

Verifying security protocols with ProVerif and StatVerif

Mark D. Ryan
University of Birmingham

FOSAD summer school, Bertinoro, Italy
3–7 September 2012

- L1: Security protocols and verification
- L2: ProVerif language, tool, theory
- L3: Example: electronic voting
- L4: Example: stateful systems

Outline

What are security protocols?

Security protocols (also known as cryptographic protocols) are **distributed procedures** that employ cryptography to achieve a **security goal**.

Example participating agents:


- *Client and Server*
- *Application and TPM*
- *VM₁ and VMM*
- *Voter and Collector*
- *Alice and Bob*

Example security goals:

- *Authentication*
- *Key agreement*
- *Secure communication*
- *Privacy*
- *Attestation*


Protocols are usually simple, but often subtle. That makes them ideal for automated reasoning.

Example: customer \leftrightarrow bank



- Mutual authentication
 - Q: How does the bank authenticate itself to the customer?
- Non-repudiation

Example: voter \leftrightarrow administrator





- Mutual authentication
- Incoercibility, verifiability
 - Anything else?

Example: user process ↔ TPM

```
#include "stdio.h"
...

int main(int argc, char **argv)
{
    ...
    y = tpm_decrypt(tpm, x);
    ...
    return 0;
}
```



- Mutual authentication
 - In a stateful context

Outline

Approaches to verifying protocols

Computational approach

(a.k.a. “provable security”)

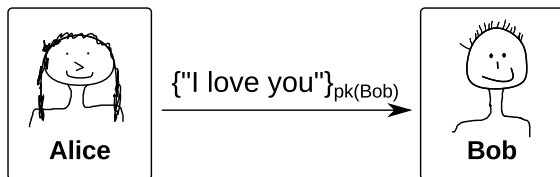
- data are bitstrings
- proof is reduction to a “hardness assumption” (e.g. the RSA assumption, which expresses hardness of factorisation)

Symbolic approach

(a.k.a. “formal methods security”)

- data are terms in a calculus
- proof is a proof in that calculus

Computational approach example



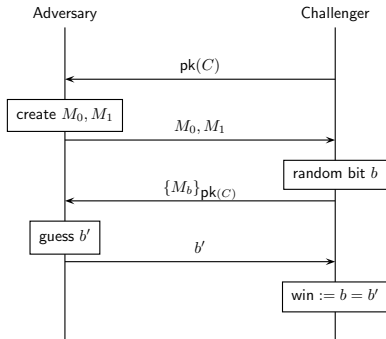
They've decided to use RSA-OAEP, so Bob has created a secret key (d, n) and a public key (e, n) such that $n = p \times q$; p, q are prime; and $d \times e = 1 \pmod{\phi(n)}$. Alice chooses a random r and sends $(x \parallel y)^e \pmod{n}$, where

$$\begin{aligned}x &= \text{"I love you"}00 \dots 0 \oplus G(r) \\y &= H(x) \oplus r\end{aligned}$$

and G, H are suitable hash functions.

Definition

The public key encryption scheme $\{\cdot\}_{pk(\cdot)}$ is IND-CPA secure if the adversary has no better than 0.5 chance of winning this game.




Theorem (Random oracle model)

If the RSA assumption holds (roughly, prime factorisation is hard) then RSA-OAEP is IND-CPA secure. (In fact, it's IND-CCA2 secure.)

Symbolic approach example

Consider the Needham-Schroeder-Lowe protocol.



where N_A and N_B are nonces chosen by Alice and Bob respectively, and Alice and Bob have the correct values for each other's public keys.

We don't know exactly what the encryption scheme is but we assume the functions $enc(\cdot, \cdot)$, $dec(\cdot, \cdot)$ and $pk(\cdot)$ which are related by the equation

$$dec(k, enc(pk(k), x)) = x$$

and no other equation.

Consider the term algebra $(N, pk, enc, dec, \langle \cdot, \cdot \rangle, fst, snd)$ where N is a set of names, $A, B, sk_A, sk_B \in N$, and there are three equations, namely:

$$\begin{aligned}fst(\langle x, y \rangle) &= x \\snd(\langle x, y \rangle) &= y \\dec(x, enc(pk(x), y)) &= y\end{aligned}$$

Suppose Alice is an oracle that, on input x , outputs $enc(y, z)$ where $y = snd(snd(dec(sk_A, x)))$ and $z = fst(snd(dec(sk_A, x)))$, and Bob is an oracle that, on input x , chooses a random r and outputs $enc(y, z)$ where $y = snd(dec(sk_B, x))$ and $z = (fst(dec(sk_B, x)), r, pk(sk_B))$. Also, Alice willingly chooses a value r and creates $enc(x, (r, pk(sk_A)))$ for any x at any time.

Suppose an attacker can:

- intercept and inject messages;
- apply any function to any message he knows;
- use the oracles Alice and Bob any number of times in any way;
- knows names a, b, c, \dots, A and B but not sk_A and sk_B .

