= Pr[x \xleftarrow{R} [\![M_i]\!]_{\Pi[\eta]} : A(\eta, x) = 1] \qquad (\text{for } 0 \leq i \leq n) \\ q_j(\eta) &= Pr[x \xleftarrow{R} [\![N_j]\!]_{\Pi[\eta]} : A(\eta, x) = 1] \qquad (\text{for } 0 \leq j \leq m) \end{aligned}$$

- Hence, since $M' = M_m$ and $N' = N_n$, we have

$$\lambda(\eta) = |\, p_m(\eta) - q_n(\eta) \,|\,.$$

- We also have $p_0 = q_0$, since M' and N' yield the same pattern, hence

$$\begin{aligned} \lambda &= |\,(p_m - p_{m-1}) + (p_{m-1} - p_{m-2}) + \ldots + (p_1 - p_0) \\ &\quad + (q_0 - q_1) + (q_1 - q_2) + \ldots + (q_{n-1} - q_n)\,| \end{aligned}$$

- By the triangle inequality for $m + n$ summands, for all $\eta \in \mathcal{F}$:

$$\begin{aligned} & \exists i \in \{1, \ldots, n\}.\ |\, p_i(\eta) - p_{i-1}(\eta) \,| \geq \lambda(\eta)/(m+n) \\ \vee\ & \exists j \in \{1, \ldots, m\}.\ |\, q_j(\eta) - q_{j-1}(\eta) \,| \geq \lambda(\eta)/(m+n) \end{aligned}$$

Finding the Large Gap (3)

- By the triangle inequality for $m+n$ summands, for all $\eta \in \mathcal{F}$:

$$\begin{array}{rl} & \exists i \in \{1,\ldots,n\}.\ \ |\,p_i(\eta) - p_{i-1}(\eta)\,| \ \geq\ \lambda(\eta)/(m+n) \\ \vee & \exists j \in \{1,\ldots,m\}.\ |\,q_j(\eta) - q_{j-1}(\eta)\,| \ \geq\ \lambda(\eta)/(m+n) \end{array}$$

- Hence, since there is a finite number $(m+n)$ possible indices, there is an index $i \in \{1,\ldots,n\}$ such that for infinitely many $\eta \in \mathcal{F}$:

$$|\,p_i(\eta) - p_{i-1}(\eta)\,| \ \geq\ \lambda(\eta)/(m+n)$$

  or there is an index $j \in \{1,\ldots,m\}$ such that for infinitely many $\eta \in \mathcal{F}$:

$$|\,q_j(\eta) - q_{j-1}(\eta)\,| \ \geq\ \lambda(\eta)/(m+n)$$

- Let i be such an index (the case for j is analogous). This means that

$$\varepsilon(\eta) = |\,p_i(\eta) - p_{i-1}(\eta)\,|$$

  is non-negligible (since λ is non-negligible and $m+n$ is a constant).

Contradicting the Security of Π (1)

Next, we construct an adversary A_0 from the adversary A such that the following lemma holds (proved shortly).

Lemma (Relating the adversaries)

$$
\begin{aligned}
p_i(\eta) &= Pr[y \xleftarrow{R} [\![M_i]\!]_{\Pi[\eta]} : A(\eta, y) = 1] && (1)\\
&= Pr[k_0, k_i \xleftarrow{R} \mathcal{K}(\eta) : A_0^{\mathcal{E}_{k_i}(\cdot), \mathcal{E}_{k_0}(\cdot)}(\eta) = 1] && (2)\\
p_{i-1}(\eta) &= Pr[y \xleftarrow{R} [\![M_{i-1}]\!]_{\Pi[\eta]} : A(\eta, y) = 1] && (3)\\
&= Pr[k_0 \xleftarrow{R} \mathcal{K}(\eta) : A_0^{\mathcal{E}_{k_0}(\perp), \mathcal{E}_{k_0}(\perp)}(\eta) = 1] && (4)
\end{aligned}
$$

From the Lemma it follows that

$$|(4) - (2)| = \boxed{Adv^0_{\Pi[\eta]}(A_0) \;=\; \varepsilon(\eta)} = |(3) - (1)|$$

which is non-negligible.

Hence, we have derived a contradiction with the type-0 security of Π.

Contradicting the Security of Π (2)

Definition (Adversary $A_0^{f,g}(\eta)$; part 1)

The adversary $A_0^{f,g}(\eta)$ is defined by the following probabilistic algorithm:

$$\begin{aligned} &\tau \xleftarrow{R} Init(\eta, M'); \\ &y \xleftarrow{R} Conv^{f,g}(\tau, M'); \\ &b \xleftarrow{R} A(\eta, y); \textbf{return } b \end{aligned}$$

where $Conv^{f,g}(\tau, M')$ is defined on the next slide.

