

Proving unlinkability using ProVerif through desynchronized bi-processes

Stéphanie Delaune

Univ Rennes, CNRS, IRISA, Spicy team

→ joint work with David Baelde and Alexandre Debant

CSF - July 10, 2023



Formal verification of cryptographic protocols

Security protocol design is critical and **error-prone** as illustrated by many attacks: FREAK, Logjam, ...



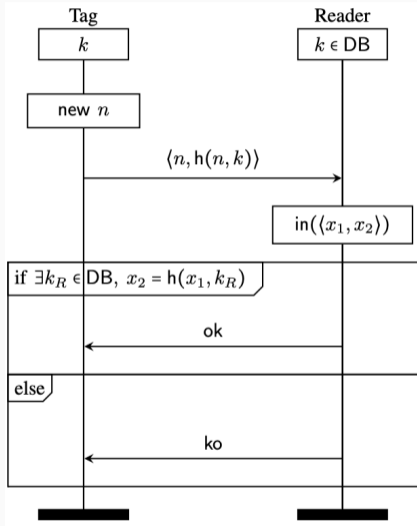
Use **formal methods** to improve confidence:

- prove the **absence of attacks** under certain **assumptions**; or
- identify **weaknesses**.

Many verification tools already exist:

- **Proverif**, Tamarin, AKISS, DeepSec, AVISPA, Squirrel, ...

Running example: Basic Hash protocol



- Each tag stores a secret key k that is never updated.
- Readers have access to a database DB containing all the keys.

ProVerif in a nutshell

→ mainly developed by Bruno Blanchet (Prosecco team, Inria Paris)

`http://proverif.inria.fr/`

An automatic tool to analyse protocols in the **symbolic model**.

- successfully used for many **large-scale case studies**: TLS 1.3, ...
- protocols are modelled using a **process algebra**;
- both reachability and **equivalence-based properties**;
- security analysis done for an **unbounded number of sessions**;
- **No miracle**: the tool may return `cannot be proved` or `never terminates`.

Unlinkability

(ISO/IEC 15408)

“Ensuring that a user may make multiple uses of a service or resource without others being able to link these uses together.”



Informally, an observer/attacker can not observe the difference between:

1. a situation where the same device/tag may be used **twice (or even more)**;
2. a situation where each device/tag is used **at most once**.

Unlinkability

(ISO/IEC 15408)

“Ensuring that a user may make multiple uses of a service or resource without others being able to link these uses together.”



Informally, an observer/attacker can not observe the difference between:

1. a situation where the same device/tag may be used **twice (or even more)**;
2. a situation where each device/tag is used **at most once**.

More formally,

$$! \text{ new } k.\text{insert DB}(k). \left(! \text{ Tag}(k) \mid ! \text{ Reader} \right)$$

$\stackrel{?}{\approx}$

$$! \text{ new } k.\text{insert DB}(k). \left(\text{ Tag}(k) \mid ! \text{ Reader} \right)$$

→ the notion of equivalence remains to be defined

ProVerif (but also Tamarin) can only prove a restricted form of equivalence, namely **diff-equivalence**, which is too limiting to establish unlinkability.

¹<https://github.com/tamarin-prover/tamarin-prover/issues/324>

ProVerif (but also Tamarin) can only prove a restricted form of equivalence, namely **diff-equivalence**, which is too limiting to establish unlinkability.

Some solutions to overcome this limitation:

- Establish unlinkability using an **indirect approach** (sufficient conditions)
e.g. [Solène Moreau PhD thesis, 21]
- Use **restrictions**: a feature available in Tamarin (2005), and in ProVerif (2022).

¹<https://github.com/tamarin-prover/tamarin-prover/issues/324>

ProVerif (but also Tamarin) can only prove a restricted form of equivalence, namely **diff-equivalence**, which is too limiting to establish unlinkability.