Claim

- The attacker can't obtain the names sk_A, sk_B, N_A, N_B .
- Alice successfully executes her protocol iff Bob successfully executes his protocol, and on the same data.

Comparing computational and symbolic approaches

- Which do you find more convincing?
- Which might be more scalable or automatable?

Computational vs. symbolic approaches

Computational	Symbolic
more concrete	more abstract
more exact	more idealistic
manual	more automatable
small scale	larger scale

Two views of verification

Provable security vs. Formal methods

- Provable security provides **stronger promises**
- But, *“proofs are so turgid that other specialists don’t even read them”* [KoblitzMenezes’04]
- Formal methods are **simpler**, specifications are **nicer** and **automated** support is available, but it is more abstract
- Both methods are based on a **model**, and therefore may **fail** to detect certain kinds of attack
- There is a gap between the abstract formal model and the actual implementation

Reconciling two views of cryptography

- Computational soundness: in some circumstances, one can prove that a proof in the symbolic model is **sound** in the computational model.

Outline

Symbolic approach: some methods/tools

ProVerif B. Blanchet, and others, 2001-

AVISPA Large EU project and team, 2005-

Scyther C. Cremers, 2008-

Casper/FDR G. Lowe, 1998

AKISS Kremer, Chadha, 2012-

Athena D. Song, 1999

NRL C. Meadows, 1994

Isabelle/HOL L. Paulson, G. Bella, 1998-

ProVerif syntax: terms

$L, M, N, \dots ::=$	
a, b, c, \dots	name
x, y, z, \dots	variable
$f(M_1, \dots, M_l)$	function application

Equational theory

Suppose we have defined nullary function ok , unary function pk , binary functions dec , $senc$, $sdec$, $sign$ and ternary functions enc and $checksign$.

equation $sdec(x, senc(x, y))$	$= y$
equation $dec(x, enc(pk(x), r, y))$	$= y$
equation $checksign(pk(x), sign(x, y))$	$= ok$
equation $getmess(sign(x, y))$	$= y$

ProVerif syntax: processes

$P, Q, R, \dots ::=$

0

null process

$P \mid Q$

parallel composition

$!P$

replication

$\text{new } n; P$

name restriction

$\text{in}(M, x); P$

message input

$\text{out}(M, N); P$

message output

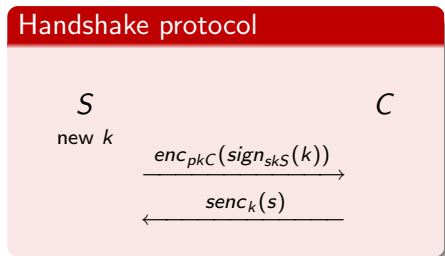
$\text{if } M = N \text{ then } P \text{ else } Q$

conditional

$\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q$

destructor application

Example: handshake protocol



C knows S 's public key. S is willing to talk to any C (does not know their public keys in advance). They want to agree a session key; they communicate on a channel that is controlled by the attacker.

Intended properties:

- 1 Secrecy: The value s is known only to C and S .
- 2 Authentication of S : if C reaches the end of the protocol with session key k , then S proposed k for use by C .
- 3 Authentication of C : if S reaches the end of the protocol and she believes she has session the key k with C , then C was indeed her interlocutor and she has session k .

Attacker model


We model a very powerful attacker, with “Dolev-Yao” capabilities:



- it completely controls the communication channels, so it is able to record, alter, delete, insert, redirect, reorder, and reuse past or current messages, and inject new messages. (The *network* is the attacker.)
- manipulate data in arbitrary ways, including applying crypto operations **provided** it has the necessary keys.
- It controls dishonest participants.

“It’s always better to assume the worst. Assume your adversaries are better than they are. Assume science and technology will soon be able to do things they cannot yet. Give yourself a margin for error. Give yourself more security than you need today.” - Bruce Schneier

Handshake protocol attack



Intended properties:

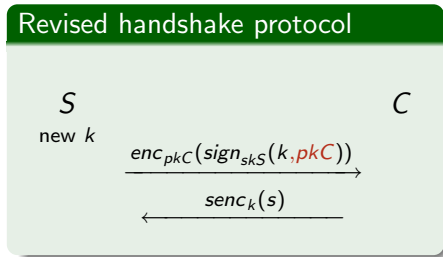
- 1 **Secrecy:** The value s is known only to C and S .
- 2 **Authentication of S :** if C reaches the end of the protocol with session key k , then S proposed k for use by C .
- 3 **Authentication of C :** if S reaches the end of the protocol and she believes she has session the key k with C , then C was indeed her interlocutor and she has session k .

Handshake protocol fixed

The attack is avoided by making the package the initiator sends include the identity of the respondent.

The three properties hold of the revised protocol, but not for the original one.


Our aim is to be able to automatically establish these facts.



