

CryptoVerif: A Computationally Sound Mechanized Prover for Cryptographic Protocols

Bruno Blanchet

CNRS, École Normale Supérieure, INRIA, Paris

April 2009

Introduction

Two approaches for the automatic proof of cryptographic protocols in a computational model:

- **Indirect approach:**
 - 1) Make a Dolev-Yao proof.
 - 2) Use a theorem that shows the soundness of the Dolev-Yao approach with respect to the computational model.

Pioneered by Abadi and Rogaway; pursued by many others.
- **Direct approach:**

Design automatic tools for proving protocols in a computational model.

Approach pioneered by Laud.

Introduction

Two models for security protocols:

- **Computational model:**
 - messages are bitstrings
 - cryptographic primitives are functions from bitstrings to bitstrings
 - the adversary is a probabilistic polynomial-time Turing machine

Proofs are done manually.

- **Formal model** (so-called “Dolev-Yao model”):
 - cryptographic primitives are ideal blackboxes
 - messages are terms built from the cryptographic primitives
 - the adversary is restricted to use only the primitives

Proofs can be done automatically.

Our goal: achieve **automatic provability** under the realistic **computational** assumptions.

Advantages and drawbacks

The indirect approach allows more reuse of previous work, but it has limitations:

- **Hypotheses** have to be added to make sure that the computational and Dolev-Yao models coincide.
- The **allowed cryptographic primitives** are often limited, and only ideal, not very practical primitives can be used.
- Using the Dolev-Yao model is actually a (big) **detour**; The computational definitions of primitives fit the computational security properties to prove. They do not fit the Dolev-Yao model.

We decided to focus on the direct approach.

An automatic prover

We have implemented an **automatic prover**:

- proves **secrecy** and **correspondence** properties.
- provides a **generic** method for specifying properties of **cryptographic primitives** which handles MACs (message authentication codes), symmetric encryption, public-key encryption, signatures, hash functions, ...
- works for **N sessions** (polynomial in the security parameter), with an **active adversary**.
- gives a bound on the **probability** of an attack (exact security).

Process calculus for games

Games are formalized in a **process calculus**:

- It is adapted from the pi calculus.
- The semantics is **purely probabilistic** (no non-determinism).
- All processes run in **polynomial time**:
 - polynomial number of copies of processes,
 - length of messages on channels bounded by polynomials.

This calculus is inspired by:

- the calculus of [Lincoln, Mitchell, Mitchell, Scedrov, 1998],
- the calculus of [Laud, 2005].

Produced proofs

As in Shoup's and Bellare&Rogaway's method, the proof is a **sequence of games**:

- The first game is the **real protocol**.
- One goes from one game to the next by syntactic transformations or by applying the definition of security of a cryptographic primitive. Between consecutive games, the difference of probability of success of an attack is negligible.
- The last game is **"ideal"**: the security property is obvious from the form of the game.
(The advantage of the adversary is typically 0 for this game.)

Example

$$A \rightarrow B : e = \{x'_k\}_{x_k}, \text{mac}(e, x_{mk}) \quad x'_k \text{ fresh}$$

A sends to B a fresh key x'_k encrypted under authenticated encryption, implemented as encrypt-then-MAC.

x'_k should remain secret.

Example (initialization)

$$A \rightarrow B : e = \{x'_k\}_{x_k}, \text{mac}(e, x_{mk}) \quad x'_k \text{ fresh}$$

$Q_0 = \text{start}(); \text{new } x_r : \text{keyseed}; \text{let } x_k : \text{key} = \text{kgen}(x_r) \text{ in}$
 $\text{new } x'_r : \text{mkeyseed}; \text{let } x_{mk} : \text{mkey} = \text{mkgen}(x'_r) \text{ in } \overline{c}\langle \rangle; (Q_A \mid Q_B)$

Initialization of keys:

- 1 The process Q_0 waits for a message on channel start to start running. The adversary triggers this process.
- 2 Q_0 **generates encryption and MAC keys**, x_k and x_{mk} respectively, using the key generation algorithms kgen and mkgen .
- 3 Q_0 returns control to the adversary by the output $\overline{c}\langle \rangle$.
 Q_A and Q_B represent the actions of A and B (see next slides).