Some solutions to overcome this limitation:

- Establish unlinkability using an **indirect approach** (sufficient conditions)
e.g. [Solène Moreau PhD thesis, 21]
- Use **restrictions**: a feature available in Tamarin (2005), and in ProVerif (2022).
Tamarin: incorrectly handled for equivalence¹, now formally justify for Type-0 (very specific class) [Paradzik, 22]
ProVerif: Need to be manipulated with a lot of care. **Restrictions for equivalence discard bi-traces!**

¹<https://github.com/tamarin-prover/tamarin-prover/issues/324>

Our contributions

We design a **transformation** (in 2 steps) allowing us to transform a ProVerif model \mathcal{M} into another one \mathcal{M}' such that:

If ProVerif succeeds on \mathcal{M}' then equivalence holds on \mathcal{M} .

Our contributions

We design a **transformation** (in 2 steps) allowing us to transform a ProVerif model \mathcal{M} into another one \mathcal{M}' such that:

If ProVerif succeeds on \mathcal{M}' then equivalence holds on \mathcal{M} .

Our transformation contains two main steps:

1. We **dissociate** the two processes that forms that bi-process.
Possible using the option: `allowDiffPatterns`
2. We generate some **axioms** (and prove them correct) to help the analysis.

The transformation has been **implemented** and sucessfully used on **several case studies**.

High-level view of ProVerif

Protocols as processes

→ a programming language with constructs for **concurrency** and **communication**

(applied-pi calculus [Abadi & Fournet, 01])

P, Q	$:=$	0	null process
		$\text{in}(c, x); P$	input
		$\text{out}(c, M); P$	output
		$\text{new } n; P$	name generation
		$\text{let } x = D \text{ in } P \text{ else } Q$	conditional
		$!P$	replication
		$(P \mid Q)$	parallel composition

Protocols as processes

→ a programming language with constructs for **concurrency** and **communication**

(applied-pi calculus [Abadi & Fournet, 01])

P, Q	$:=$	0	null process
		$\text{in}(c, x); P$	input
		$\text{out}(c, M); P$	output
		$\text{new } n; P$	name generation
		$\text{let } x = D \text{ in } P \text{ else } Q$	conditional
		$!P$	replication
		$(P \mid Q)$	parallel composition
		$\text{event}(e); P$	event
		$\text{insert } \text{tbl}(M); P$	insertion
		$\text{get } \text{tbl}(x) \text{ st. } D \text{ in } P \text{ else } Q$	lookup
		\dots	

Messages/Computations as terms

Terms are built over a set of **names** \mathcal{N} , and function symbols $\Sigma_c \cup \Sigma_d$ equipped with an **equational theory** E and **rewriting rules** for destructors.

Example:

- **constructor symbols**: $\Sigma_c = \{\langle \rangle, \text{proj}_1, \text{proj}_2, \text{h}, \text{true}\}$;
- $E = \{\text{proj}_1(\langle x_1, x_2 \rangle) = x_1, \text{proj}_2(\langle x_1, x_2 \rangle) = x_2\}$;
- **destructor symbols**: $\Sigma_d = \{\text{eq}\}$;
- rewriting rule: $\text{eq}(x, x) \rightarrow \text{true}$.
- all the function symbols are public (available to the attacker);

Going back to Basic Hash

We consider:


- $T(k) = \text{new } n; \text{out}(c, \langle n, h(n, k) \rangle)$.
- $R =$
 $\text{in}(c, y); \text{get } db(k) \text{ st. } \text{eq}(h(\text{proj}_1(y), k), \text{proj}_2(y)) \text{ in out}(c, \text{ok}) \text{ else out}(c, \text{ko}).$

The real system corresponds to the following process:

$$! R \mid (! \text{new } k; \text{insert } \text{keys}(k); ! T(k))$$


Semantics (some selected rules)

Labelled transition system over **configurations**:



Semantics (some selected rules)

Labelled transition system over configurations:



- OUT** $(\{\text{out}(c, M); P\} \uplus \mathcal{P}; \Phi; \mathcal{S}) \xrightarrow{\text{out}(c, w_i)} (\{P\} \uplus \mathcal{P}; \Phi \cup \{w_i \mapsto M\}; \mathcal{S})$ with $i = |\Phi|$
- IN** $(\{\text{in}(c, x); P\} \uplus \mathcal{P}; \Phi; \mathcal{S}) \xrightarrow{\text{in}(c, R)} (\{P\{x \mapsto M\}\} \uplus \mathcal{P}; \Phi; \mathcal{S})$ with $R\Phi \Downarrow =_{\text{E}} M$
- GET-T** $(\{\text{get } \text{tbl}(x) \text{ st. } D \text{ in } P \text{ else } Q\} \uplus \mathcal{P}; \Phi; \mathcal{S}) \xrightarrow{\tau} (\{P\{x \mapsto M\}\} \uplus \mathcal{P}; \Phi; \mathcal{S})$
with $\text{tbl}(M) \in \mathcal{S}$, and $D\{x \mapsto M\} \Downarrow =_{\text{E}} \text{true}$
- ...

Trace equivalence

Static equivalence between frames: $\Phi \sim_s \Phi'$.

Any **test** that holds in Φ also holds in Φ' (and conversely).

Example: $\{w_1 \mapsto \langle n, h(n, k) \rangle; w_2 \mapsto k\} \not\sim_s \{w_1 \mapsto \langle n, h(n, k) \rangle; w_2 \mapsto k'\}$

→ with the test $h(\text{proj}_1(w_1), w_2) \stackrel{?}{=} \text{proj}_2(w_1)$.

Trace equivalence

Static equivalence between frames: $\Phi \sim_s \Phi'$.

Any **test** that holds in Φ also holds in Φ' (and conversely).

Example: $\{w_1 \mapsto \langle n, h(n, k) \rangle; w_2 \mapsto k\} \not\sim_s \{w_1 \mapsto \langle n, h(n, k) \rangle; w_2 \mapsto k'\}$

→ with the test $h(\text{proj}_1(w_1), w_2) \stackrel{?}{=} \text{proj}_2(w_1)$.

Trace equivalence between configurations: $K \approx_t K'$.

For any execution trace $K \xrightarrow{\text{tr}} (\mathcal{P}; \Phi; \mathcal{S})$ there exists an execution $K' \xrightarrow{\text{tr}} (\mathcal{P}'; \Phi'; \mathcal{S}')$ such that $\Phi \sim_s \Phi'$ (and conversely)

Example:

$!R \mid (!\text{new } k; \text{insert } \text{keys}(k); !T(k)) \approx_t !R \mid (!\text{new } k; \text{insert } \text{keys}(k); T(k))$

→ an equivalence that ProVerif (and also Tamarin) is **not able to prove** directly.

Going back to diff-equivalence

How it works (or not)?

- form a **bi-process** B using the operator **choice** $[M_L, M_R]$;
- both sides of the bi-processes have to evolve **simultaneously** to be declared in diff-equivalence (and this implies $\text{fst}(B) \approx_t \text{snd}(B)$)

→ the semantics is given by a labelled transition system over bi-configurations $(\mathcal{P}; \Phi; \mathcal{S})$ where messages and computations may contain the **choice** operator.

Going back to diff-equivalence

How it works (or not)?

- form a **bi-process** B using the operator $\text{choice}[M_L, M_R]$;
- both sides of the bi-processes have to evolve **simultaneously** to be declared in diff-equivalence (and this implies $\text{fst}(B) \approx_t \text{snd}(B)$)

→ the semantics is given by a labelled transition system over bi-configurations $(\mathcal{P}; \Phi; \mathcal{S})$ where messages and computations may contain the **choice** operator.

Example - Basic Hash protocol

$$B = !R \mid (!\text{new } k; !\text{new } kk; \text{insert } db(\text{choice}[k, kk]); T(\text{choice}[k, kk]))$$

We have that:

- $\text{fst}(B) = !R \mid !\text{new } k; !\text{insert } db(k); T(k)$ (* real situation *)
- $\text{snd}(B) = !R \mid !!\text{new } kk; \text{insert } db(kk); T(kk)$ (* ideal situation *)

Why diff-equivalence is too strong?

$$B = !R \mid (!\text{new } k; !\text{new } kk; \text{insert } db(\text{choice}[k, kk])); T(\text{choice}[k, kk])$$

Let's consider a scenario with:

- 1 reader;
- 2 tags: $T(\text{choice}[k, kk_1])$,
and $T(\text{choice}[k, kk_2])$.

