

Verifying Security Protocols in Tamarin

Ralf Sasse
Institute of Information Security
ETH Zurich

Tamarin Day 4, v.1
Jan 28, 2016

Roadmap

- 1 Observational Equivalence: High-Level Overview**
- 2 Observational Equivalence Details**
- 3 Case Studies**

Outline

- 1 Observational Equivalence: High-Level Overview**
- 2 Observational Equivalence Details
- 3 Case Studies

Motivation

- Security protocol design is **critical** and **error-prone**
- **Symbolic analysis** methods make a difference

Two types of properties:

- **Trace properties:**
 - ★ (Weak) secrecy as reachability
 - ★ Authentication as correspondence

- **Observational equivalence**



Why observational equivalence?

Consider classic **Dolev-Yao** intruder for deterministic public-key encryption:

$$\frac{enc(x, pk(k)) \quad k}{x}$$

- Intruder can only decrypt if he knows the secret key

Why observational equivalence?

Consider classic **Dolev-Yao** intruder for deterministic public-key encryption:

$$\frac{enc(x, pk(k)) \quad k}{x}$$

- Intruder can only decrypt if he knows the secret key

Now consider a simple **voting** system:

- Voter chooses $v = \text{"Yes"}$ or $v = \text{"No"}$
- Encrypt v using server's public key $pk(k)$: $c = enc(v, pk(k))$
- Send c to server

Why observational equivalence?

Consider classic **Dolev-Yao** intruder for deterministic public-key encryption:

$$\frac{enc(x, pk(k)) \quad k}{x}$$

- Intruder can only decrypt if he knows the secret key

Now consider a simple **voting** system:

- Voter chooses $v = \text{"Yes"}$ or $v = \text{"No"}$
- Encrypt v using server's public key $pk(k)$: $c = enc(v, pk(k))$
- Send c to server

Is the vote **secret**?

- Dolev-Yao: **Yes**, intruder does not know server's secret key

Why observational equivalence?

Consider classic **Dolev-Yao** intruder for deterministic public-key encryption:

$$\frac{enc(x, pk(k)) \quad k}{x}$$

- Intruder can only decrypt if he knows the secret key

Now consider a simple **voting** system:

- Voter chooses $v = \text{"Yes"}$ or $v = \text{"No"}$
- Encrypt v using server's public key $pk(k)$: $c = enc(v, pk(k))$
- Send c to server

Is the vote **secret**?

- Dolev-Yao: **Yes**, intruder does not know server's secret key
- Reality: **No**, encryption is deterministic and there are only two choices
 - ★ **Attack**: encrypt "Yes", and compare to c

Observational Equivalence vs Reachability

- **Reachability**-based (weak) secrecy is insufficient
- **Stronger** notion: intruder cannot distinguish
 - ★ a system where the voter votes “Yes” from
 - ★ a system where the voter votes “No”

Observational Equivalence vs Reachability

- **Reachability**-based (weak) secrecy is insufficient
- **Stronger** notion: intruder cannot distinguish
 - ★ a system where the voter votes “Yes” from
 - ★ a system where the voter votes “No”
- **Observational equivalence** between two systems

Observational Equivalence vs Reachability

- **Reachability**-based (weak) secrecy is insufficient
- **Stronger** notion: intruder cannot distinguish
 - ★ a system where the voter votes “Yes” from
 - ★ a system where the voter votes “No”
- **Observational equivalence** between two systems
- Can be used to express:
 - ★ Strong secrecy
 - ★ Privacy notions
 - ★ Game-based notions, e.g., ciphertext indistinguishability

Running Example

- **Auction** system
- Property: **strong secrecy** of bids

Running Example

- **Auction** system
- Property: **strong secrecy** of bids
- Property **violated**: **Shout-out auction**
 - ★ Broadcast bid (e.g., A or B)
 - ★ Send “A” in first system
 - ★ Send “B” in second system
 - ★ Observer knows if he is observing first or second system

Running Example

- **Auction system**
- Property: **strong secrecy** of bids
- Property **violated**: **Shout-out auction**
 - ★ Broadcast bid (e.g., A or B)
 - ★ Send “A” in first system
 - ★ Send “B” in second system
 - ★ Observer knows if he is observing first or second system
- Property **holds**: using **shared symmetric key**
 - ★ Shared symmetric key k between bidder and auctioneer
 - ★ Send “ $\{A\}_k$ ” in first system
 - ★ Send “ $\{B\}_k$ ” in second system
 - ★ Observer has no access to k , does not know which system he is observing

Symbolic Model for Observational Equivalence

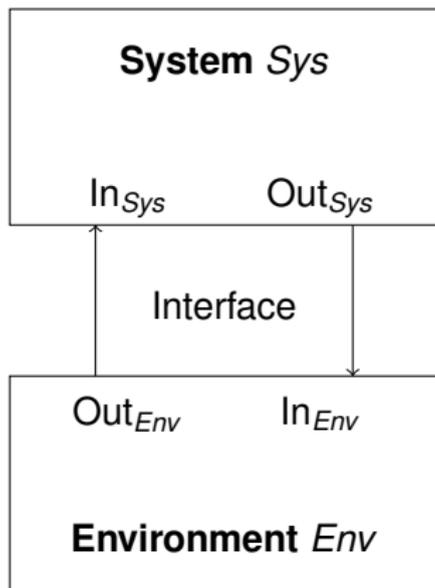
- **Symbolic abstractions** of cryptographic operators
- Enables **high** degree of **automation** for observational equivalence, existing tools:
 - ★ APTE, AKISS, ProVerif, SPEC, Maude-NPA
- Limitations: **No tool** can prove observational equivalence for protocols with **combination** of
 - ★ Mutable state
 - ★ Unbounded sessions
 - ★ Diffie-Hellman
- Examples of applications beyond **state-of-the-art**:
 - ★ Protocols using hardware security modules
 - ★ Key exchange protocols based on Diffie-Hellman

Contribution: Tamarin

- Novel definition of observational equivalence
 - ★ **General** definition within context of multiset rewriting setting
- **Sound algorithm**, first to support all of:
 - ★ Observational equivalence
 - ★ Mutable state
 - ★ Unbounded number of sessions
 - ★ Diffie-Hellman exponentiation
- Implementation, **practicality** demonstrated by
 - ★ non-trivial protocols
- **Largely automated** proofs
- Approach is **not complete**, but **succeeds broadly**

System and environment

- We separate **environment** and **system**
 - ★ System: agents running according to protocol
 - ★ Environment: adversary acting according to its capabilities
- Environment can observe:
 - ★ Output of the system
 - ★ If system reacts at all



Defining observational equivalence

- **Two** system **specifications** given as set of rules
 - ★ One rule per role action (send/receive)
 - ★ Running example shout-out auction:

System 1: $\overline{\text{Out}_{\text{Sys}}(A)}$

System 2: $\overline{\text{Out}_{\text{Sys}}(B)}$

- Interface and environment/adversary rule(s):

$$\frac{\text{Out}_{\text{Sys}}(X)}{\text{In}_{\text{Env}}(X)} \qquad \frac{\text{Out}_{\text{Env}}(X)}{\text{In}_{\text{Sys}}(X)} \qquad \frac{\text{In}_{\text{Env}}(X) \quad K(X)}{\text{Out}_{\text{Env}}(\text{true})}$$

- ★ $K(X)$ represents that environment knows term X
 - ★ last rule models comparisons by the adversary
- Each specification yields a labeled transition system
- Observational equivalence is a kind of **bisimulation** accounting for the adversaries' **viewpoint** and **capabilities**
 - ★ Our definition can be instantiated for various adversaries

Diff terms

- General definition **difficult** to verify: requires inventing simulation relation
- Idea: **specialize** for cryptographic protocols
 - ★ Consider strong bid secrecy:
 - ▶ both systems differ in **secret bid only**, i.e.
 - ▶ both specifications contain **same rule(s)** which differ only in **some terms**
 - ★ Exploit this similarity in description and proof
- Approach: two systems described by one **specification** – using **diff-terms**