Example (role of B)

$$A \rightarrow B : e = \{x'_k\}_{x_k}, \text{mac}(e, x_{mk}) \quad x'_k \text{ fresh}$$

$Q_B = !^{i' \leq n} c_B(x'_m : \text{bitstring}, x_{ma} : \text{macstring});$
 $\text{if } \text{verify}(x'_m, x_{mk}, x_{ma}) \text{ then}$
 $\text{let } i_{\perp}(k2b(x''_k)) = \text{dec}(x'_m, x_k) \text{ in } \overline{c_B}\langle \rangle$

Role of B:

- 1 n copies, as for Q_A .
- 2 The process Q_B waits for the message on channel c_B .
- 3 It verifies the MAC, decrypts, and stores the key in x''_k .

Example (role of A)

$$A \rightarrow B : e = \{x'_k\}_{x_k}, \text{mac}(e, x_{mk}) \quad x'_k \text{ fresh}$$

$Q_A = !^{i \leq n} c_A(); \text{new } x'_k : \text{key}; \text{new } x''_r : \text{coins};$
 $\text{let } x_m : \text{bitstring} = \text{enc}(k2b(x'_k), x_k, x''_r) \text{ in}$
 $\overline{c_A}\langle x_m, \text{mac}(x_m, x_{mk}) \rangle$

Role of A:

- 1 $!^{i \leq n}$ represents n copies, indexes by $i \in [1, n]$
The protocol can be run n times (polynomial in the security parameter).
- 2 The process is triggered when a message is sent on c_A by the adversary.
- 3 The process **chooses a fresh key x'_k and sends the message** on channel c_A .

Example (summary)

$$A \rightarrow B : e = \{x'_k\}_{x_k}, \text{mac}(e, x_{mk}) \quad x'_k \text{ fresh}$$

$Q_0 = \text{start}(); \text{new } x_r : \text{keyseed}; \text{let } x_k : \text{key} = \text{kgen}(x_r) \text{ in}$
 $\text{new } x'_r : \text{mkeyseed}; \text{let } x_{mk} : \text{mkey} = \text{mkgen}(x'_r) \text{ in } \overline{c}\langle \rangle; (Q_A \mid Q_B)$

$Q_A = !^{i \leq n} c_A(); \text{new } x'_k : \text{key}; \text{new } x''_r : \text{coins};$
 $\text{let } x_m : \text{bitstring} = \text{enc}(k2b(x'_k), x_k, x''_r) \text{ in}$
 $\overline{c_A}\langle x_m, \text{mac}(x_m, x_{mk}) \rangle$

$Q_B = !^{i' \leq n} c_B(x'_m : \text{bitstring}, x_{ma} : \text{macstring});$
 $\text{if } \text{verify}(x'_m, x_{mk}, x_{ma}) \text{ then}$
 $\text{let } i_{\perp}(k2b(x''_k)) = \text{dec}(x'_m, x_k) \text{ in } \overline{c_B}\langle \rangle$

Arrays

Arrays replace **lists** often used in cryptographic proofs.

They avoid the need for explicit list insertion instructions, which would be hard to guess for an automatic tool.

A variable defined under a replication is implicitly an **array**:

$$Q_A = !^{i \leq n} c_A(); \text{new } x'_k[i] : \text{key}; \text{new } x''_r[i] : \text{coins};$$

$$\text{let } x_m[i] : \text{bitstring} = \text{enc}(k2b(x'_k[i]), x_k, x''_r[i]) \text{ in}$$

$$\overline{c_A}(x_m[i], \text{mac}(x_m[i], x_{mk}))$$

Requirements:

- Only variables with the current indexes can be assigned.
- Variables may be defined at several places, but only one definition can be executed for the same indexes.

(if ... then let $x = M$ in P else let $x = M'$ in P' is ok)

So each array cell can be **assigned at most once**.

Indistinguishability as observational equivalence

Two processes (games) Q_1, Q_2 are **observationally equivalent** when the adversary has a negligible probability of distinguishing them:

$$Q_1 \approx Q_2$$

In the formal definition, the adversary is represented by an acceptable evaluation context $C ::= [] \mid C \mid Q \mid C \mid \text{newChannel } c; C$.

