Adaptive Soundness of Static Equivalence

Steve Kremer
Joint work with Laurent Mazaré

LSV, ENS Cachan & CNRS & INRIA Futurs

January 16, 2007

Context and outline

- In [AbadiRogaway00] soundness result for a symbolic pattern based equivalence: passive adversary, symmetric encryption and pairs
- In [MicciancioPanjwani05]: similar model, but adaptive adversary
- In [BaudetCortierKremer05], [AbadiBaudetWarinschi06]: soundness of static equivalence, framework for arbitrary equational theories, passive adversary

In this talk: static equivalence + adaptive adversary

- adaptive soundness of static equivalence
- general results and "combination" proof technique
- adaptive soundness results for equational theories: xor, modular exponentiation, symmetric encryption, symmetric encryption + xor, symmetric encryption + modular exponentiation
- an application of adaptive soundness: dynamic group key exchange protocols

Model

- arbitrary signature + sorts

$$
\begin{array}{lll}
T & ::= & \text{term of sort } s \\
& \mid \quad x & \text{variable } x \text{ of sort s} \\
& \mid \quad a & \text{name } a \text{ of sort s} \\
& \mid \quad f(T_1, \ldots, T_k) & \text{application of symbol } f \in \mathcal{F}
\end{array}
$$

  + concrete implementation $[\![T]\!]$
  (bitstrings, poly-time algorithms, randomly chosen names)

- arbitrary equational theories

Example

$$E = \{\mathrm{dec}(\mathrm{enc}(x, y), y) = x, x \oplus y = y \oplus x, x \oplus 0 = x, x \oplus x = 0\}$$

Frames and static equivalence

Terms are organized into frames:

$$\varphi_1 = \{x_1 \mapsto \mathrm{enc}(k_1, k_2),\ x_2 \mapsto \mathrm{enc}(k_4, k_3),\ x_3 \mapsto k_3\}$$

Definition (static equivalence)

$\varphi_1 \approx_E \varphi_2$ if

- $\mathrm{dom}(\varphi_1) = \mathrm{dom}(\varphi_2)$
- for all M, N which do not use names appearing in the frames we have: $M\varphi_1 =_E N\varphi_1$ iff $M\varphi_2 =_E N\varphi_2$.

Example

$$\phi_1 = \{x \mapsto \mathrm{enc}(0, k)\} \qquad \phi_2 = \{x \mapsto \mathrm{enc}(1, k)\}$$
$$\phi_3 = \{x \mapsto \mathrm{enc}(0, k),\ y \mapsto k\} \qquad \phi_4 = \{x \mapsto \mathrm{enc}(0, k'),\ y \mapsto k\}$$

Frames and static equivalence

Terms are organized into frames:

$$\varphi_1 = \{x_1 \mapsto \mathrm{enc}(k_1, k_2),\ x_2 \mapsto \mathrm{enc}(k_4, k_3),\ x_3 \mapsto k_3\}$$

Definition (static equivalence)

$\varphi_1 \approx_E \varphi_2$ if

- $\mathrm{dom}(\varphi_1) = \mathrm{dom}(\varphi_2)$
- for all M, N which do not use names appearing in the frames we have: $M\varphi_1 =_E N\varphi_1$ iff $M\varphi_2 =_E N\varphi_2$.

Example

$$\phi_1 = \{x \mapsto \mathrm{enc}(0, k)\} \quad \approx_E \quad \phi_2 = \{x \mapsto \mathrm{enc}(1, k)\}$$
$$\phi_3 = \{x \mapsto \mathrm{enc}(0, k),\ y \mapsto k\} \quad \not\approx_E \quad \phi_4 = \{x \mapsto \mathrm{enc}(0, k'),\ y \mapsto k\}$$

Computational indistinguishability

The concrete interpretation of terms, or frames yields a (family of) distribution(s) $[\![\phi]\!]_{A_\eta}$

Definition (indistinguishability)

$[\![\varphi_1]\!] \approx [\![\varphi_2]\!]$ if $\quad \mathrm{Adv}^{\mathsf{IND}}(\mathcal{A}, [\![\varphi_1]\!], [\![\varphi_2]\!])(\eta) =$

$$|\mathbb{P}\left[\phi_1 \leftarrow [\![\varphi_1]\!]; \mathcal{A}(\eta, \phi_1) = 1\right] - \mathbb{P}\left[\phi_2 \leftarrow [\![\varphi_2]\!]; A(\eta, \phi_2) = 1\right]|$$

is a negligible function of η.