Diff terms

- General definition **difficult** to verify: requires inventing simulation relation
- Idea: **specialize** for cryptographic protocols
 - ★ Consider strong bid secrecy:
 - ▶ both systems differ in **secret bid only**, i.e.
 - ▶ both specifications contain **same rule(s)** which differ only in **some terms**
 - ★ Exploit this similarity in description and proof
- Approach: two systems described by one **specification** – using **diff-terms**
 - ★ Running example

$$\overline{\text{Out}_{\text{Sys}}(A)}$$

$$\overline{\text{Out}_{\text{Sys}}(B)}$$

- ★ Is equivalent to one rule with a **diff-term**

$$\overline{\text{Out}_{\text{Sys}}(\mathbf{diff}(A, B))}$$

Approximating observational equivalence using mirroring

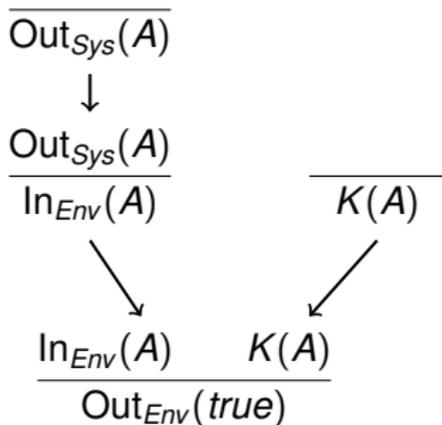
- Both systems contain the same rules modulo diff-terms
- Idea: assume that each rule simulates itself
- **Mirrors** each execution into the other system
- If the mirrors are **valid executions**, we have **observational equivalence** (sound approximation)
- We represent executions using **dependency graphs**
 - ★ Computed via backwards constraint solving

Dependency graphs and mirrors

Bidder picks A, observer compares to public value A

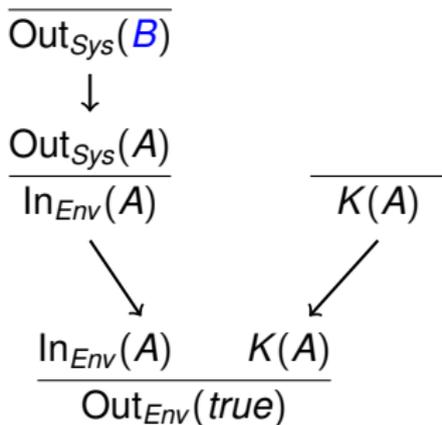
Dependency graphs and mirrors

Bidder picks A, observer compares to public value A



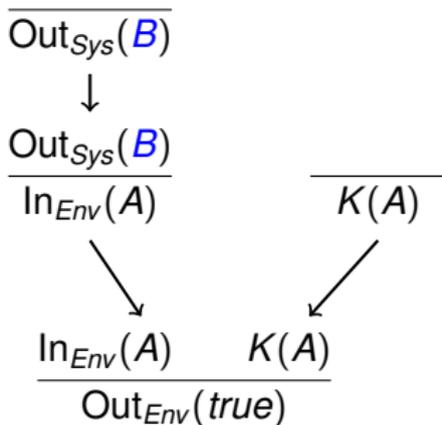
Dependency graphs and mirrors

Bidder picks **B**, observer compares to public value A



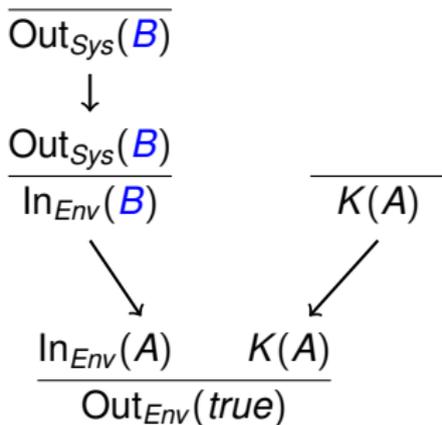
Dependency graphs and mirrors

Bidder picks B , observer compares to public value A



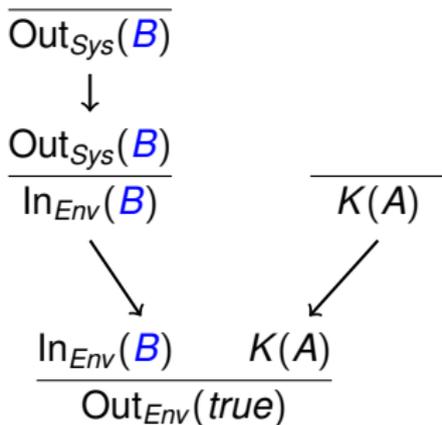
Dependency graphs and mirrors

Bidder picks B , observer compares to public value A



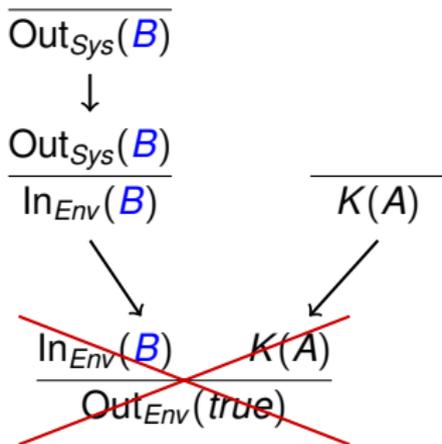
Dependency graphs and mirrors

Bidder picks B , observer compares to public value A



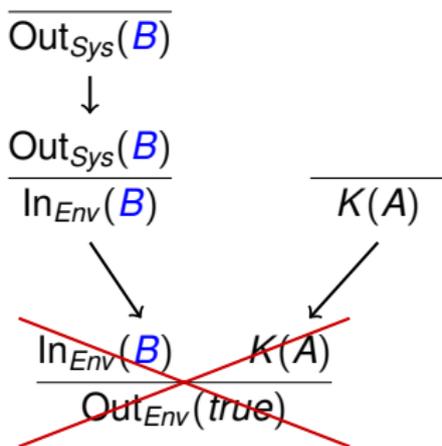
Dependency graphs and mirrors

Bidder picks B , observer compares to public value A



Dependency graphs and mirrors

Bidder picks B , observer compares to public value A



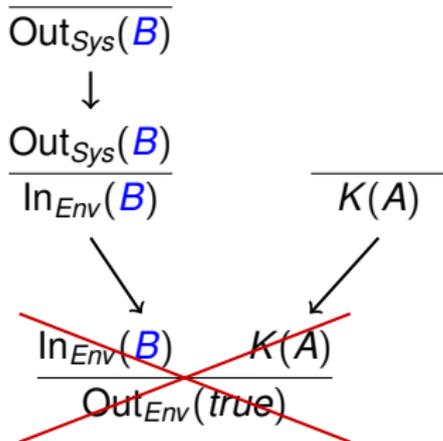
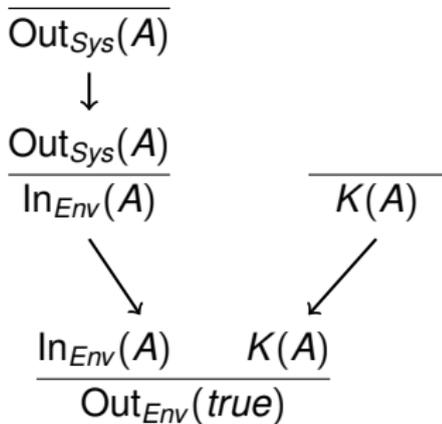
- Dependency graph mirror for bidder choice B is **invalid**
 - ★ Adversary choices stay fixed, comparison is with A