Coding protocols as processes

Original handshake protocol:


```
let Server =  
  in (ch, pkC');  
  new k;  
  out (ch, enc(pkC', sign(skS, k)));  
  in (ch, m);  
  let x = sdec(k, m) in  
  0.
```



Coding protocols as processes

Original handshake protocol:

```
let Client =  
  in (ch, m);  
  let sig = dec(skC, m) in  
  if checksign(pkS, sig) = ok then  
  let k' = getmess(sig) in  
  new s;  
  out (ch, senc(k', s)).
```



The handshake protocol in full - 1

free ch.

(* Public key cryptography *)

fun pk/1.

fun enc/2. fun dec/2.

equation dec(x, enc(pk(x), y)) = y.

(* Signatures *)

fun sign/2. fun checksign/2. fun getmess/1. fun ok/0.

equation checksign(pk(x), sign(x,y)) = ok.

equation getmess(sign(x,y)) = y.

(* Shared-key cryptography *)

fun senc/2. fun sdec/2.

equation sdec(senc(x,y),x) = y.

query attacker:s.

The handshake protocol in full - 2

```
let Server =
  in (ch, pkC');
  new k;
  out (ch, enc(pkC', sign(skS, k)));
  in (ch, m);
  let x = sdec(k, m) in 0.

let Client =
  in (ch, m);
  let sig = dec(skC, m) in
  if checksign(pkS, sig) = ok then
  let k' = getmess(sig) in
  new s;
  out (ch, senc(k', s)).

process
  new skC; new skS;
  let pkC = pk(skC) in out (ch, pkC);
  let pkS = pk(skS) in out (ch, pkS);
  (!Client) | (!Server)
```

Output from ProVerif

```
Starting query not attacker:s_19[m = v_97,!1 = v_98]
goal reachable: attacker:x_297 -> attacker:s_19[m = enc(pk(skC_12[]),
    sign(skS_13[],k_21[pkC' = pk(x_297),!1 = @sid_298])),!1 = v_299]

new skC_12 creating skC_12_327 at {1}
new skS_13 creating skS_13_328 at {2}
out(ch, pk(skC_12_327)) at {4}
out(ch, pk(skS_13_328)) at {6}
in(ch, pk(a_324)) at {15}in copy a_325
new k_21 creating k_21_329 at {16}in copy a_325
out(ch, enc(pk(a_324),sign(skS_13_328,k_21_329))) at {17}in copy a_325
in(ch, enc(pk(skC_12_327),sign(skS_13_328,k_21_329))) at {8}in copy a_326
new s_19 creating s_19_330 at {12}in copy a_326
out(ch, senc(getmess(dec(skC_12_327,enc(pk(skC_12_327),
    sign(skS_13_328,k_21_329))))),s_19_330)) at {13}in copy a_326
The attacker has the message s_19_330.

A trace has been found.
RESULT not attacker:s_19[m = v_97,!1 = v_98] is false.
```

Three kinds of property

- Weak secrecy
(attacker cannot derive a value)
- Correspondence properties
(e.g. authentication properties)
- Equivalence properties
(attacker can't detect difference between instances of a system)
 - strong secrecy

Correspondence properties I

By annotating processes with events $\bar{f}\langle M \rangle$, **relationships** between the **order** of events and their **parametrisation** M can be studied.

Annotated server process

```
let Server =  
  in (ch, pkC');  
  new k;  
  event startedS(pair(pkC',k));  
  out (ch, enc(pkC',sign(skS,k)));  
  in (ch, m);  
  if pkC' = pkC then  
    event compS(k).
```

- event $\text{startedS}(\text{pair}(\text{pkC}',k))$ means *Server* started the protocol with *Client* having pub key pkC' , and k is the session key.
- event $\text{compS}(k)$ means *Server* completed the protocol with session key k .

Since event $\text{compS}(k)$ is under a conditional it can only occur when the protocol completes with *Client*.

Correspondence properties II (Handshake protocol)

```
let Server =  
  in (ch, pkC');  
  new k;  
  event startedS(pair(pkC',k));  
  out (ch, enc(pkC', sign(skS, k) ));  
  in (ch, m);  
  if pkC' = pkC then  
    event compS(k).
```

```
let Client =  
  in (ch, m);  
  let sig = dec(skC, m) in  
  if checksign(pkS, sig) = ok then  
  let k' = getmess(sig) in  
  event startedC(k');  
  out (ch, senc(k', s));  
  event compC(pair(pkC,k')).
```

Authentication properties

Authentication of server to client:

$$\text{ev} : \text{compC}(x, y) \implies \text{ev} : \text{startedS}(x, y).$$

Authentication of client to server:

$$\text{ev} : \text{compS}(y) \implies \text{ev} : \text{startedC}(y).$$

Equivalence properties

```
free ch.  
fun pk/1. fun enc/3. fun dec/2.  
equation dec(x, enc(pk(x), y, z)) = z.  
  
let Alice =  
  new n;  
  let msg = choice[enc(pk(skB), n, I_love_you), n] in  
  out (ch, msg).  
  
process  
  new skB; out (ch, pk(skB)); !Alice
```

Process calculus semantics

Applied pi calculus

Mobile values, new names
and secure communication

[Abadi/Fournet 2001]

ProVerif calculus (theory)