(Adaptive) Soundness of Static Equivalence

Soundness of static equivalence

An implementation A_η is $\approx_E$-sound iff $\varphi_1 \approx_E \varphi_2$ implies $[\![\varphi_1]\!] \approx [\![\varphi_2]\!]$

Adaptive Soundness of static equivalence

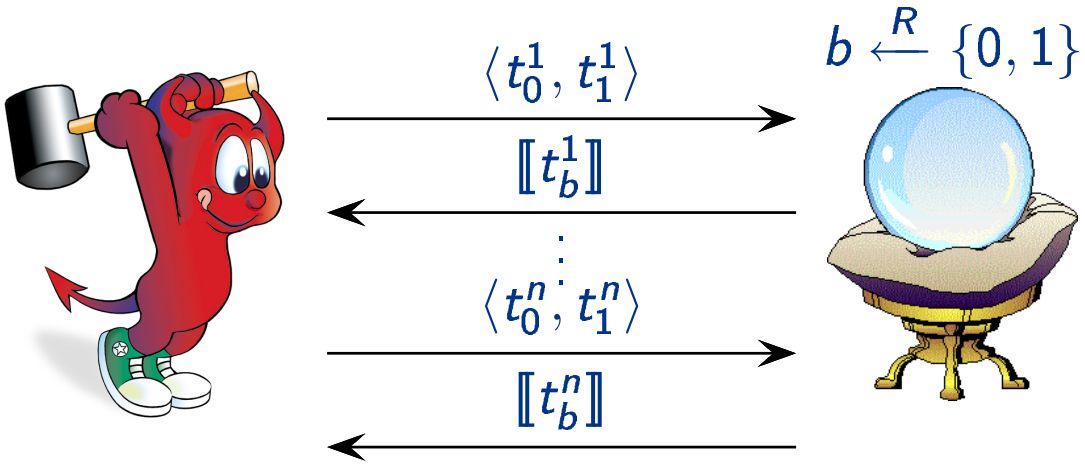

An implementation A_η is $\approx_E$-ad-sound if for any sequence $\langle t_0^i, t_1^i \rangle$ we have that $\{x_i \mapsto t_0^i\} \approx_E \{x_i \mapsto t_1^i\}$ implies

$\mathrm{Adv}^{\mathrm{ADPT}}_{\mathcal{A},A_\eta}(\eta) = \left| \mathbb{P}\left[\mathcal{A}^{\mathcal{O}^1_{LR,A_\eta}} = 1\right] - \mathbb{P}\left[\mathcal{A}^{\mathcal{O}^0_{LR,A_\eta}} = 1\right] \right|$ is negligeable in η

(Adaptive) Soundness of Static Equivalence

Soundness of static equivalence

An implementation A_η is unconditionally $\approx_E$-sound iff $\varphi_1 \approx_E \varphi_2$ implies $[\![\varphi_1]\!] = [\![\varphi_2]\!]$