Invalid mirrors and attacks

Bidder picks A/B , observer compares to public value A

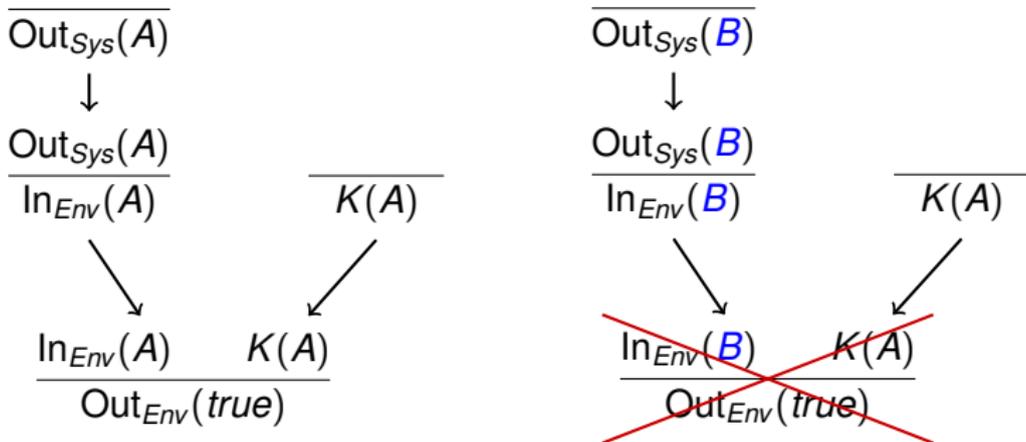
Invalid mirrors and attacks

Bidder picks A/B, observer compares to public value A



Invalid mirrors and attacks

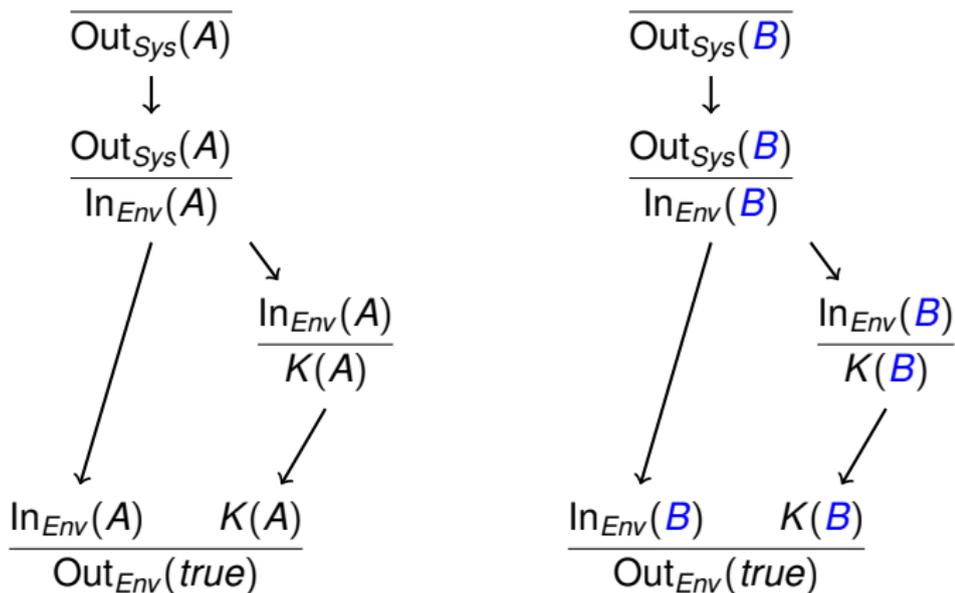
Bidder picks A/B, observer compares to public value A



- **Counter example** to observational equivalence of the given systems

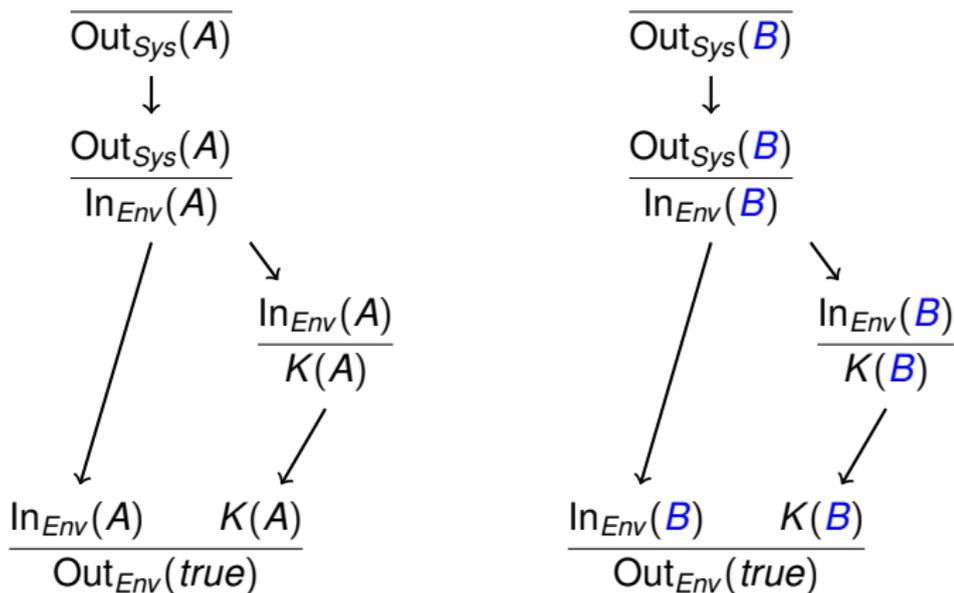
Valid mirror

Observer compares system output to itself



Valid mirror

Observer compares system output to itself



- **All** mirrors need to be valid for observational equivalence

Dependency graph equivalence

A **diff**-system is **dependency graph equivalent** if mirrors of all dependency graphs rooted in any rule on both sides are valid.

- Sound but incomplete approximation
- Efficient and sufficient in practice

Dependency graph equivalence

A **diff**-system is **dependency graph equivalent** if mirrors of all dependency graphs rooted in any rule on both sides are valid.

- Sound but incomplete approximation
- Efficient and sufficient in practice

Input:

- Protocol specification
- Property: equivalence given two choices for some term(s)
 - ★ Example: random value vs expected value

Dependency graph equivalence

A **diff**-system is **dependency graph equivalent** if mirrors of all dependency graphs rooted in any rule on both sides are valid.

- Sound but incomplete approximation
- Efficient and sufficient in practice

Input:

- Protocol specification
- Property: equivalence given two choices for some term(s)
 - ★ Example: random value vs expected value

Output:

- **Yes**, observational equivalent
- **No**, dependency graph with invalid mirror
- Non-termination possible

Tamarin advantages

Approach implemented in the Tamarin tool:

- Tamarin supports **verification** with:
 - ★ equational theories (DH), induction, loops, mutable state
- Security protocol model is based on rewriting
- Restricted FOL for security properties
- Equational theories modeling algebraic properties
- Constraint-solving algorithm for analysis of unbounded number of sessions
- Performance good despite undecidability

Tamarin extension for observational equivalence

Implemented algorithm:

- Extended constraint solving
- (**Normal**) dependency graphs
 - ★ Important for state space reduction and termination
- Equivalence of dependency graphs by mirroring

Some Simple Examples

- Indistinguishability of **probabilistic encryption**
 - ★ Adversary cannot distinguish random value from encryption
 - ★ Automatically verified in 0.2 seconds

- **Decisional Diffie-Hellman**
 - ★ Given algebraic properties of DH exponentiation as equational theory
 - ★ Adversary cannot distinguish g^{ab} from random g^c
 - ▶ Given g^a and g^b
 - ★ Automatically verified in 15.2 seconds

Case studies

- **Feldhofer's RFID** protocol
 - ★ Adversary cannot determine which RFID tag is communicating with reader
 - ★ Automatically verified in 1.6 seconds
- **Signed Diffie-Hellman** key exchange
 - ★ Real-or-random secrecy of session key
 - ★ Needs manual guidance in one subcase
 - ★ Automatically completed proof in 2.5 minutes
- **TPM_Envelope**
 - ★ Real-or-random secrecy
 - ★ Finds attack for deterministic encryption
 - ▶ Despite previous proof wrt trace-based secrecy
 - ★ We recommend to use probabilistic encryption

Future Work

- Implement **more precise approximation**
 - ★ Currently rules must match 1–1 due to mirroring
 - ★ General definition does allow matching using different, even multiple, rules

- Protocols with loops: need induction

- Further case studies
 - ★ Signed Diffie-Hellman with Perfect Forward Secrecy
 - ★ NAXOS, authenticated key exchange with PFS

Related Work

	APTE	AKISS	ProVerif	ProVerifDH	SPEC	Maude-NPA	Tamarin	Extension
Unbounded sessions			x	x		x	x	x
Mutable state	x	x			x	?	x	x
Diffie-Hellman	x	x		x	x	x	x	x
Definable crypto	x	x	x	x		x	x	x
Verification	x	x	x	x	x	x	x	x
Obs. equiv.	x	x	x		x	/		x

- APTE, AKISS
 - ★ Limited to bounded number of sessions
- ProVerif
 - ★ No mutable state support
 - ★ DH support only without observational equivalence
- SPEC
 - ★ Fixed crypto primitives, bounded number of sessions
- StatVerif, SAPIC
 - ★ Support mutable state, but no observational equivalence
- Maude-NPA
 - ★ Creates synchronous product of two similar protocols
 - ★ Suffers from termination issues - only finds attacks

Summary

- Extended multiset rewriting approach to observational equivalence
- Result is well-suited for cryptographic protocol analysis
- Shown algorithm's **effectiveness and scope** on case studies
- **High** degree of **automation**

Outline

- 1 Observational Equivalence: High-Level Overview
- 2 Observational Equivalence Details**
- 3 Case Studies

Modeling

Define functions $one(\cdot)$ and $two(\cdot)$, and a fact symbol M .

$$Env = \{ \begin{array}{l} E_{null} : \text{---} \rightarrow M(null), \\ E_{one} : M(x) \text{---} \rightarrow M(one(x)), \\ E_{two} : M(x) \text{---} \rightarrow M(two(x)), \\ E_{check} : M(x), \text{In}_{Env}(x) \text{---} \rightarrow \text{Out}_{Env}(true) \end{array} \}$$

Using the final rule, E_{check} , the system can compare a constructed term with the value stored in the $\text{In}_{Env}(\cdot)$ fact.