An efficient cryptographic
protocol verifier based on
Prolog rules

[Blanchet 2001]

Automatic verification of
correspondences for security
protocols

[Blanchet 2008]

ProVerif calculus (tool)

The ProVerif tool

Semantic configuration: E, \mathcal{P}

E : a finite set of names currently in use

\mathcal{P} : a finite multiset of closed processes

Transitions $E, \mathcal{P} \rightarrow E', \mathcal{P}'$

Semantics of ProVerif calculus

$E, \mathcal{P} \cup \{0\} \rightarrow E, \mathcal{P}$	(Red Nil)
$E, \mathcal{P} \cup \{!P\} \rightarrow E, \mathcal{P} \cup \{P, !P\}$	(Red Repl)
$E, \mathcal{P} \cup \{P \mid Q\} \rightarrow E, \mathcal{P} \cup \{P, Q\}$	(Red Par)
$E, \mathcal{P} \cup \{\nu a.P\} \rightarrow E \cup \{a'\}, \mathcal{P} \cup \{P\{a'/a\}\}$ where $a' \notin E$	(Red Res)
$E, \mathcal{P} \cup \{\overline{N}(M).Q, N(x).P\} \rightarrow E, \mathcal{P} \cup \{Q, P\{M/x\}\}$	(Red I/O)
$E, \mathcal{P} \cup \{\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{P\{M'/x\}\}$ if $g(M_1, \dots, M_n) \rightarrow M'$	(Red Destr 1)
$E, \mathcal{P} \cup \{\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{Q\}$ if there exists no M' such that $g(M_1, \dots, M_n) \rightarrow M'$	(Red Destr 2)
$E, \mathcal{P} \cup \{\text{if } M = M \text{ then } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{P\}$	(Red Cond 1)
$E, \mathcal{P} \cup \{\text{if } M = N \text{ then } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{Q\}$ if $M \neq N$	(Red Cond 2)
$E, \mathcal{P} \cup \{\text{event}(M).P\} \rightarrow E, \mathcal{P} \cup \{P\}$	(Red Event)

Predicate *attacker*(·)

The ProVerif tool performs logical reasoning (called resolution) on logic formulae (called Horn clauses) that describe knowledge obtained by the attacker.

The principal predicate in the Horn clauses is

$$\textit{attacker}(\cdot)$$

attacker:*M* means that the attacker has the means to obtain the term *M*. The attacker's means are:

- his own private computation, represented by applying functions to terms he already has;
- the “services” (or oracle) represented by protocol participants, which accept data as input and produce new data as output.

Declaring the constructors and destructors

```
fun pk/1. fun enc/3.  
reduc dec(x, enc(pk(x), y, z)) = z
```

produces the clauses

```
attacker:x -> attacker:pk(x);  
attacker:x & attacker:y & attacker:z  
  -> attacker:enc(x, y, z);  
attacker:x & attacker:enc(pk(x), y, z))  
  -> attacker:z;
```

Services offered by participants

```
let Server =  
  in (ch, x);  
  new n;  
  out (ch, enc(k, (x,n) ));
```

produces the clause

```
attacker:x → attacker:enc(k[], (x,n[x = x]));
```

Services offered by participants

```
let Client =  
  in (ch, m);  
  let sig = dec(skC, m) in  
  if checksign(pkS, sig) = ok then  
  let k' = getmess(sig) in  
  new s;  
  out (ch, senc(k', s)).
```

produces the clause

```
attacker:enc(pk(skC[]), sign(skS[], y))  
  -> attacker:senc(y,  
    s[m = enc(pk(skC[]), sign(skS[], y))]);
```

The predicate $message(\cdot, \cdot)$

$attacker:N$ means N can be made available on a public channel.

$message:M, N$ means N can be made available on a channel called M .

Inputs and outputs on private channels are encoded using $message$.

These clauses link the two:

```
message:x,y & attacker:x -> attacker:y  
attacker:x & attacker:y -> message:x,y
```

Equations vs. reduces (constructors/destructors)

When should we use equations and when should we use reduces?

AES cipher

```
equation dec(k,enc(k,m)) = m.  
equation enc(k,dec(k,m)) = m.
```

```
att:x & att:y -> att:enc(x,y)  
att:x & att:y -> att:dec(x,y)  
att:x & att:enc(x,y) -> att:y  
att:x & att:dec(x,y) -> att:y
```

AES scheme

```
reduc dec(k,enc(k,v,m)) = m.  
reduc iv(enc(k,v,m)) = v.
```

```
att:x & att:y & att:z  
-> att:enc(x,y,z)  
att:x & att:enc(x,y,z) -> att:z  
att:enc(x,y,z) -> att:y
```