Adaptive Soundness of static equivalence

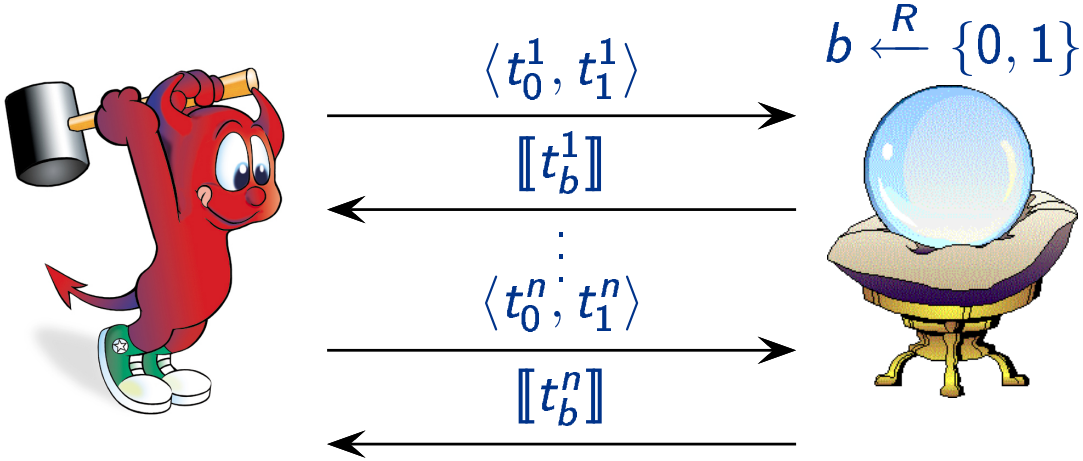

An implementation A_η is unconditionally $\approx_E$-ad-sound if for any sequence $\langle t_0^i, t_1^i \rangle$ we have that $\{x_i \mapsto t_0^i\} \approx_E \{x_i \mapsto t_1^i\}$ implies $\mathrm{Adv}^{\mathrm{ADPT}}_{\mathcal{A},A_\eta}(\eta) = \left| \mathbb{P}\left[\mathcal{A}^{\mathcal{O}^1_{LR,A_\eta}} = 1\right] - \mathbb{P}\left[\mathcal{A}^{\mathcal{O}^0_{LR,A_\eta}} = 1\right] \right|$ is 0

Adaptive soundness is strictly stronger!

Proposition

If A_η is $\approx_E$-ad-sound then A_η is also $\approx_E$-sound but the converse is false in general.

Proof idea: Consider the following signature without any equations

$$\begin{array}{rcl} 0, 1 & : & \mathrm{Bit} \\ \mathrm{cons} & : & \mathrm{Bit} \times \mathrm{BS} \to \mathrm{BS} \\ \mathrm{eq} & : & \mathrm{BS} \times \mathrm{Nonce} \to \mathrm{Bool} \end{array}$$

$[\![eq]\!](bs, N)$ outputs 1 if $bs = N$, 0 otherwise. We do have $\approx$-soundness, but not $\approx$-ad-soundness

Proposition

A_η is unconditionnaly $\approx_E$-ad-sound iff A_η is unconditionnaly $\approx_E$-sound.

Combining signatures

Definition (Disjoint signatures)

Let $\Sigma_1 = (\mathcal{S}_1, \mathcal{F}_1)$ and $\Sigma_2 = (\mathcal{S}_2, \mathcal{F}_2)$ are disjoint iff $\mathcal{F}_1 \cap \mathcal{F}_2 = \emptyset$ and $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$.

Definition (Signature combination, Layered signatures)

Let $\Sigma_1 = (\mathcal{S}_1, \mathcal{F}_1)$ and $\Sigma_2 = (\mathcal{S}_2, \mathcal{F}_2)$ be two disjoint signatures. We say that a subsort relation S is a signature combination for Σ_1 and Σ_2 if $\mathsf{S} \subseteq \mathcal{S}_2 \times \mathcal{S}_1$.
$\Sigma = (\mathcal{S}_1 \cup \mathcal{S}_2, \mathcal{F}_1 \cup \mathcal{F}_2)$ is a $(\Sigma_1, \Sigma_2)_{\mathsf{S}}$-layered signature.

Example (Encryption and pseudo-random generators)

Let $\Sigma_1 = (\{\mathrm{Data}\}, \{\mathrm{enc}, \mathrm{dec}\})$ and $\Sigma_2 = (\{\mathrm{Rand}\}, \{\mathrm{prg}\})$ and $\mathrm{Rand}\,\mathsf{S}\,\mathrm{Data}$. Then $\Sigma_1 \cup \Sigma_2$ is $(\Sigma_1, \Sigma_2)_{\mathsf{S}}$-layered.