Modeling Example

Each fact is associated with a **recipe** of its derivation

- Order the premises and conclusions of a rule: $id : l \rightarrow [a] \rightarrow r$ as $seq^{\leq}(l)$ and $seq^{\leq}(r)$
- For fact $F \in r$, where k is the index of F in $seq^{\leq}(r)$ we get

$$recipe(F) = id_k(newvars(F), [recipe(l_1), \dots, recipe(l_n)])$$

- Abusing notation, for a rule with name id we get:

$$recipe(id) = id([newvars(r_1), \dots, newvars(r_m)], [recipe(l_1), \dots, recipe(l_n)])$$

Semantics

The semantics of a set of multiset rewrite rules P are given by a **labeled transition relation** $\rightarrow_P \subseteq \mathcal{G}^\# \times (\mathcal{G}^\# \times \rho) \times \mathcal{G}^\#$, defined by the transition rule:

$$\frac{ri = id : l \rightarrow [a] \rightarrow r \in_E \text{ginsts}(P) \quad lfacts(l) \subseteq^\# S \quad pfacts(l) \subseteq S}{S \xrightarrow[\text{recipe}(id)]{\text{set}(a)}_P ((S \setminus^\# lfacts(l)) \cup^\# mset(r))}$$

Example: Pairs

Two systems, where the first system outputs a pair of identical values

$$S_A = \{ A : -[\] \rightarrow \text{Out}_{\text{Sys}}((x, x)) \}$$

and the second system may output two different values

$$S_B = \{ B : -[\] \rightarrow \text{Out}_{\text{Sys}}((x, y)) \}.$$

Example: Pairs

Two systems, where the first system outputs a pair of identical values

$$S_A = \{ A : -[] \rightarrow \text{Out}_{\text{Sys}}((x, x)) \}$$

and the second system may output two different values

$$S_B = \{ B : -[] \rightarrow \text{Out}_{\text{Sys}}((x, y)) \}.$$

In S_A , we have that

$$\emptyset \xrightarrow{A(\{\{m/x\}, \emptyset)} \{ \text{Out}_{\text{Sys}}((m, m)) \}.$$

Example: Pairs

Two systems, where the first system outputs a pair of identical values

$$S_A = \{ A : \text{---} \rightarrow \text{Out}_{\text{Sys}}((x, x)) \}$$

and the second system may output two different values

$$S_B = \{ B : \text{---} \rightarrow \text{Out}_{\text{Sys}}((x, y)) \}.$$

In S_A , we have that

$$\emptyset \xrightarrow{A(\{\{m/x\}, \emptyset\})} \{\text{Out}_{\text{Sys}}((m, m))\}.$$

In S_B , we can either take a similar transition

$$\emptyset \xrightarrow{B(\{\{m/x, m/y\}, \emptyset\})} \{\text{Out}_{\text{Sys}}((m, m))\}$$

or alternatively

Example: Pairs

Two systems, where the first system outputs a pair of identical values

$$S_A = \{ A : -[\] \rightarrow \text{Out}_{\text{Sys}}((x, x)) \}$$

and the second system may output two different values

$$S_B = \{ B : -[\] \rightarrow \text{Out}_{\text{Sys}}((x, y)) \}.$$

In S_A , we have that

$$\emptyset \xrightarrow{A(\{\{m/x\}\}, \emptyset)} \{\text{Out}_{\text{Sys}}((m, m))\}.$$

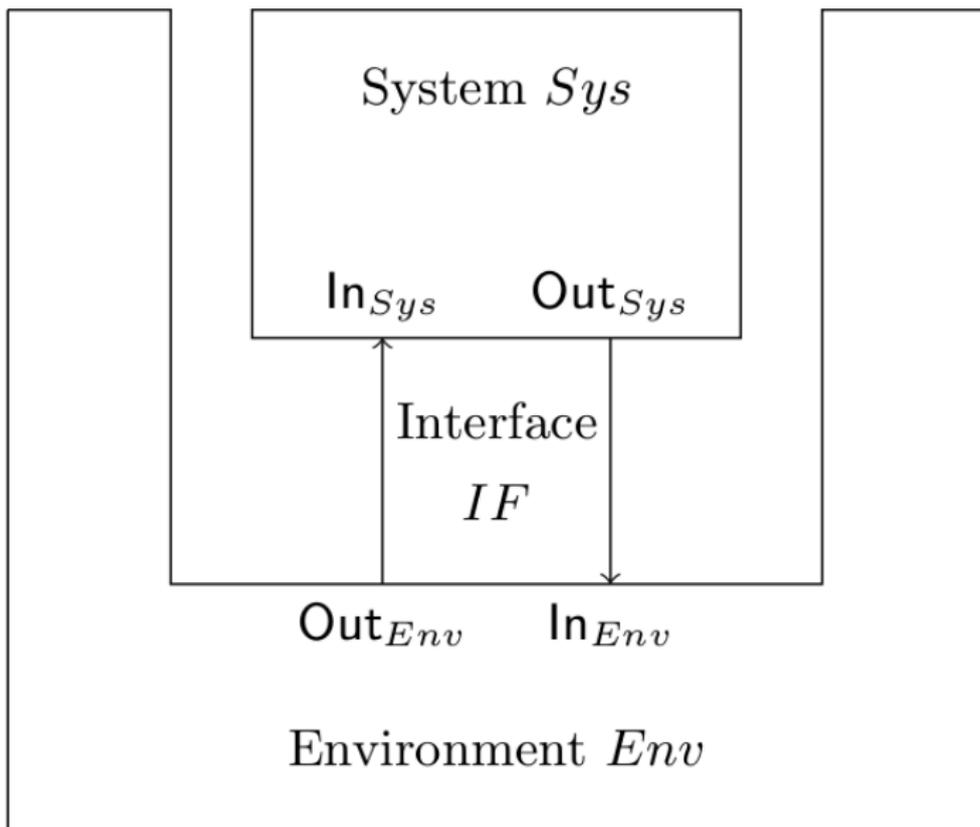
In S_B , we can either take a similar transition

$$\emptyset \xrightarrow{B(\{\{m/x, m/y\}\}, \emptyset)} \{\text{Out}_{\text{Sys}}((m, m))\}$$

or alternatively

$$\emptyset \xrightarrow{B(\{\{m/x, n/y\}\}, \emptyset)} \{\text{Out}_{\text{Sys}}((m, n))\}.$$

System Model



Interface rules

$$\begin{aligned} OUT &= \{ OUT : \text{Out}_{\text{Sys}}(M) \dashv [O] \rightarrow \text{In}_{\text{Env}}(M) \} \\ IN &= \{ IN : \text{Out}_{\text{Env}}(M) \dashv [I] \rightarrow \text{In}_{\text{Sys}}(M) \} \\ IF &= OUT \cup IN \end{aligned}$$

Special recipe for In_{Env}

- Environment is not supposed to see the internal system choices.
- Reflected by defining recipe of $\text{In}_{Env}(M)$ fact differently:

$$\text{recipe}(\text{In}_{Env}(M)) = \text{OUT}_1([], x)$$

where x is a new variable.

- The recipe associated to the rule is: $\text{recipe}(\text{OUT}) = \text{OUT}([], x)$.
- Replaces the $\text{Out}_{Sys}(M)$ fact's recipe with a variable, as the recipe is considered internal to the system.