Equations vs. reduces (process)

reduc/equation $\text{dec}(x, \text{enc}(x, y)) = y$.

```
new k; new s;  
in (c, x);  
let y = dec(k, x) in  
out (c, s)
```

reduc: $\text{att}:\text{enc}(k[], y) \rightarrow \text{att}:s[]$

equation: $\text{att}:\text{enc}(k[], y) \rightarrow \text{att}:s[]$

$\text{att}:x \ \& \ \text{att}:y \rightarrow \text{att}:\text{dec}(x, y)$

$\text{att}:x \rightarrow \text{att}:s[]$

Definitions

- A *rule* (or Horn clause) is a set of hypothesis facts and a conclusion fact, written $F_1, \dots, F_n \rightarrow F$.
- A fact is *unselectable* if it is of the form $\text{attacker}(x)$, for some variable x .
- A selection function $\text{sel} : \text{rules} \rightarrow \text{facts}$ satisfies:

$$\text{sel}(F_1, \dots, F_n \rightarrow F) = \begin{cases} \emptyset & \text{if } F_1, \dots, F_n \text{ are unselectable} \\ \{F_i\} & F_i \text{ is selectable,} \\ & \text{where } 1 \leq i \leq n \end{cases}$$

First phase: saturation

$\text{saturate}(\mathcal{R}_0) =$

- $\mathcal{R} \leftarrow \emptyset$.
For each $R \in \mathcal{R}_0$, $\mathcal{R} \leftarrow \text{elim}(\text{simplify}(R) \cup \mathcal{R})$.
- Repeat until a fixpoint is reached
for each $R \in \mathcal{R}$ such that $\text{sel}(R) = \emptyset$,
for each $R' \in \mathcal{R}$, for each $F_0 \in \text{sel}(R')$ such that $R \circ_{F_0} R'$ is defined,
 $\mathcal{R} \leftarrow \text{elim}(\text{simplify}(R \circ_{F_0} R') \cup \mathcal{R})$.
- Return $\{R \in \mathcal{R} \mid \text{sel}(R) = \emptyset\}$.

Second phase: backwards depth-first search

$$\text{deriv}(R, \mathcal{R}, \mathcal{R}_1) = \begin{cases} \emptyset & \text{if } \exists R' \in \mathcal{R}, R' \sqsupseteq R \\ \{R\} & \text{otherwise, if } \text{sel}(R) = \emptyset \\ \bigcup \{ \text{deriv}(\text{simplify}'(R' \circ_{F_0} R), \\ \{R\} \cup \mathcal{R}, \mathcal{R}_1) \mid R' \in \mathcal{R}_1, \\ F_0 \in \text{sel}(R) \text{ such that} \\ R' \circ_{F_0} R \text{ is defined} \} & \text{otherwise} \end{cases}$$

$$\text{derivable}(F, \mathcal{R}_1) = \text{deriv}(F \Rightarrow F, \emptyset, \mathcal{R}_1)$$

Correctness

If $attacker(M)$ cannot be derived from $clauses(P)$ then M is secret in P .

ProVerif abstractions

- Processes are treated as if they can be executed multiple times.

$$\text{Clauses}(\text{new } \tilde{n}; P) \approx \text{Clauses}(\text{new } \tilde{n}; !P)$$

- A value output on a channel is treated as if output multiple times.
- New values are not necessarily new.

$$\text{new } p; (!(\text{new } n; \text{in}(c, x); \text{out}(p, (x, n)) \mid \\ \text{in}(p, (=a, y)); \text{in}(p, (=b, =y)); \text{out}(c, \text{secret}))))$$

These abstractions are safe (i.e. they increase the power of the attacker; they do not lose attacks), but they can produce false attacks.

ProVerif tricks (non-termination)

- Put *new* before *in*(\cdot, \cdot).
- Replace *new n*
with *in(pubch, n)*
- Replace *new n; ... if $x_n \langle \rangle T$ then ...*
with *new n_0 ; let $n = \text{prop}(n_0)$ in ... if *deprop*(x_n) then ...*
where we have the reduc *deprop(prop(x)) = x*.
- Don't have queries with variables:
Replace *query attacker:M*.
with *attacker:M -> mybad;;*
query mybad:.
- Look at infinite series of clauses being produced. Consider inserting a clause that subsumes them all.
- Consider fiddling with *nounif*.

ProVerif tricks (false attacks)

General

- Put *new* after *in*(\cdot , \cdot).
- Use different private channels for different purposes. (Can't test occurrences on private channels, but can test presence.)

Equivalences

- Swapping for equivalences (see later).
 $out(c, choice[a, b]) \mid out(c, choice[b, a])$
- Skolem term depends on private data.
 $new\ p; (in(c, y); new\ n; out(p, (choice[const, y], n)))$
 $\mid (in(p, x); new\ m; out(c, m))$

Equivalence properties

```
free ch.  
fun pk/1. fun enc/3. fun dec/2.  
equation dec(x, enc(pk(x), y, z)) = z.  
  
let Alice =  
  new n;  
  let msg = choice[enc(pk(skB), n, I_love_you), n] in  
  out (ch, msg).  
  
process  
  new skB; out (ch, pk(skB)); !Alice
```

Predicate $attacker2(\cdot, \cdot)$

Predicate $attacker2:M,N$

Using the bprocess, the attacker can obtain M from the left process and N from the right process.