Why diff-equivalence is too strong?

$$B = !R \mid (!\text{new } k; !\text{new } kk; \text{insert } db(\text{choice}[k, kk]); T(\text{choice}[k, kk]))$$

Let's consider a scenario with:

- 1 reader;
- 2 tags: $T(\text{choice}[k, kk_1])$,
and $T(\text{choice}[k, kk_2])$.

DB	left	right
line 1	k	kk_1
line 2	k	kk_2

The frame contains: $w_1 = \langle n_1, h(n, \text{choice}[k, kk_1]) \rangle$.

Why diff-equivalence is too strong?

$$B = !R \mid (!\text{new } k; !\text{new } kk; \text{insert } db(\text{choice}[k, kk]); T(\text{choice}[k, kk]))$$

Let's consider a scenario with:

- 1 reader;
- 2 tags: $T(\text{choice}[k, kk_1])$,
and $T(\text{choice}[k, kk_2])$.

DB	left	right
line 1	k	kk_1
line 2	k	kk_2

The frame contains: $w_1 = \langle n_1, h(n, \text{choice}[k, kk_1]) \rangle$.

On line 2, with w_1 in input for process R , the bi-process B will diverge.

→ Thus, Proverif returns **cannot be proved** on this example.

→ support for **axioms**, **lemmas**, and **restrictions** as in Tamarin.

Syntax: This gives the user the possibility to write correspondence queries of the form:

$$\text{event}(e_1) \wedge \dots \wedge \text{event}(e_n) \Rightarrow \psi$$

with $\psi, \psi' = \text{true} \mid \text{false} \mid \text{event}(e) \mid M = N \mid M \neq N \mid \psi \wedge \psi' \mid \psi \vee \psi'$

Semantics: An execution trace T satisfies ρ (noted $T \vdash \rho$) if whenever T contains instances of $\text{event}(e_i)$ at some timepoint τ_i for each i , then T also satisfies ψ .

Desynchronized bi-processes

We consider an extension of standard bi-processes using the `allowDiffPatterns` option available in ProVerif since 2018.

→ this allow us to use `choice[xL, xR]` for variable bindings in `let`, `get`, and `input`.

Desynchronized bi-processes

We consider an extension of standard bi-processes using the `allowDiffPatterns` option available in ProVerif since 2018.

→ this allow us to use `choice[xL, xR]` for variable bindings in `let`, `get`, and `input`.

Example: $B = \text{in}(c, \text{choice}[x^L, x^R]); \text{out}(c, \langle x^L, x^R \rangle).$

Warning !

Closed bi-processes of this form have a well-defined semantics in Proverif and we can study diff-equivalence (convergence of all the bi-traces). However, this does **not** imply:

$$\text{fst}(B) \approx_t \text{snd}(B)$$

Our transformation

Main Goal

Transform a ProVerif model \mathcal{M} of unlinkability into another model \mathcal{M}' such that:

- diff-equivalence is verified with ProVerif on the transformed model \mathcal{M}' ; and
- diff-equivalence on \mathcal{M}' implies trace equivalence for the original model \mathcal{M} .

Main Goal

Transform a ProVerif model \mathcal{M} of unlinkability into another model \mathcal{M}' such that:

- diff-equivalence is verified with ProVerif on the transformed model \mathcal{M}' ; and
- diff-equivalence on \mathcal{M}' implies trace equivalence for the original model \mathcal{M} .

Two main steps

1. **duplicate the get instructions** in \mathcal{M} to dissociate the two parts of the bi-process;
2. add **some axioms** to help ProVerif to reason on our new model.

Desynchronizing the two parts of the biprocess

Instead of performing a get instruction to access a bi-record in the keys table, we perform **two get instructions in a row** to access two records in the keys table.