- Observational equivalence is an equivalence relation.
- It is **contextual**: $Q_1 \approx Q_2$ implies $C[Q_1] \approx C[Q_2]$ where C is any acceptable evaluation context.

Arrays (continued)

find performs an **array lookup**:

$$!^{i \leq N} \dots \text{let } x = M \text{ in } P$$

$$| !^{i' \leq N'} c(y : T) \text{find } j \leq N \text{ suchthat defined}(x[j]) \wedge y = x[j] \text{ then } \dots$$

Note that **find** is here used outside the scope of x .

This is the only way of getting access to values of variables in other sessions.

When several array elements satisfy the condition of the **find**, the returned index is chosen randomly, with uniform probability.

Proof technique

We transform a game G_0 into an observationally equivalent one using:

- **observational equivalences** $L \approx R$ given as **axioms** and that come from security assumptions on primitives. These equivalences are used inside a context:

$$G_1 \approx C[L] \approx C[R] \approx G_2$$

- **syntactic transformations**: simplification, expansion of assignments, ...

We obtain a **sequence of games** $G_0 \approx G_1 \approx \dots \approx G_m$, which implies $G_0 \approx G_m$.

If some equivalence or trace property holds with overwhelming probability in G_m , then it also holds with overwhelming probability in G_0 .

MACs: security definition

A MAC scheme:

- (Randomized) key generation function *mkgen*.
- MAC function *mac(m, k)* takes as input a message *m* and a key *k*.
- Verification function *verify(m, k, t)* such that

$$\text{verify}(m, k, \text{mac}(m, k)) = \text{true}.$$

A MAC guarantees the integrity and authenticity of the message because only someone who knows the secret key can build the mac.

More formally, an adversary \mathcal{A} that has oracle access to *mac* and *verify* has a **negligible probability to forge a MAC** (UF-CMA):

$$\max_{\mathcal{A}} \Pr[\text{verify}(m, k, t) \mid k \xleftarrow{R} \text{mkgen}; (m, t) \leftarrow \mathcal{A}^{\text{mac}(\cdot, k), \text{verify}(\cdot, k, \cdot)}]$$

is negligible, when the adversary \mathcal{A} has not called the *mac* oracle on message *m*.

MACs: formal implementation

$\text{verify}(m, \text{mkgen}(r), \text{mac}(m, \text{mkgen}(r))) = \text{true}$

```
!^{N''} new r : mkeyseed; (
  !^N (x : bitstring) → mac(x, mkgen(r)),
  !^N (m : bitstring, t : macstring) → verify(m, mkgen(r), t))
≈
!^{N''} new r : mkeyseed; (
  !^N (x : bitstring) → mac(x, mkgen(r)),
  !^N (m : bitstring, t : macstring) →
    find j ≤ N suchthat defined(x[j]) ∧ (m = x[j]) ∧
      verify(m, mkgen(r), t) then true else false)
```

The prover understands such specifications of primitives.

They can be reused in the proof of many protocols.

MACs: intuitive implementation

By the previous definition, up to negligible probability,

- the adversary cannot forge a correct MAC
- so when verifying a MAC with *verify(m, k, t)* and $k \xleftarrow{R} \text{mkgen}$ is used only for generating and verifying MACs, the verification can succeed **only if *m* is in the list (array) of messages whose *mac* has been computed** by the protocol
- so we can replace a call to *verify* with an array lookup:
if the call to *mac* is *mac(x, k)*, we replace *verify(m, k, t)* with

**find $j \leq N$ suchthat defined(x[j]) ∧
 (m = x[j]) ∧ verify(m, k, t) then true else false**

MACs: formal implementation

$\text{verify}(m, \text{mkgen}(r), \text{mac}(m, \text{mkgen}(r))) = \text{true}$

```
!^{N''} new r : mkeyseed; (
  !^N (x : bitstring) → mac(x, mkgen(r)),
  !^N (m : bitstring, t : macstring) → verify(m, mkgen(r), t))
≈
!^{N''} new r : mkeyseed; (
  !^N (x : bitstring) → mac'(x, mkgen'(r)),
  !^N (m : bitstring, t : macstring) →
    find j ≤ N suchthat defined(x[j]) ∧ (m = x[j]) ∧
      verify'(m, mkgen'(r), t) then true else false)
```

The prover understands such specifications of primitives.