Simple equivalence example

Process:

```
param verboseClauses = short.

free ch, vote.

(* Public key cryptography *)
fun pk/1.
fun enc/3. fun dec/2.
equation dec(x, enc(pk(x), y, z) ) = z.

let Arole =
  new n;
  let msg = choice[ enc(pk(skB), n, vote), n ] in
  out (ch, msg).

process
  new skB; let pkB=pk(skB) in out (ch, pk(skB));
  ! Arole
```

Clauses:

```
Clause 0: att2:x,y -> att2:pk(x),pk(y)
Clause 1: att2:x,x' & att2:y,y' & att2:z,z' -> att2:enc(x,y,z),enc(x',y',z')
Clause 2: att2:x,x' & att2:y,y' -> att2:dec(x,y),dec(x',y')
Clause 3: att2:u,x & att2:v,enc(pk(x),y,z) -> att2:dec(u,v),z
Clause 4: att2:x,u & att2:enc(pk(x),y,z),v -> att2:z,dec(u,v)
Clause 5: att2:x,x' & att2:enc(pk(x),y,z),enc(pk(x'),y',z') -> att2:z,z'
Clause 13: att2:vote[],vote[]
Clause 14: att2:ch[],ch[]
Clause 15: att2:new_name[!att = v_73],new_name[!att = v_73]
Clause 16: x <> z & att2:x,y & att2:x,z -> bad:
Clause 17: x <> z & att2:y,x & att2:z,x -> bad:
Clause 18: att2:pk(skB[]),pk(skB[])
Clause 19: att2:enc(pk(skB_7[]),n_9[!1 = sid_101],vote[]),n_9[!1 = sid_101]
```

Simple non-equivalence example

Process:

```
param verboseClauses = short.  
  
free ch, a,b.  
  
let Arole = out (ch, choice[a,b]).  
let Brole = out (ch, choice[b,a]).  
  
process  
  Arole | Brole
```

Clauses:

```
Clause 7: attacker2:b[],b[]  
Clause 8: attacker2:a[],a[]  
Clause 9: attacker2:ch[],ch[]  
Clause 10: attacker2:new_name[!att = v_25],new_name[!att = v_25]  
Clause 11: x <> y & attacker2:x,z & attacker2:y,z -> bad:  
Clause 12: x <> y & attacker2:z,x & attacker2:z,y -> bad:  
Clause 13: attacker2:a[],b[]  
Clause 14: attacker2:b[],a[]
```

Outline

Applied pi calculus

Applied pi calculus: grammar

Terms

$L, M, N, T, U, V ::=$	
$a, b, c, k, m, n, s, t, r, \dots$	name
x, y, z	variable
$g(M_1, \dots, M_l)$	function

Equational theory

Suppose we have defined nullary function *ok*, unary function *pk*, binary functions *dec*, *senc*, *sdec*, *sign* and ternary functions *enc* and *checksign*.

$$\text{equation } sdec(x, senc(x, y)) = y$$

$$\text{equation } dec(x, enc(pk(x), r, y)) = y$$

$$\text{equation } checksign(pk(x), y, sign(x, y)) = ok$$

Equations to model the cryptography

1 Encryption and signatures

$$\begin{aligned} sdec(x, senc(x, y)) &= y \\ dec(x, enc(pk(x), r, y)) &= y \\ checksign(pk(x), sign(x, y)) &= ok \end{aligned}$$

2 Blind signatures

$$unblind(r, sign(x, blind(r, y))) = y$$

3 Designated verifier proof of re-encryption

The term $dvp(x, reencrypt(r, x), r, pkv)$ represents a proof designated for the owner of pkv that x and $reencrypt(x, r)$ have the same plaintext.

$$\begin{aligned} checkdvp(dvp(x, reencrypt(r, x), r, pkv), x, reencrypt(r, x), pkv) &= ok \\ checkdvp(dvp(x, y, z, skv), x, y, pk(skv)) &= ok \end{aligned}$$

4 Zero knowledge proofs of knowledge...

Applied pi calculus: grammar

Processes

$P, Q, R, \dots ::=$	processes	$A, B, C, \dots ::=$	extended processes
0	null process	P	plain process
$P \mid Q$	parallel composition	$A \mid B$	parallel composition
$!P$	replication	$\nu n.A$	name restriction
$\nu n.P$	name restriction	$\nu x.A$	variable restriction
$u(x)$	message input	$\{M/x\}$	active substitution
$\bar{u}\langle M \rangle$	message output		
$\text{if } M = N \text{ then}$ $ P \text{ else } Q$	conditional		

Example

$$\nu k.(\bar{c}\langle \text{senc}(k, a) \rangle. \bar{c}\langle \text{senc}(k, b) \rangle \mid \{h(k)/x\})$$

Machine-readable syntax

Math. syntax	Machine syntax
0	0
$P \mid Q$	P Q
$\neg P$!P
$\nu n.P$	new n; P
$u(x).P$	in(u,x); P
$\bar{u}\langle M \rangle.P$	out(u,M); P
<i>if M = N then P else Q</i>	if M = N then P else Q
$\nu x.(\{M/x\} \mid P)$	let x = M in P