Hybrid frames and functions

Let Σ be a $(\Sigma_1, \Sigma_2)_S$-layered signature with eq. theories E_1 and E_2.

Definition ($st_2(t)$)

For t in Σ, $st_2(t) = \{t|_p \mid sort(t|_p) \in \mathcal{S}_2, p = p' \cdot i \Rightarrow sort(t|_{p'}) \notin \mathcal{S}_2\}$

Example (Encryption and pseudo-random generators)

Let $t = \mathrm{enc}(\mathrm{prg}(r), \mathrm{prg}(\mathrm{prg}(r')))$. $st_2(t)$ contains $\mathrm{prg}(r)$ and $\mathrm{prg}(\mathrm{prg}(r'))$.

Definition (Hybrid functions, Hybrid Frames)

A (E_1, E_2)-hybrid function is a function σ which given terms on Σ, $(t_i)_{1 \leq i \leq n}$ outputs a function from $st_2((t_i)_{1 \leq i \leq n})$ to terms on Σ_2 such that:

- if $st_2(t_i) = \{(s_i^j)_j\}$ then $\{x_{i,j} \mapsto s_i^j\}_{i,j} \approx_{E_2} \{x_{i,j} \mapsto s_i^j \sigma(t_1, ..., t_i)\}_{i,j}$
- let $(u_i)_{1 \leq i \leq n}$ be a sequence on Σ, $\{x_i \mapsto t_i\}_i \approx_{E_1 \cup E_2} \{x_i \mapsto u_i\}_i \Rightarrow \{x_i \mapsto \sigma(t_1, ..., t_i)(t_i)\}_i \approx_{E_1} \{x_i \mapsto \sigma(u_1, ..., u_i)(u_i)\}_i$

If $\phi = \{x_i \mapsto t_i\}$, the hybrid frame $\phi\sigma = \{x_i \mapsto \sigma(t_1, ..., t_i)(t_i)\}$.

Combining soundness results

Proposition 3

We consider the families of computational algebras (A^1_η) for Σ_1 and (A^2_η) for Σ_2 respecting S, *i.e.* $(s_2, s_1) \in S$ implies that $[\![s_2]\!]_{A^2_\eta} \subseteq [\![s_1]\!]_{A^1_\eta}$.

Let F be a set of frames over $\Sigma_1 \cup \Sigma_2$ and σ be a (E_1, E_2)-hybrid function. If $A^1_\eta \times A^2_\eta$ is $\approx_{E_1}$-ad-sound for $F\sigma$ and A^2_η is $\approx_{E_2}$-ad-sound for frames on Σ_2, then $A^1_\eta \times A^2_\eta$ is $\approx_{E_1 \cup E_2}$-ad-sound for frames in F.

We will use this result for combining symmetric encryption with xor, respectively modular exponentiation

Exclusive Or

Consider the following signature modelling XOR

$$\begin{array}{rcl} \oplus & : & \mathrm{Data}_{\oplus} \times \mathrm{Data}_{\oplus} \to \mathrm{Data}_{\oplus} \\ 0_{\oplus}, 1_{\oplus} & : & \mathrm{Data}_{\oplus} \end{array}$$

and the equational theory

$$\begin{array}{rcl} (x \oplus y) \oplus z & = & x \oplus (y \oplus z) \\ x \oplus y & = & y \oplus x \end{array} \qquad \begin{array}{rcl} x \oplus x & = & 0_{\oplus} \\ x \oplus 0_{\oplus} & = & x \end{array}$$

Proposition

The usual implementation for the XOR theory is unconditionally $\approx_{E_{\oplus}}$-ad-sound.

This follows directly from unconditional $\approx_{E_{\oplus}}$-soundness shown in [BCK05].