They can be reused in the proof of many protocols.

MACs: formal implementation

The prover applies the previous rule automatically in **any (polynomial-time) context**, perhaps containing **several occurrences** of *mac* and of *verify*:

- Each occurrence of *mac* is replaced with *mac'*.
- Each occurrence of *verify* is replaced with a **find** that looks in all arrays of computed MACs (one array for each occurrence of function *mac*).

Syntactic transformations

- **Single assignment renaming**: when a variable is assigned at several places, rename it with a distinct name for each assignment. (Not completely trivial because of array references.)
- **Expansion of assignments**: replacing a variable with its value. (Not completely trivial because of array references.)
- **Move new**: move restrictions downwards in the game as much as possible, when there is no array reference to them. (Moving **new** $x : T$ under a **if** or a **find** duplicates it. A subsequent single assignment renaming will distinguish cases.)

IND-CPA symmetric encryption

We consider a non-deterministic, length-revealing encryption scheme that satisfies INDistinguishability under Chosen Plaintext Attacks (IND-CPA).

$$\text{dec}(\text{enc}(m, \text{kgen}(r), r'), \text{kgen}(r)) = i_{\perp}(m)$$

$$\begin{aligned} & !^{N'} \text{new } r : \text{keyseed}; !^N (x : \text{bitstring}) \rightarrow \\ & \quad \text{new } r' : \text{coins}; \text{enc}(x, \text{kgen}(r), r') \end{aligned}$$

\approx

$$\begin{aligned} & !^{N'} \text{new } r : \text{keyseed}; !^N (x : \text{bitstring}) \rightarrow \\ & \quad \text{new } r' : \text{coins}; \text{enc}'(Z(x), \text{kgen}'(r), r') \end{aligned}$$

$Z(x)$ is the bitstring of the same length as x containing only zeroes (for all $x : \text{nonce}$, $Z(x) = Z_{\text{nonce}}, \dots$).

Simplification and elimination of collisions

Terms are simplified according to equalities that come from:

- **Assignments**: **let** $x = M$ **in** P implies that $x = M$ in P
- **Tests**: **if** $M = N$ **then** P implies that $M = N$ in P
- **Definitions of cryptographic primitives**
- When a **find** guarantees that $x[j]$ is **defined**, equalities that hold at definition of x also hold under the find (after substituting j for the array indexes at the definition of x)
- **Elimination of collisions**: if x is created by **new** $x : T$, $x[i] = x[j]$ implies $i = j$, up to negligible probability (when T is large)

Proof of security properties: one-session secrecy

One-session secrecy: the adversary cannot distinguish any of the secrets from a random number with one test query.

Criterion for proving one-session secrecy of x :

x is defined by **new** $x[i] : T$ and there is a set of variables S such that only variables in S depend on x .

The output messages and the control-flow do not depend on x .

Proof strategy: advice

- One tries to execute each transformation given by the definition of a cryptographic primitive.
- When it fails, it tries to analyze why the transformation failed, and **suggests syntactic transformations** that could make it work.
- One tries to execute these syntactic transformations. (If they fail, they may also suggest other syntactic transformations, which are then executed.)
- We retry the cryptographic transformation, and so on.

Proof of security properties: secrecy

Secrecy: the adversary cannot distinguish the secrets from independent random numbers with several test queries.

Criterion for proving secrecy of x : same as one-session secrecy, plus $x[i]$ and $x[i']$ do not come from the same copy of the same restriction when $i \neq i'$.

Proof of the example: initial game

$$\begin{aligned}
 Q_0 &= \text{start}(); \text{new } x_r : \text{keyseed}; \text{let } x_k : \text{key} = \text{kgen}(x_r) \text{ in} \\
 &\quad \text{new } x'_r : \text{mkeyseed}; \text{let } x_{mk} : \text{mkey} = \text{mkgen}(x'_r) \text{ in } \overline{c}\langle \rangle; (Q_A \mid Q_B) \\
 Q_A &= !^{i \leq n} c_A(); \text{new } x'_k : \text{key}; \text{new } x''_r : \text{coins}; \\
 &\quad \text{let } x_m : \text{bitstring} = \text{enc}(k2b(x'_k), x_k, x''_r) \text{ in} \\
 &\quad \overline{c}_A\langle x_m, \text{mac}(x_m, x_{mk}) \rangle \\
 Q_B &= !^{i' \leq n} c_B(x'_m : \text{bitstring}, x_{ma} : \text{macstring}); \\
 &\quad \text{if } \text{verify}(x'_m, x_{mk}, x_{ma}) \text{ then} \\
 &\quad \text{let } i_{\perp}(k2b(x''_k)) = \text{dec}(x'_m, x_k) \text{ in } \overline{c}_B\langle \rangle
 \end{aligned}$$