Applied pi calculus: Operational semantics I

PAR-0	$A \equiv A \mid 0$
PAR-A	$A \mid (B \mid C) \equiv (A \mid B) \mid C$
PAR-C	$A \mid B \equiv B \mid A$
REPL	$!P \equiv P \mid !P$
NEW-0	$\nu n.0 \equiv 0$
NEW-C	$\nu u.\nu w.A \equiv \nu w.\nu u.A$
NEW-PAR	$A \mid \nu u.B \equiv \nu u.(A \mid B)$ where $u \notin fv(A) \cup fn(A)$
ALIAS	$\nu x.\{M/x\} \equiv 0$
SUBST	$\{M/x\} \mid A \equiv \{M/x\} \mid A\{M/x\}$
REWRITE	$\{M/x\} \equiv \{N/x\}$ where $M =_E N$

COMM $\bar{c}\langle M \rangle.P \mid c(x).Q \rightarrow P \mid Q\{M/x\}$

THEN *if* $N = N$ *then* P *else* $Q \rightarrow P$

THEN *if* $L = M$ *then* P *else* $Q \rightarrow Q$
for ground terms L, M where $L \neq_E M$

Syntactic secrecy

- Secrecy of M is preserved if an adversary cannot construct M from the outputs of the protocol.
- Formalise the adversary as a process I running in parallel. If I cannot output M , then secrecy is preserved.

Can-output

P can output the term M if there exists an evaluation context $C[_]$ and channel $c \notin \text{bn}(C)$ and process R such that the reduction $P \rightarrow^* C[\bar{c}\langle M \rangle.R]$ with no alpha-renaming of the names in $\text{fn}(M)$.

Syntactic secrecy

A closed plain process P preserves the *syntactic secrecy* of M , if there is no plain processes I where $(\text{fn}(I) \cup \text{bn}(I)) \cap \text{bn}(P) = \emptyset$ such that $P \mid I$ can output M .

Security properties

The applied pi calculus can model the following:

- Reachability properties (e.g. **secrecy**)
- Correspondence assertions (e.g. **authentication**)
- Observational equivalence (e.g. **strong secrecy**, for instance, **ballot secrecy**)

Examples:

- *Certified email* [AbadiBlanchet05]
- *Privacy* properties [DelauneKremerRyan09], and *election verifiability* properties [SmythRyanKremer10] in e-voting
- *Trusted computing protocols* [ChenRyan09,MukhamedovGordonRyan09], and *attestation* protocols [SmythRyanChen07,Backes08]
- *Web services interoperability* [BhargavanFournetGordonTse]
- *Integrity of file systems* on untrusted storage [ChaudhuriBlanchet08]

Labelled semantics: $A \xrightarrow{\alpha} B$

- $A \xrightarrow{c(M)} B$ means that the process A performs an input of the term M from the environment on the channel c , and the resulting process is B .
- $A \xrightarrow{\bar{c}\langle u \rangle} B$ means that the process A outputs the free u (which may be a variable, or a channel name).
- $A \xrightarrow{\nu u.\bar{c}\langle u \rangle} B$ means A outputs u that is restricted in A , and becomes free in B . Again, u is a channel name or a variable representing a term.

Applied pi calculus: operational semantics IV

IN $c(x).P \xrightarrow{c(M)} P\{M/x\}$

OUT-ATOM $\bar{c}\langle u \rangle.P \xrightarrow{\bar{c}\langle u \rangle} P$

OPEN-ATOM
$$\frac{A \xrightarrow{\bar{c}\langle u \rangle} A' \quad u \neq c}{\nu u.A \xrightarrow{\nu u.\bar{c}\langle u \rangle} A'}$$

SCOPE
$$\frac{A \xrightarrow{\alpha} A' \quad u \text{ does not occur in } \alpha}{\nu u.A \xrightarrow{\alpha} \nu u.A'}$$

PAR
$$\frac{A \xrightarrow{\alpha} A' \quad bv(\alpha) \cap fv(B) = bn(\alpha) \cap fn(B) = \emptyset}{A \mid B \xrightarrow{\alpha} A' \mid B}$$

STRUCT
$$\frac{A \equiv B \quad B \xrightarrow{\alpha} B' \quad B' \equiv A'}{A \xrightarrow{\alpha} A'}$$

Operational semantics: example

$$\begin{array}{c} \frac{}{\bar{c}\langle x \rangle . P \xrightarrow{\bar{c}\langle x \rangle} P} \text{OUT-ATOM} \\ \frac{}{\bar{c}\langle x \rangle . P \mid \{M/x\} \xrightarrow{\bar{c}\langle x \rangle} P \mid \{M/x\}} \text{PAR} \\ \frac{}{\bar{c}\langle M \rangle . P \equiv \nu x . (\bar{c}\langle x \rangle . P \mid \{M/x\}) \xrightarrow{\nu x . \bar{c}\langle x \rangle} P \mid \{M/x\} \equiv P \mid \{M/x\}} \text{OPEN-ATOM} \\ \frac{}{\bar{c}\langle M \rangle . P \xrightarrow{\nu x . \bar{c}\langle x \rangle} P \mid \{M/x\}} \text{STRUCT} \end{array}$$

Correspondence properties III

Correspondence property

A *correspondence property* is a formula of the form:

$$\bar{f}\langle M \rangle \rightsquigarrow \bar{g}\langle N \rangle.$$

A correspondence property asserts if event f has been executed then the event g must have been previously executed and any relationship between the event parameters must be satisfied.