Pairs and recipes

$$Env = \{ C : \text{In}_{Env}(x, x) \dashv\vdash \text{Out}_{Env}(true) \}.$$

Then in $S_A \cup IF \cup Env$ we have

$$\begin{aligned} \emptyset & \xrightarrow{A(\{\{m/x\}, \emptyset\})} \{\text{Out}_{Sys}((m, m))\} \\ & \xrightarrow{O} \{\text{In}_{Env}((m, m))\} \\ & \xrightarrow{C(\emptyset, [\text{OUT}_1(\emptyset, z)])} \{\text{Out}_{Env}(true)\}. \end{aligned}$$

In $S_B \cup IF \cup Env$ similarly:

$$\begin{aligned} \emptyset & \xrightarrow{B(\{\{m/x, m/y\}, \emptyset\})} \{\text{Out}_{Sys}((m, m))\} \\ & \xrightarrow{O} \{\text{In}_{Env}((m, m))\} \\ & \xrightarrow{C(\emptyset, [\text{OUT}_1(\emptyset, z)])} \{\text{Out}_{Env}(true)\}. \end{aligned}$$

To adversary these look the same.

Pairs and recipes

However, in $S_B \cup IF \cup Env$ we also have the following transitions:

$$\emptyset \xrightarrow{B(\{\{m/x, n/y\}\}, \emptyset)} \{\text{Out}_{\text{Sys}}((m, n))\}$$

$$\xrightarrow[OUT(\emptyset, z)]{O} \{\text{In}_{\text{Env}}((m, n))\}.$$

Note that the first transition has a different recipe as we instantiate y differently, but again the output replaces this recipe with a new variable z , which is the same on both sides.

The two system will turn out to be not observational equivalent shortly.

High-level definition

S_A and S_B are observational equivalent wrt interface and environment, given empty initial states of the relation:

- Interface and environment rules must make the same choices on both sides, with same recipe.
- Application of system rules S_A (resp. S_B) must be matchable by some number of system rule applications of S_B (resp. S_A), with some recipes, and end in states in the relation.

Observational equivalence - definition

Two sets of multiset rewrite rules S_A and S_B are **observational equivalent** with respect to an environment Env (and interface IF) if there is a **relation** between the initial states in $S_A \cup IF \cup Env$ (**left system**) and $S_B \cup IF \cup Env$ (**right system**), and for all pairs of states in that relation:

- If the **left system** can make a **move** with an environment or interface rule, the **right system** can **match it precisely**
 - ★ Resulting states are in the relation
- If the **left system** can make a **move** with an S_A rule, the **right system** can **match it**, possibly using **multiple steps**
 - ★ resulting states are in the relation

The same holds in the other direction.

Algorithm

- 1: **function** VERIFY(S)
- 2: $RU \leftarrow L(S) \cup R(S) \cup IF \cup Env$

Algorithm

```
1: function VERIFY(S)
2:    $RU \leftarrow L(S) \cup R(S) \cup IF \cup Env$ 
3:   while  $RU \neq \emptyset$  do
4:     choose  $r \in RU$ ,  $RU \leftarrow (RU \setminus \{r\})$ 
```

Algorithm

```
1: function VERIFY(S)
2:    $RU \leftarrow L(S) \cup R(S) \cup IF \cup Env$ 
3:   while  $RU \neq \emptyset$  do
4:     choose  $r \in RU$ ,  $RU \leftarrow (RU \setminus \{r\})$ 
5:     compute  $DG \leftarrow dgraphs(r)$  by constraint solving
```

Algorithm

```
1: function VERIFY(S)
2:    $RU \leftarrow L(S) \cup R(S) \cup IF \cup Env$ 
3:   while  $RU \neq \emptyset$  do
4:     choose  $r \in RU$ ,  $RU \leftarrow (RU \setminus \{r\})$ 
5:     compute  $DG \leftarrow dgraphs(r)$  by constraint solving
6:     if  $\exists dg \in DG$  s.t. mirrors(dg) lacks ground instances
7:       then return "potential attack found: ", dg
```

Algorithm

```
1: function VERIFY(S)
2:    $RU \leftarrow L(S) \cup R(S) \cup IF \cup Env$ 
3:   while  $RU \neq \emptyset$  do
4:     choose  $r \in RU$ ,  $RU \leftarrow (RU \setminus \{r\})$ 
5:     compute  $DG \leftarrow dgraphs(r)$  by constraint solving
6:     if  $\exists dg \in DG$  s.t. mirrors(dg) lacks ground instances
7:       then return "potential attack found: ", dg
8:   return "verification successful"
```

Example: Pairs

In $S_B \cup IF \cup Env$ we have

$$\emptyset \xrightarrow{B(\{\{m/x, n/y\}, \emptyset\})} \{\text{Out}_{\text{Sys}}((m, n))\}$$

$$\xrightarrow{O} \{\text{In}_{\text{Env}}((m, n))\}.$$

$$\text{OUT}(\emptyset, z)$$

The only way for $S_A \cup IF \cup Env$ to simulate this would be

$$\emptyset \xrightarrow{A(\{\{m/x\}, \emptyset\})} \{\text{Out}_{\text{Sys}}((m, m))\}$$

$$\xrightarrow{O} \{\text{In}_{\text{Env}}((m, m))\}.$$

$$\text{OUT}(\emptyset, z)$$

and further rule A transitions, which add more $\text{Out}_{\text{Sys}}((m, m))$ to S .
Output rule can only be used once, so just one $\text{In}_{\text{Env}}((m, m))$ in S . Then

$$S \xrightarrow{C(\emptyset, [\text{OUT}_1(\emptyset, z)])} \{\text{Out}_{\text{Env}}(\text{true})\},$$

but in state $\{\text{In}_{\text{Env}}((m, n))\}$ no such transition is possible, so $S_A \not\approx_{\text{Env}} S_B$.

Example result

The previous example showed that if a rule is applicable by the environment on one side, and not on the other, the two sides are distinguishable. Thus the two systems are not observationally equivalent.

Next we see an example where the recipe used by the environment is crucial for distinction.

Example: different Environment

$$S_A = \{ A : -[] \rightarrow \text{Out}_{\text{Sys}}((x, x)) \}$$

$$S_B = \{ B : -[] \rightarrow \text{Out}_{\text{Sys}}((x, y)) \}$$

Environment Env' with **persistent** $M(\cdot)$:

$$Env' = \{ \begin{array}{l} E_{fst} : \text{In}_{Env}((x, y)) \rightarrow M(x), \\ E_{snd} : \text{In}_{Env}((x, y)) \rightarrow M(y), \\ E_{cmp} : M(x), M(x) \rightarrow \text{Out}_{Env}(true) \end{array} \},$$

We expect this environment to be able to distinguish S_A and S_B , but need to reason differently than before.

Example continued

Let us start with $S_B \cup IF \cup Env'$:

$$\begin{array}{l}
 \emptyset \xrightarrow{\quad} \{\text{Out}_{\text{Sys}}((m, n))\} \\
 \quad \quad B(\{\{m/x, n/y\}\}, \emptyset) \\
 \xrightarrow{O} \{\text{In}_{\text{Env}}((m, n))\} \\
 \quad \quad \text{OUT}(\emptyset, z) \\
 \xrightarrow{\quad} \{\text{M}(m)\} \\
 \quad \quad E_{\text{fst}}(\emptyset, [\text{OUT}_1(\emptyset, z)]) \\
 \xrightarrow{\quad} \{\text{M}(m), \text{M}(n)\} . \\
 \quad \quad E_{\text{snd}}(\emptyset, [\text{OUT}_1(\emptyset, z)])
 \end{array}$$

In $S_A \cup IF \cup Env'$ this can be simulated as follows:

$$\begin{array}{l}
 \emptyset \xrightarrow{\quad} \{\text{Out}_{\text{Sys}}((m, m))\} \\
 \quad \quad A(\{\{m/x\}\}, \emptyset) \\
 \xrightarrow{O} \{\text{In}_{\text{Env}}((m, m))\} \\
 \quad \quad \text{OUT}(\emptyset, z) \\
 \xrightarrow{\quad} \{\text{M}(m)\} \\
 \quad \quad E_{\text{fst}}(\emptyset, [\text{OUT}_1(\emptyset, z)]) \\
 \xrightarrow{\quad} \{\text{M}(m), \text{M}(m)\} . \\
 \quad \quad E_{\text{snd}}(\emptyset, [\text{OUT}_1(\emptyset, z)])
 \end{array}$$

Example continued

Moreover, we can compare the first and second value of the tuple with

$$\{M(m), M(m)\} \xrightarrow{r_1} \{M(m), M(m), \text{Out}_{Env}(true)\},$$

where

$$r_1 = E_{cmp}([], [E_{fst,1}([], [OUT_1([], z)]), E_{snd,1}([], [OUT_1([], z)])])$$

This transition cannot be matched by $S_B \cup IF \cup Env'$. Note however that the following transition **is** possible for $S_B \cup IF \cup Env'$:

$$\{M(m), M(n)\} \xrightarrow{r_2} \{M(m), M(n), \text{Out}_{Env}(true)\},$$

where

$$r_2 = E_{cmp}([], [E_{fst,1}([], [OUT_1([], z)]), E_{fst,1}([], [OUT_1([], z)])]).$$

The recipes are crucial, as otherwise we could compare the first value with itself and it would seem these are observationally equivalent.