Proof of the example: remove assignments x_{mk}

```

Q0 = start(); new xr : keyseed; let xk : key = kgen(xr) in
  new x'r : mkeyseed;  $\overline{c}$ ⟨⟩; (QA | QB)

QA = !i≤ncA(); new x'k : key; new x''r : coins;
  let xm : bitstring = enc(k2b(x'k), xk, x''r) in
     $\overline{c}_A$ ⟨xm, mac(xm, mkgen(x'r))⟩

QB = !i'≤ncB(x'm : bitstring, xma : macstring);
  if verify(x'm, mkgen(x'r), xma) then
    let i⊥(k2b(x''k)) = dec(x'm, xk) in  $\overline{c}_B$ ⟨⟩
    
```

Proof of the example: simplify

```

Q0 = start(); new xr : keyseed; let xk : key = kgen(xr) in
  new x'r : mkeyseed;  $\overline{c}$ ⟨⟩; (QA | QB)

QA = !i≤ncA(); new x'k : key; new x''r : coins;
  let xm : bitstring = enc(k2b(x'k), xk, x''r) in
     $\overline{c}_A$ ⟨xm, mac'(xm, mkgen'(x'r))⟩

QB = !i'≤ncB(x'm : bitstring, xma : macstring);
  find j ≤ n suchthat defined(xm[j]) ∧ x'm = xm[j] ∧
    verify'(x'm, mkgen'(x'r), xma) then
    let x''k = x'k[j] in  $\overline{c}_B$ ⟨⟩
    
```

$dec(x'_m, x_k) = dec(enc(k2b(x'_k[j]), x_k, x''_r[j]), x_k) = i_{\perp}(k2b(x'_k[j]))$

Proof of the example: security of the MAC

```

Q0 = start(); new xr : keyseed; let xk : key = kgen(xr) in
  new x'r : mkeyseed;  $\overline{c}$ ⟨⟩; (QA | QB)

QA = !i≤ncA(); new x'k : key; new x''r : coins;
  let xm : bitstring = enc(k2b(x'k), xk, x''r) in
     $\overline{c}_A$ ⟨xm, mac'(xm, mkgen'(x'r))⟩

QB = !i'≤ncB(x'm : bitstring, xma : macstring);
  find j ≤ n suchthat defined(xm[j]) ∧ x'm = xm[j] ∧
    verify'(x'm, mkgen'(x'r), xma) then
    let i⊥(k2b(x''k)) = dec(x'm, xk) in  $\overline{c}_B$ ⟨⟩
    
```

Proof of the example: remove assignments x_k

```

Q0 = start(); new xr : keyseed; new x'r : mkeyseed;  $\overline{c}$ ⟨⟩; (QA | QB)

QA = !i≤ncA(); new x'k : key; new x''r : coins;
  let xm : bitstring = enc(k2b(x'k), kgen(xr), x''r) in
     $\overline{c}_A$ ⟨xm, mac'(xm, mkgen'(x'r))⟩

QB = !i'≤ncB(x'm : bitstring, xma : macstring);
  find j ≤ n suchthat defined(xm[j]) ∧ x'm = xm[j] ∧
    verify'(x'm, mkgen'(x'r), xma) then
    let x''k = x'k[j] in  $\overline{c}_B$ ⟨⟩
    
```


Proof of the example: security of the encryption

```

Q0 = start(); new xr : keyseed; new x'r : mkeyseed;  $\bar{c}()$ ; (QA | QB)

QA = !i ≤ n cA(); new x'k : key; new x''r : coins;
  let xm : bitstring = enc'(Z(k2b(x'k)), kgen'(xr), x''r) in
     $\bar{c}_A(x_m, mac'(x_m, mkgen'(x'_r)))$ 

QB = !i' ≤ n cB(x'm : bitstring, xma : macstring);
  find j ≤ n suchthat defined(xm[j]) ∧ x'm = xm[j] ∧
    verify'(x'm, mkgen'(x'_r), xma) then
    let x''k = x'k[j] in  $\bar{c}_B()$ 