Validity of correspondence property

Let E be an equational theory, and A_0 an extended process. We say that A_0 *satisfies the correspondence property* $\bar{f}\langle M \rangle \rightsquigarrow \bar{g}\langle N \rangle$ if for all execution paths

$$A_0 \xrightarrow{* \alpha_1} \rightarrow^* A_1 \xrightarrow{* \alpha_2} \rightarrow^* \dots \rightarrow^* \alpha_n \rightarrow^* A_n,$$

and all index $i \in \mathbb{N}$, substitution σ and variable e such that $\alpha_i = \nu e. \bar{f}\langle e \rangle$ and $e\varphi(A_i) =_E M\sigma$, there exists $j \in \mathbb{N}$ and e' such that $\alpha_j = \nu e'. \bar{g}\langle e' \rangle$, $e'\varphi(A_j) =_E N\sigma$ and $j < i$.

Equivalence properties

Equivalence defines indistinguishability between two processes and allows us to consider properties that cannot be expressed as secrecy or correspondence properties.

Example: electronic voting

We cannot model vote privacy as syntactic secrecy, because

- the identities of the voters are not secret;
- the votes (names of candidates) are not secret.

What is secret is the link between the voter and the vote.

Privacy in electronic voting

A voting protocol respects **privacy** if

$$S[V_A\{a/v\} \mid V_B\{b/v\}] \approx S[V_A\{b/v\} \mid V_B\{a/v\}].$$

Observational equivalence

We write $A \Downarrow c$ when A can evolve to a process that can send a message on c , that is, when $A \rightarrow^* C[\bar{c}\langle M \rangle.P]$ for some term M and some evaluation context $C[-]$ that does not bind c .

Observational equivalence

Observational equivalence (\approx) is the largest symmetric relation \mathcal{R} between closed extended processes with the same domain such that $A \mathcal{R} B$ implies:

- 1 if $A \Downarrow c$, then $B \Downarrow c$.
- 2 if $A \rightarrow^* A'$ then, for some B' , we have $B \rightarrow^* B'$ and $A' \mathcal{R} B'$;
- 3 $C[A] \mathcal{R} C[B]$ for all closing evaluation contexts $C[-]$.

The definition universally quantifies over evaluation contexts to capture all possible adversary behaviour. This makes the definition of observational equivalence hard to use in practice.

Labelled bisimilarity I

Labelled bisimilarity is **more suitable for reasoning**. It relies on an equivalence relation between frames; intuitively, two frames are statically equivalent if no 'test' $M = N$ can tell them apart

Static equivalence

Two closed frames $\varphi \equiv \nu \tilde{m}.\sigma$ and $\psi \equiv \nu \tilde{n}.\tau$ are statically equivalent, denoted $\varphi \approx_s \psi$, if $\text{dom}(\varphi) = \text{dom}(\psi)$ and for all terms M, N such that $(\tilde{m} \cup \tilde{n}) \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$, we have $M\sigma =_E N\sigma$ holds if and only if $M\tau =_E N\tau$ holds.

Examples

- $\nu m.\{m/x\} \approx_s \nu n.\{n/x\}$; they are structurally equivalent.
- $\nu m.\{m/x\} \approx_s \nu n.\{\text{hash}(n)/x\}$.
- $\{m/x\} \not\approx_s \{\text{hash}(m)/x\}$. LHS satisfies $x = m$.
- $\nu s.\{\text{pair}(s, s)/x\} \not\approx_s \nu s.\{s/x\}$.
LHS satisfies $\text{pair}(\text{fst}(x), \text{snd}(x)) = x$.

Labelled bisimilarity II

Static equivalence examines the current state of the processes (as represented by their frames), and not the processes' dynamic behaviour (that is, the ways in which they may execute in the future). The dynamic part is captured as follows.

Labelled bisimilarity

Labelled bisimilarity (\approx_l) is the largest symmetric relation \mathcal{R} on closed extended processes such that $A \mathcal{R} B$ implies:

- 1 $A \approx_s B$;
- 2 if $A \rightarrow A'$ then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ;
- 3 if $A \xrightarrow{\alpha} A'$ and $\text{fv}(\alpha) \subseteq \text{dom}(A)$ and $\text{bn}(\alpha) \cap \text{fn}(B) = \emptyset$; then $B \rightarrow^* \xrightarrow{\alpha} \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' .

Abadi & Fournet state that observational equivalence and labelled bisimilarity coincide.