Example: coin returning vending machine

Now we add equational reasoning to the mix. We use 1 € and 2 € coins. Represent these by functions *one* and *two*, with constant *null*. How to return 3 €?

Example: coin returning vending machine

Now we add equational reasoning to the mix. We use 1 € and 2 € coins. Represent these by functions *one* and *two*, with constant *null*. How to return 3 €? Which coins and which order?

Example: coin returning vending machine

Now we add equational reasoning to the mix. We use 1 € and 2 € coins. Represent these by functions *one* and *two*, with constant *null*. How to return 3 €? Which coins and which order?

Specification:

$$S_A = \{ A : \text{---} \rightarrow \text{Out}_{\text{Sys}}(\text{two}(\text{one}(\text{null}))) \}.$$

Example: coin returning vending machine

Now we add equational reasoning to the mix. We use 1 € and 2 € coins. Represent these by functions *one* and *two*, with constant *null*. How to return 3 €? Which coins and which order?

Specification:

$$S_A = \{ A : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{two}(\text{one}(\text{null}))) \}.$$

Implementation:

$$S_B = \{ \begin{array}{l} B_1 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{two}(\text{one}(\text{null}))), \\ B_2 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{one}(\text{one}(\text{one}(\text{null})))), \\ B_3 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{one}(\text{two}(\text{null}))) \end{array} \}.$$

Example: coin value

Environment to check coin value:

$$Env = \{ \begin{array}{l} E_{null} : -[] \rightarrow M(null), \\ E_{one} : M(x) -[] \rightarrow M(one(x)), \\ E_{two} : M(x) -[] \rightarrow M(two(x)), \\ E_{check} : M(x), In_{Env}(x) -[] \rightarrow Out_{Env}(true) \end{array} \}.$$

Example: coin value

Environment to check coin value:

$$Env = \{ \begin{array}{l} E_{null} : -[] \rightarrow M(null), \\ E_{one} : M(x) -[] \rightarrow M(one(x)), \\ E_{two} : M(x) -[] \rightarrow M(two(x)), \\ E_{check} : M(x), In_{Env}(x) -[] \rightarrow Out_{Env}(true) \end{array} \}.$$

This environment will distinguish based on the exact order of coins.

Example: coin value

Environment to check coin value:

$$Env = \{ \begin{array}{l} E_{null} : -[] \rightarrow M(null), \\ E_{one} : M(x) -[] \rightarrow M(one(x)), \\ E_{two} : M(x) -[] \rightarrow M(two(x)), \\ E_{check} : M(x), In_{Env}(x) -[] \rightarrow Out_{Env}(true) \end{array} \}.$$

This environment will distinguish based on the exact order of coins.

We only want to check that the same value is returned.

Example: coin value

Environment to check coin value:

$$Env = \{ \begin{array}{l} E_{null} : -[] \rightarrow M(null), \\ E_{one} : M(x) -[] \rightarrow M(one(x)), \\ E_{two} : M(x) -[] \rightarrow M(two(x)), \\ E_{check} : M(x), In_{Env}(x) -[] \rightarrow Out_{Env}(true) \end{array} \}.$$

This environment will distinguish based on the exact order of coins.

We only want to check that the same value is returned.

Add the equation $two(x) = one(one(x))$. Then the specification and implementation are observationally equivalent for the given environment.

How to prove ObsEq: Bi-Systems

- Requiring the tool to guess which rule on one side is simulated by what rule on the other is difficult
- Instead use **diff**-terms and thus same-format rules
- Additionally we require that each rule is simulated by itself
- This is sufficient to prove ObsEq, but not necessary

What is covered?

- Real-or-random test
- Privacy properties of voting and auctions
- Ciphertext indistinguishability
- Authenticated key-exchange
- All previous example in this lecture

Bi-System yields two systems

For bi-system S we get

- Left instance $L(S)$
- Right instance $R(S)$

where for all terms $diff[M, N]$ we get M in $L(S)$ and N in $R(S)$.

Example: pairs again

Single bi-system S combines S_A and S_B as

$$S = \{ \quad AB : \text{---} \rightarrow \text{Out}_{\text{Sys}}((x, \text{diff}[x, y])) \quad \},$$

where $L(S) = S_A$ and $R(S) = S_B$.

Example: coins again

$$S = \{ \begin{array}{l} AB_1 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{two}(\text{one}(\text{null}))]]), \\ AB_2 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{one}(\text{one}(\text{one}(\text{null})))]]), \\ AB_3 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{one}(\text{two}(\text{null}))]]) \}. \end{array}$$

Example: coins again

$$S = \{ \begin{array}{l} AB_1 : -[\] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{two}(\text{one}(\text{null}))]]), \\ AB_2 : -[\] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{one}(\text{one}(\text{one}(\text{null})))]]), \\ AB_3 : -[\] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{one}(\text{two}(\text{null}))]]) \}. \end{array}$$

Keeping the environment Env from earlier results in the bi-system S not satisfying observational equivalence.

Example: coins again

$$S = \{ \begin{array}{l} AB_1 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{two}(\text{one}(\text{null}))]]), \\ AB_2 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{one}(\text{one}(\text{one}(\text{null})))]]), \\ AB_3 : -[] \rightarrow \text{Out}_{\text{Sys}}(\text{diff}[\text{two}(\text{one}(\text{null})), \\ \text{one}(\text{two}(\text{null}))]]) \}. \end{array}$$

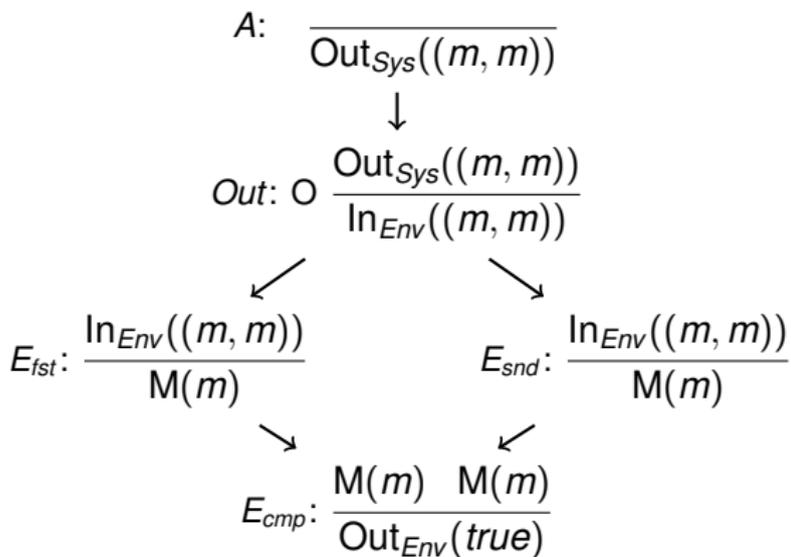
Keeping the environment Env from earlier results in the bi-system S not satisfying observational equivalence.

If we add the equation $\text{two}(x) = \text{one}(\text{one}(x))$, then S satisfies observational equivalence.