```

Proof of the example: secrecy

```

Q0 = start(); new xr : keyseed; new x'r : mkeyseed;  $\bar{c}()$ ; (QA | QB)

QA = !i ≤ n cA(); new x'k : key; new x''r : coins;
  let xm : bitstring = enc'(Zk, kgen'(xr), x''r) in
     $\bar{c}_A(x_m, mac'(x_m, mkgen'(x'_r)))$ 

QB = !i' ≤ n cB(x'm : bitstring, xma : macstring);
  find j ≤ n suchthat defined(xm[j]) ∧ x'm = xm[j] ∧
    verify'(x'm, mkgen'(x'_r), xma) then
    let x''k = x'k[j] in  $\bar{c}_B()$ 

```

Preserves the one-session secrecy of x''_k but not its secrecy.

Proof of the example: simplify

```

Q0 = start(); new xr : keyseed; new x'r : mkeyseed;  $\bar{c}()$ ; (QA | QB)

QA = !i ≤ n cA(); new x'k : key; new x''r : coins;
  let xm : bitstring = enc'(Zk, kgen'(xr), x''r) in
     $\bar{c}_A(x_m, mac'(x_m, mkgen'(x'_r)))$ 

QB = !i' ≤ n cB(x'm : bitstring, xma : macstring);
  find j ≤ n suchthat defined(xm[j]) ∧ x'm = xm[j] ∧
    verify'(x'm, mkgen'(x'_r), xma) then
    let x''k = x'k[j] in  $\bar{c}_B()$ 

```

$Z(k2b(x'_k)) = Z_k$

Experiments

Tested on the following protocols (original and corrected versions):

- Otway-Rees (shared-key)
 - Yahalom (shared-key)
 - Denning-Sacco (public-key)
 - Woo-Lam shared-key and public-key
 - Needham-Schroeder shared-key and public-key
 - Full domain hash signature (with D. Pointcheval)
 - Encryption schemes of Bellare-Rogaway'93 (with D. Pointcheval)
- Shared-key encryption is implemented as encrypt-then-MAC, using a IND-CPA encryption scheme.
 (For Otway-Rees, we also considered a SPRP encryption scheme,
 a IND-CPA + INT-CTXT encryption scheme,
 a IND-CCA2 + IND-PTXT encryption scheme.)

Public-key encryption is assumed to be IND-CCA2.

We prove secrecy of session keys and correspondence properties.

Results (1)

In most cases, the prover succeeds in proving the desired properties when they hold, and obviously it always fails to prove them when they do not hold.

Only cases in which the prover fails although the property holds:

- Needham-Schroeder public-key when the exchanged key is the nonce N_A .
- Needham-Schroeder shared-key: fails to prove that $N_B[i] \neq N_B[i'] - 1$ with overwhelming probability, where N_B is a nonce
- Showing that the encryption scheme $\mathcal{E}(m, r) = f(r) \| H(r) \oplus m \| H'(m, r)$ is IND-CCA2.

Conclusion

Hopefully a promising approach.

Future extensions:

- Extension to other cryptographic primitives, in particular Diffie-Hellman, full support of XOR.
- More game transformations.
- More case studies.

More information: <http://www.cryptoverif.ens.fr/>

I warmly thank David Pointcheval for his advice and explanations of the computational proofs of protocols. This project would not have been possible without him.

This work was partly supported by the ANR project ARA SSIA FormaCrypt.

Results (2)

- Some public-key protocols need manual proofs.
(Give the cryptographic proof steps and single assignment renaming instructions.)
- Runtime: 7 ms to 35 s, average: 5 s on a Pentium M 1.8 GHz.
- A detailed case study of Kerberos V, with and without its public-key extension PKINIT (AsiaCCS'08, with with Aaron D. Jaggard, Andre Scedrov, and Joe-Kai Tsay).
- Starts being used by others: Verification of F# implementations, including TLS, by Microsoft Research and the MSR-INRIA lab.