Modular exponentiation

Consider the following signature modelling modular exponentiation

$$\begin{array}{rcll} \exp & : & R \to G & \text{exponentiation} \\ * & : & G \times G \to G & \text{mult in } \mathbb{G} \end{array} \qquad \begin{array}{rcll} +, \cdot & : & R \times R \to R & \text{add, mult} \\ - & : & R \to R & \text{inverse} \\ 0_R, 1_R & : & R & \text{constants} \end{array}$$

and the equational theory

$$\begin{array}{ll} x + y = y + x & x \cdot y = y \cdot x \\ (x + y) + z = x + (y + z) & x \cdot (y + z) = x \cdot y + x \cdot z \\ (x \cdot y) \cdot z = x \cdot (y \cdot z) & x + (-x) = 0_R \\ 0_R + x = x & 1_R \cdot x = x \end{array}$$

$$\exp(x) * \exp(y) = \exp(x + y)$$

Some restrictions on frames:

- only elements of sort $\mathbb{G}$
- products have to be power-free, *i.e* x^n is forbidden for $n > 1$
- products must not contain more than ℓ elements for a fixed bound ℓ

Soundness of modular exponentiation

Concrete implementation

- Instance Generator IG: a polynomial-time algorithm that outputs a cyclic group $\mathbb{G}$ of prime order q
- Operations are the usual implementations over $\mathbb{Z}_q$

An instance generator satisfies the DDH assumption if
$\left|\mathbb{P}\left[(g,q) \leftarrow IG(\eta) \ : \ a,b \leftarrow \mathbb{Z}_q \ : \ \mathcal{A}(g^a, g^b, g^{ab}) = 1\right] - \right.$
$\left.\mathbb{P}\left[(g,q) \leftarrow IG(\eta) \ : \ a,b,c \leftarrow \mathbb{Z}_q \ : \ \mathcal{A}(g^a, g^b, g^c) = 1\right]\right|$ is negligeable.

Soundness results

$$\begin{array}{ccc} \approx_{E_{\mathsf{DH}}}\text{-ad-sound} & \Longleftrightarrow & \mathsf{DDH} \\ & & \Updownarrow \\ \approx_{E_{\mathsf{DH}}}\text{-sound} & \Longleftrightarrow & \mathsf{DDH} \end{array}$$

Symmetric encryption

Consider the following signature modelling symmetric encryption

$$\begin{array}{rcl@{\qquad}rcl} \mathrm{enc}, \mathrm{dec} & : & \mathrm{Data} \times \mathrm{Data} \to \mathrm{Data} & \mathrm{samekey} & : & \mathrm{Data} \times \mathrm{Data} \to \mathrm{Data} \\ \mathrm{pair} & : & \mathrm{Data} \times \mathrm{Data} \to \mathrm{Data} & \mathrm{tenc}, \mathrm{tpair} & : & \mathrm{Data} \to \mathrm{Data} \\ \pi_l, \pi_r & : & \mathrm{Data} \to \mathrm{Data} & 0, 1 & : & \mathrm{Data} \end{array}$$

and the equational theory

$$\begin{array}{rcl@{\qquad}rcl} \mathrm{dec}(\mathrm{enc}(x,y),y) & = & x & \mathrm{samekey}(\mathrm{enc}(x,y),\mathrm{enc}(z,y)) & = & 1 \\ \pi_l(\mathrm{pair}(x,y)) & = & x & \mathrm{tenc}(\mathrm{enc}(x,y)) & = & 1 \\ \pi_r(\mathrm{pair}(x,y)) & = & y & \mathrm{tpair}(\mathrm{pair}(x,y)) & = & 1 \end{array}$$

Restrictions on frames

- no key cycles
- no destructor symbols
- if k is used as plaintext in t_i, then k cannot be used at a key position in t_j for $j < i$ (avoids selective decommitment)

Soundness of symmetric encryption

Concrete implementation

We suppose length-conceiling semantic secure encryption (IND-CPA):

$$\mathrm{Adv}^{cpa}_{\mathcal{SE},\mathcal{A}}(\eta) = \left|\mathbb{P}\left[\mathcal{A}^{LR^1_{\mathcal{SE}}}(\eta) = 1\right] - \mathbb{P}\left[\mathcal{A}^{LR^0_{\mathcal{SE}}}(\eta) = 1\right]\right|$$

where $LR^b_{\mathcal{SE}}$ is a left-right encryption oracle.