The Lemma requires that

$$\begin{aligned} Pr[y \xleftarrow{R} [\![M_i]\!]_{\Pi[\eta]} : \quad A(\eta, y) = 1] &= Pr[k_0, k_i \xleftarrow{R} \mathcal{K}(\eta) : A_0^{\mathcal{E}_{k_i}(\cdot), \mathcal{E}_{k_0}(\cdot)}(\eta) = 1] \\ Pr[y \xleftarrow{R} [\![M_{i-1}]\!]_{\Pi[\eta]} : A(\eta, y) = 1] &= Pr[k_0 \xleftarrow{R} \mathcal{K}(\eta) : \quad A_0^{\mathcal{E}_{k_0}(\perp), \mathcal{E}_{k_0}(\perp)}(\eta) = 1] \end{aligned}$$

We derive the following requirements on $Conv^{f,g}(\tau, M')$:

$$\begin{aligned} Convert(\tau, M_i) &= Conv^{\mathcal{E}_{k_i}(\cdot), \mathcal{E}_{k_0}(\cdot)}(\tau, M') \\ Convert(\tau, M_{i-1}) &= Conv^{\mathcal{E}_{k_0}(\perp), \mathcal{E}_{k_0}(\perp)}(\tau, M') \end{aligned}$$

Contradicting the Security of Π (3)

Definition (Adversary $A_0^{f,g}(\eta)$; part 2)

$$
\begin{array}{lcl}
Conv^{f,g}(\tau, K) & = & (\tau(K), \text{"}key\text{"}) \\
Conv^{f,g}(\tau, i) & = & (i, \text{"}bool\text{"}) \\
Conv^{f,g}(\tau, (M_1, M_2)) & = & (Conv^{f,g}(\tau, M_1), Conv^{f,g}(\tau, M_2), \text{"}pair\text{"}) \\
Conv^{f,g}(\tau, \{\!|M_1|\!\}_K) & = & \textbf{if } K \in rec(M') \cup \{K_1, \ldots, K_{i-1}\} \textbf{ then} \\
 & & \quad (\mathcal{E}_{\tau(K)}(Conv^{f,g}(\tau, M_1)), \text{"}ciphertext\text{"}) \\
 & & \textbf{else if } K = K_i \textbf{ then} \\
 & & \quad (f(Conv^{f,g}(\tau, M_1)), \text{"}ciphertext\text{"}) \\
 & & \textbf{else} \quad (^* \ K \in \{K_{i+1}, \ldots, K_n\} \ ^*) \\
 & & \quad (g(\perp), \text{"}ciphertext\text{"})
\end{array}
$$

Requirements on $Conv^{f,g}(\tau, M')$:

$$
\begin{array}{lcll}
Convert(\tau, M_i) & = & Conv^{\mathcal{E}_{k_i}(\cdot), \mathcal{E}_{k_0}(\cdot)}(\tau, M') & \checkmark \\
Convert(\tau, M_{i-1}) & = & Conv^{\mathcal{E}_{k_0}(\perp), \mathcal{E}_{k_0}(\perp)}(\tau, M') & \checkmark
\end{array}
$$

Conclusion (Summary)

First step towards bridging the gap between the two views:

- Suitable design of the formal and computational models, so that a relation is possible.
- Soundness result: if messages in the formal view are indistinguishable, then so are their representations in the computational model.

$\Rightarrow$ Verification in the formal model carries over to verification in the computational model.

$\Rightarrow$ In this setting, there can't be any "attacks" in the cryptographic model, that would not be possible in the formal model.

Conclusion (Extensions)

Still lots of challenging problems to solve:

- Active adversary, in particular interaction with honest agents.
  - ★ “trace-mapping” approach [Micciancio, Warinschi & Cortier]
  - ★ based on universal composability [Canetti & Herzog]
  - ★ cryptographic library based on reactive simulatability framework [Backes, Pfitzmann & Waidner]
- More cryptographic primitives like Diffie-Hellman.
- Cryptographic systems for which tagging/typing is not possible.
- Weaker restrictions on the encryption scheme.

Bibliography

- Martin Abadí and Philip Rogaway: *Reconciling Two Views of Cryptography (The Computational Soundness of Formal Encryption)*, Journal of Cryptology 15(2), pp 103-127, 2002.
- Daniele Micciancio, Bogdan Warinschi: *Soundness of Formal Encryption in the Presence of Active Adversaries*. TCC 2004: 133-151.
- Véronique Cortier, Bogdan Warinschi: *Computationally Sound, Automated Proofs for Security Protocols*. ESOP 2005: 157-171.
- Ran Canetti and Jonathan Herzog. *Universally Composable Symbolic Analysis of Mutual Authentication and Key-Exchange Protocols*. TCC 2006: 380-403. Also: Cryptology ePrint Archive, Report 2004/334, http://eprint.iacr.org/2004/334.
- Michael Backes, Birgit Pfitzmann, Michael Waidner: *A composable cryptographic library with nested operations*. CCS 2003: 220-230. Also: Cryptology ePrint Archive, Report 2003/015, http://eprint.iacr.org/2003/015.