→ This allows us to choose two **different** records for the left and for the right.

Desynchronizing the two parts of the biprocess

Instead of performing a get instruction to access a bi-record in the keys table, we perform **two get instructions in a row** to access two records in the keys table.

→ This allows us to choose two **different** records for the left and for the right.

Example:

```
in(c, choice[xL, xR]);
get db(choice[yL, -]) st. eq(proj2(xL), h(proj1(xL, yL))) in
  get db(choice[-, yR]) st. eq(proj2(xR), h(proj1(xR, yR))) in out(c, choice[ok, ok])
  else out(c, choice[ok, ko])
else
  get db(choice[-, yR]) st. eq(proj2(xR), h(proj1(xR, yR))) in out(c, choice[ko, ok])
  else out(c, choice[ko, ko])
```

Refining the analysis in the failure branches

We illustrate this on a very simple example.

Before, ...

```
B = insert tbl(ok);  
    get tbl(x) st. true in out(c, ok)  
        else out(c, choice[okL, okR])
```

... and ProVerif can **not proved** equivalence (whereas it holds).

Refining the analysis in the failure branches

We illustrate this on a very simple example.

After, ...

```
B = event(Inserted(ok)); insert tbl(ok);  
    get tbl(x) st. true in out(c, ok)  
    else event(Fail()); out(c, choice[okL, okR])
```

... together with the following axiom:

$$\text{event(Fail())} \wedge \text{event(Inserted(choice[y^L, y^R]))} \Rightarrow \text{false.}$$

→ On this model, ProVerif is able to conclude that equivalence holds.

Refining the analysis in the failure branches

We illustrate this on a very simple example.

After, ...

```
B = event(Inserted(ok)); insert tbl(ok);  
    get tbl(x) st. true in out(c, ok)  
    else event(Fail()); out(c, choice[okL, okR])
```

... together with the following axiom:

$$\text{event(Fail())} \wedge \text{event(Inserted(choice}[y^L, y^R]))} \Rightarrow \text{false.}$$

→ On this model, ProVerif is able to conclude that equivalence holds.

Going back to the Basic Hash protocol

$$\begin{aligned} \text{event(FailL}(x^L)) \wedge \text{event(Inserted(choice}[y^L, y^R]))} &\Rightarrow \text{proj}_2(x^L) \neq h(\text{proj}_1(x^L), y^L) \\ \text{event(FailR}(x^R)) \wedge \text{event(Inserted(choice}[y^L, y^R]))} &\Rightarrow \text{proj}_2(x^R) \neq h(\text{proj}_1(x^R), y^R) \end{aligned}$$

Theorem

Let $\mathcal{M} = (B_0, \emptyset, \mathcal{A}_X, \mathcal{L})$ be a ProVerif standard model (B_0 is separated), and $\mathcal{M}' = (B', \emptyset, \mathcal{A}_X \cup \mathcal{A}_{X'}, \mathcal{L})$ be the model obtained after applying our transformation.

Moreover, we assume that:

- for all $\varrho \in \mathcal{A}_X$, we have that $\text{traces}(B_0) \vdash \varrho$;
- for all $\varrho \in \mathcal{A}_X$, we have that $\text{traces}(B') \vdash \varrho$;
- ProVerif returns `diff-equivalence is true` on \mathcal{M}' .

We conclude that $\text{fst}(B_0) \approx_t \text{snd}(B_0)$.

Case studies

Implementation

The two steps of the transformation have been implemented (\approx 2k Ocaml LoC).

Case studies

Basic Hash, Hash-Lock, Feldhofer, a variant of LAK, OSK.

→ ProVerif is able to conclude on all these examples !



Our approach significantly improves automation regarding unlinkability.

Future Work

- better integration in ProVerif;
- beyond unlinkability (e.g. privacy in e-voting);
- other difficulty: dealing with mutable states (GSVerif).

→ if you are interested in this subject (going beyond diff-equivalence using ProVerif), you should attend Vincent's talk (tomorrow afternoon).

Questions ?