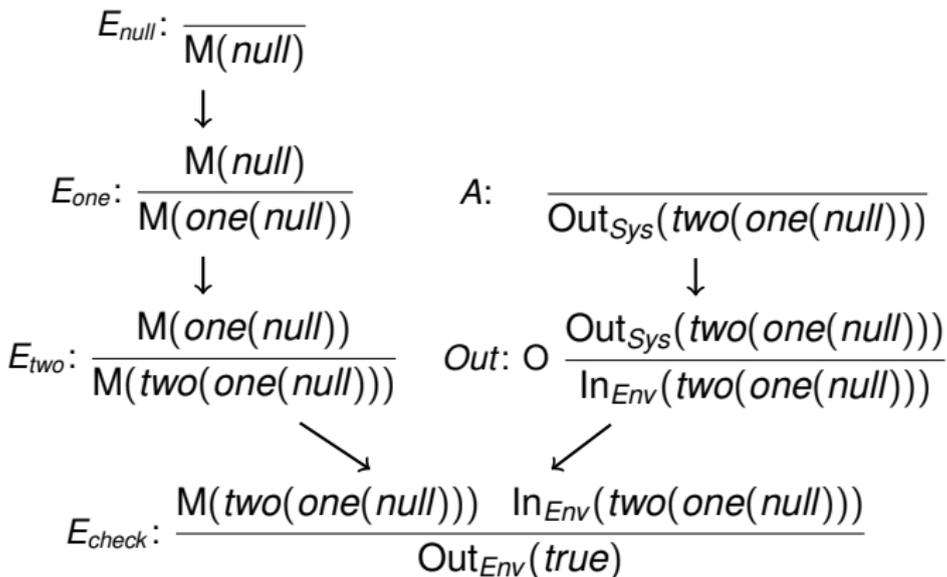
Dependency Graphs

- To simplify reasoning, our algorithm works with **dependency graphs** rather than with the labeled transition system.
- Dependency graphs are a data structure that formalize the entire structure of a system execution, including which facts originate from which rules, similar to recipes.
- They are well-suited for automated analysis using constraint solving, as they cover whole system state.
- Dependency graphs naturally give rise to an equivalence relation that implies observational equivalence; however, it is substantially simpler to verify.

Dependency graph example for pairs



Dependency graph example for coins



Mirroring dependency graphs

- **mirrors** of dependency graphs: mirrors contain all dependency graphs on the other side of the bi-system with the same structure,
 - ★ the same edges
 - ★ where the nodes are instances (potentially different due to the diff-terms) of the same rules.
- If the set of mirrors contains all “necessary instances” (for new diff variables), each transition is enabled on the other side as well.
- This guarantees observational equivalence.

Mirroring dependency graphs

- **mirrors** of dependency graphs: mirrors contain all dependency graphs on the other side of the bi-system with the same structure,
 - ★ the same edges
 - ★ where the nodes are instances (potentially different due to the diff-terms) of the same rules.
- If the set of mirrors contains all “necessary instances” (for new diff variables), each transition is enabled on the other side as well.
- This guarantees observational equivalence.
- Mirroring DG are sufficient, but not necessary for observational equivalence, see Theorem 1.

Reminder: Algorithm

```
1: function VERIFY(S)
2:    $RU \leftarrow L(S) \cup R(S) \cup IF \cup Env$ 
3:   while  $RU \neq \emptyset$  do
4:     choose  $r \in RU$ ,  $RU \leftarrow (RU \setminus \{r\})$ 
5:     compute  $DG \leftarrow dgraphs(r)$  by constraint solving
6:     if  $\exists dg \in DG$  s.t.  $mirrors(dg)$  lacks ground instances
7:       then return "potential attack found: ",  $dg$ 
8:   return "verification successful"
```

Implementation

Execution trace of

$$S_0, (l_1 \xrightarrow[\text{rec}_1]{a_1} r_1), S_1, \dots, S_{k-1}, (l_k \xrightarrow[\text{rec}_k]{a_k} r_k), S_k$$

is the sequence of the multisets of its action labels $[a_1, \dots, a_k]$.

Implementation

Execution trace of

$$S_0, (l_1 \xrightarrow[rec_1]{a_1} r_1), S_1, \dots, S_{k-1}, (l_k \xrightarrow[rec_k]{a_k} r_k), S_k$$

is the sequence of the multisets of its action labels $[a_1, \dots, a_k]$.

The adversary's message deduction capabilities are captured by the following set of rules.

$$\begin{aligned}
 MD = & \{ \text{Out}(t) \text{---} \square \rightarrow K(t), K(t) \text{---} [K(t)] \rightarrow \text{In}(t), \\
 & \text{Fr}(x: fr) \text{---} \square \rightarrow K(x: fr), \square \text{---} \square \rightarrow K(x: pub) \} \\
 & \cup \{ K(t_1), \dots, K(t_n) \text{---} \square \rightarrow K(f(t_1, \dots, t_n)) \mid f \in \Sigma_{Fun}^n \}
 \end{aligned}$$

Implementation

Added new message deduction rule:

$$IEquality : K^\downarrow(x), K^\uparrow(x) \multimap \square \rightarrow .$$

Implementation

Added new message deduction rule:

$$IEquality : K^\downarrow(x), K^\uparrow(x) \multimap \square \rightarrow .$$

Compare different sources of x .

Implementation

Added new message deduction rule:

$$IEquality : K^\downarrow(x), K^\uparrow(x) \multimap \square \rightarrow .$$

Compare different sources of x .

- If a side can construct the same value twice this rule is applicable
- If the same source is used twice, the other side can trivially mirror it
- But, if one side can decrypt a value and compare it, but the other cannot, this will be exposed.

Outline

- 1 Observational Equivalence: High-Level Overview
- 2 Observational Equivalence Details
- 3 Case Studies**

Case studies

- Probabilistic encryption - next
- DDH - see theory file and paper (if time permits)

Case studies

- Probabilistic encryption - next
- DDH - see theory file and paper (if time permits)
- TPM_Envelope - find attack in lab
- Signed Diffie-Hellman - analyze in lab

Probabilistic encryption

Equational theory:

$$pdec(penc(m, pk(k), r), k) \simeq m.$$

This equation gives rise to the following **decryption rule** for probabilistic encryption for the adversary, which is automatically generated by TAMARIN:

$$D_{penc} : K(penc(m, pk(k), r)), K(k) \dashv\vdash K(m).$$

Probabilistic encryption (ctd.)

We now express, as a bi-system, that a probabilistic encryption cannot be distinguished from a random value:

$$S = \left\{ \begin{array}{l} GEN : Fr(k) \xrightarrow{[]} Key(k), Out(pk(k)) \\ ENC : Key(k), Fr(r_1), Fr(r_2), In(x) \xrightarrow{[]} \\ \quad Out(diff[r_1, penc(x, pk(k), r_2)]) \end{array} \right\}.$$

Probabilistic encryption (ctd.)

We summarize below how TAMARIN automatically proves this property. The algorithm VERIFY (line 2) first constructs the set RU of rules to be analyzed,

$$RU = \{ L(GEN), R(GEN), L(ENC), R(ENC), \\ \text{FRESH}_{\text{Sys}}, \text{FRESH}_{\text{Env}}, \text{IEquality}, \text{Dpenc} \},$$

together with the remaining rules in IF and Env .

Algorithm VERIFY iterates over all rules (lines 3–4) until either an attack is found (line 7) or all rules have been checked and the verification is complete (line 8), which happens here.

Probabilistic encryption (ctd.)

VERIFY works like this:

- VERIFY first generates dependency graphs with the rule as the root (line 5).
- For each resulting dependency graph, it looks for a mirror (line 6) that contains all instances required by the definition of normal dependency graph equivalence.
- It always finds a mirror and verification therefore succeeds.
- TAMARIN analyzes left-diff and right-diff instantiations independently, we present them together.
- No explicit dependency graphs shown (run the tool to see those); however, we do explain how they are mirrored in each case so that the verification succeeds.

Probabilistic encryption (ctd.)

$$S = \{ \begin{array}{l} GEN : Fr(k) \dashrightarrow Key(k), Out(pk(k)) \\ ENC : Key(k), Fr(r_1), Fr(r_2), In(x) \dashrightarrow \\ Out(diff[r_1, penc(x, pk(k), r_2)]) \end{array} \}.$$

- As rule *GEN* does not contain a diff-term, the left diff-instantiation of this rule is identical to the right diff-instantiation.
- The rule has only a single fresh fact as its premise and thus any dependency graph with this rule at its root contains only those two rule instances and is trivially mirrored by itself.

Probabilistic encryption (ctd.)

$$S = \left\{ \begin{array}{l} GEN : Fr(k) \dashrightarrow Key(k), Out(pk(k)) \\ ENC : Key(k), Fr(r_1), Fr(r_2), In(x) \dashrightarrow \\ Out(diff[r_1, penc(x, pk(k), r_2)]) \end{array} \right\}.$$

- The rule *ENC* has the same premises in the left- and right-hand side system and is therefore identical for the purpose of dependency graph computation with the *ENC* rule as root. (Note that outputs will be considered using the equality rule below.)
- The two fresh premises will result in identical dependency graphs, while the key and message reception input are independent.
- Hence both of them will have identical dependency graphs as premises, and the resulting dependency graphs are identical (up to the outputs) and therefore mirror each other.