We also consider a variant IND-CPA': the oracle only accepts a single call. Given a list of of pairs of bitstrings $(\langle bs^i_0, bs^i_1\rangle)_i$ the oracle returns the list $(\mathcal{E}(bs^i_b, k))_i$.

Soundness results

$$\begin{array}{ccc} \approx_{E_{\mathrm{sym}}}\text{-ad-sound} & \begin{array}{c}\Leftarrow\\ \not\Rightarrow\end{array} & \text{IND-CPA} \\ \not\Uparrow \quad \Downarrow & & \not\Uparrow \quad \Downarrow \\ \approx_{E_{\mathrm{sym}}}\text{-sound} & \begin{array}{c}\Leftarrow\\ \not\Rightarrow\end{array} & \text{IND-CPA}' \end{array}$$

Combinations

We show a stronger soundness result for symmetric encryption: given a $(\Sigma_{\text{sym}}, \Sigma_2)$-layered signature, IND-CPA guarantees soundness $\approx_{E_{\text{sym}}}$-ad-soundess.
Using our combination technique we obtain

- a soundness result for symmetric encryption + modular exponentation
- a soundness result for symmetric encryption + xor

Applying the combination result

- The difficult part is to exhibit a hybrid function σ
- When considering modular exponentiation, the IND-CPA scheme uses group elements as keys, together with a randomness extractor

An application: dynamic group key exchange protocols

A protocol is specified by four functions $(\mathcal{S}, \mathcal{J}, \mathcal{L}, \mathcal{K})$:

- $\mathcal{S}(U_1, \ldots U_n)$: setup for users $(U_1, \ldots, U_n)$ returns terms modelling the execution of the setup protocol
- $\mathcal{J}(U), \mathcal{L}(U)$: join, leave of user U returns the terms modelling the execution of the join/leave sub-protocol
- $\mathcal{K}$: the key function returns the current group key

We consider static corruption (once at the beginning of the protocol).

In the symbolic model, the above functions yield a transition system associating the frame of exchanged messages to each state.

A DKE protocol is symbolically secure if for any reachable state s

$$\phi(s) \cup \{x \mapsto \mathcal{K}(s)\} \approx_E \phi(s) \cup \{x \mapsto r\}$$

Soundness result for DKE

In the concrete model, we suppose an oracle modeling corruption, $\mathcal{S}$, $\mathcal{T}$ and $\mathcal{L}$.
Moreover, there is a final *Test* call which returns either the key or a random key (depending on the challenge bit).

A DKE protocol is secure in the concrete model if the advantage

$$\mathrm{Adv}_{\mathcal{A},A_\eta}^{(\mathcal{S},\mathcal{T},\mathcal{L},\mathcal{K})}(\eta) = \mathbb{P}\left[\mathcal{A}^{\mathcal{O}_1} = 1\right] - \mathbb{P}\left[\mathcal{A}^{\mathcal{O}_0} = 1\right]$$

is negligible for any adversary.

Soundness result for DKE and modular exponentiation

Let (A_η) be a family of computational algebras and $\Pi = (\mathcal{S}, \mathcal{T}, \mathcal{L}, \mathcal{K})$ be a DKE. If (A_η) is $\approx_{E_{\mathsf{DH}}}$-ad-sound and Π is secure in the symbolic model, then Π is secure in the concrete model.

Conclusion

- Definition of a framework for reasoning about soundness of static equivalence in the presence of an adaptive aadversary
- Proof technique for combining soundness results
- Illustration on several equational theories and combinations of them
- Application to the case of DKE protocols
  Decision procedures for the symbolic verifiaction of DKE protocols?