Probabilistic encryption (ctd.)

The fresh rules $\text{FRESH}_{\text{Sys}}$ and $\text{FRESH}_{\text{Env}}$ have no premises. Hence the dependency graphs with them as root are just their instances, which mirror each other in the left- and right-hand system.

Probabilistic encryption (ctd.)

- For an equality rule instance of *IEquality* as the root of a dependency graph, the two premises are the same instance of a variable x .
- If both of the premises are adversary generated, then the resulting dependency graphs are the same in the left- and right-hand system, and thus will mirror themselves trivially.

Probabilistic encryption (ctd.)

- For an equality rule instance of *IEquality* as the root of a dependency graph, the two premises are the same instance of a variable x .
- If both of the premises are adversary generated, then the resulting dependency graphs are the same in the left- and right-hand system, and thus will mirror themselves trivially.
- Alternatively, if one of the premises uses the output of an instance of either the *ENC* rule or the *GEN* rule, then there is no dependency graph with a matching second premise.
- This is because all system outputs, $pk(k)$ for *GEN* and $r1$ or $penc(x, pk(k), r2)$ for *ENC*, contain a fresh value, k , $r1$, respectively $r2$, that is never available to the intruder.
- As this will never allow a complete dependency graph to be derived, no mirroring dependency graph is needed.

Probabilistic encryption (ctd.)

- For the decryption rule generated for the probabilistic encryption, this rule is never applicable on either side as the adversary never receives the keys needed for decrypting system generated encryptions.
- As there is no dependency graph, no mirroring one is needed.

Probabilistic encryption (ctd.)

- For the decryption rule generated for the probabilistic encryption, this rule is never applicable on either side as the adversary never receives the keys needed for decrypting system generated encryptions.
- As there is no dependency graph, no mirroring one is needed.

(One might mistakenly think that this rule might apply to intruder-generated terms. However, this is not the case due to the restrictions on how the adversary may combine its knowledge (K^\downarrow vs K^\uparrow) and, in any case, both sides would use the same dependency graphs as premise, so the result would be the same.)

Probabilistic encryption (ctd.)

For all other adversary rules, it is obvious that they result in identical dependency graphs on both sides.

More precisely:

- Construction rules have adversary knowledge input and thus the same dependency graphs as premises.
- For the deconstruction rules, the only relevant one is the previous decryption rule, as that is the only one that can use information coming out of the system; all other rules can only be used on adversary-generated terms and thus have the same dependency graphs as premises.

Probabilistic encryption (ctd.)

This completes our summary of TAMARIN's verification of observational equivalence for this example. TAMARIN automatically constructs the proof in under 0.2 seconds.

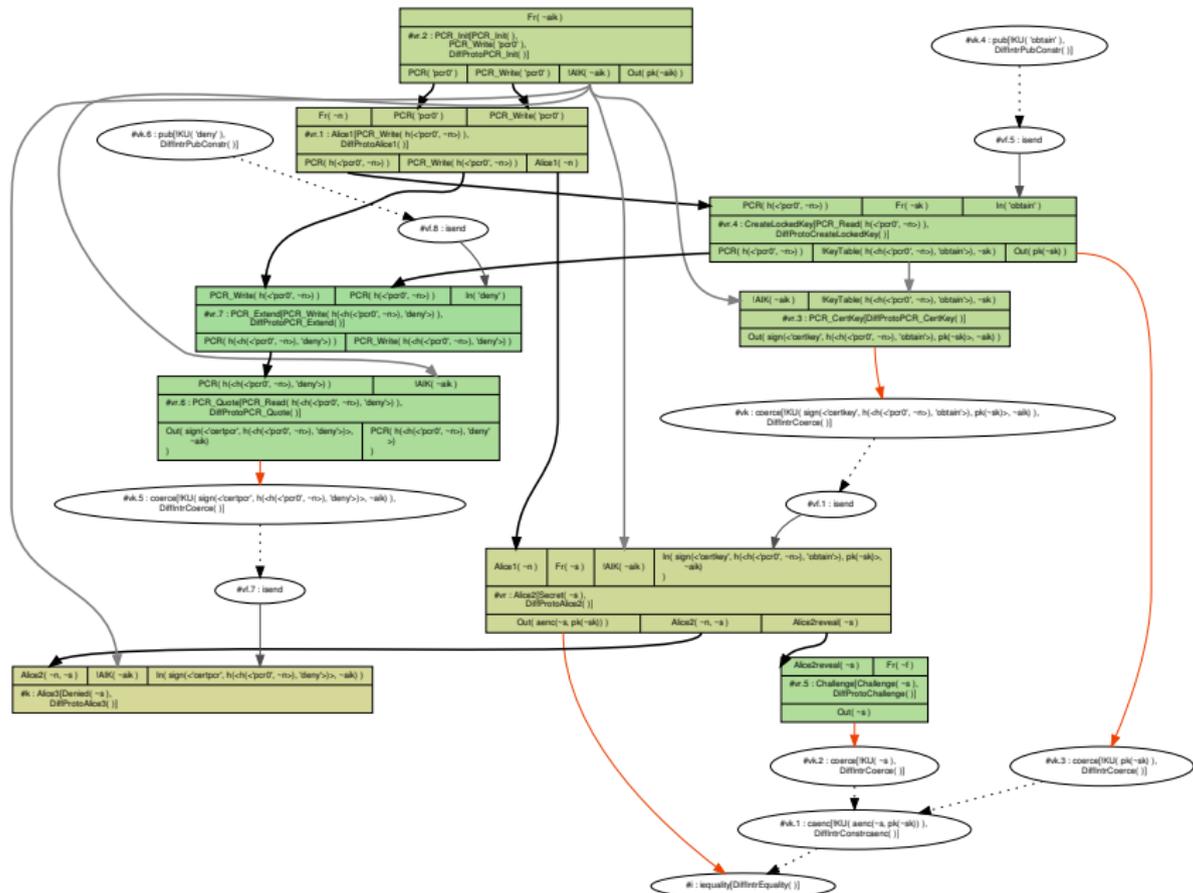
DDH

Demo if time permits

Simplified TPM_Envelope specification

- Init* : Fr(*aik*) \rightarrow PCR('pcr0'), AIK(*aik*), Out(pk(*aik*))
Ext : PCR(*x*), In(*y*) \rightarrow PCR(*h*(*x*, *y*))
CertK : AIK(*aik*), KT(*lock*, *sk*) \rightarrow
 Out(sign('certk', *lock*, pk(*sk*)), *aik*)
Quote : PCR(*x*), AIK(*aik*) \rightarrow
 PCR(*x*), Out(sign('certcpr', *x*), *aik*)
Unbind : PCR(*x*), KT(*x*, *sk*), In(aenc(*m*, pk(*sk*))) \rightarrow
 PCR(*x*), Out(*m*)
A1 : Fr(*n*), PCR(*x*) \rightarrow PCR(*h*(*x*, *n*)), A1(*n*)
A2 : Fr(*s*), A1(*n*), AIK(*aik*),
 In(sign('certk', *h*(*h*('pcr0', *n*), 'obtain'), pk), *aik*)
 \rightarrow Out(aenc(*s*, pk)), A2(*n*, *s*), A2ror(*s*)
A3 : In(sign('certpcr', *h*(*h*('pcr0', *n*), 'deny')), *aik*),
 A2(*n*, *s*), AIK(*aik*) \rightarrow [Denied(*s*)]
CLKey : Fr(*sk*), PCR(*x*), In(*lock*) \rightarrow
 PCR(*x*), KT(*h*(*x*, *lock*), *sk*), Out(pk(*sk*))
ROR : A2ror(*s*), Fr(*f*) \rightarrow Out(diff[*s*, *f*])

TPM_Envelope attack



Conclusions

- Observational equivalence can be automatically proven
- Diff-terms allow concise description
- Approach is sound (mirroring DG are sufficient, not necessary)
- Compares well to other approaches

Conclusions

- Observational equivalence can be automatically proven
- Diff-terms allow concise description
- Approach is sound (mirroring DG are sufficient, not necessary)
- Compares well to other approaches

(-: Happy Proving :-)