Deciding, typing and composing (α, β)-privacy

Laouen Fernet[1,*], Sebastian Mödersheim[1] and Luca Viganò[2]

[1] DTU Compute, Technical University of Denmark, DK * lpkf@dtu.dk
[2] Informatics, King's College London, UK

1. Problem: privacy in security protocols

Current trend of **increasing digitalization**: more and more applications use private information to provide various services.

Health Payment Transport Voting

We need **strong guarantees** that digital applications respect privacy. We focus on applications written as **security protocols**: participants exchange messages, often using cryptography.

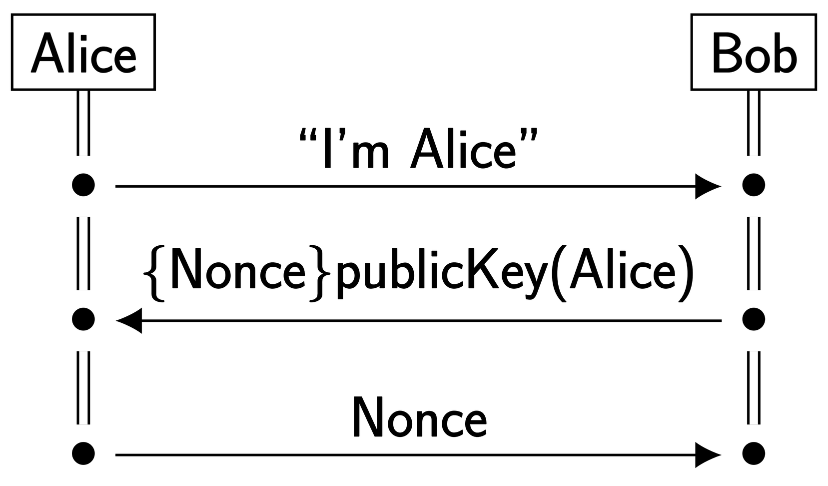

Example of a simple security protocol

We use (α, β)**-privacy** to characterize privacy with logical formulas.

α is the **payload**: information intentionally disclosed.

β is the **technical information**: intruder knowledge.

Example: $\alpha \equiv x_1, \ldots, x_n \in$ Agent $\rightarrow$ unlinkability goal
If $\beta \Rightarrow x_1 =$ Alice or $\beta \Rightarrow x_2 = x_3$, then it is a **violation of privacy**: the intruder has learned more than allowed.

Input

```
...
* x in {a,b,i}.  # Pick an agent
* y in {yes,no}. # Flip a coin
receive M.
try N = dcrypt(inv(pk(s)),M) in
  if y = yes then
    new R. send crypt(pk(x),pair(yes,N),R)
  else
    new R. send crypt(pk(x),no,R)
...
```

2. Objective: automated verification

Specification of a protocol: **transition system** where executing an **atomic transaction** leads to the next state. In each state, a pair (α, β) defines the privacy goals and intruder knowledge.

Our objective: **decide privacy expressed as a reachability property**.

Main challenge: verify an **infinite state space**.

1. The intruder has infinitely many choices when sending messages.
   $\rightarrow$ We use a **symbolic representation** with constraint systems.
2. Some transaction can always be executed.
   $\rightarrow$ We only look at a **bounded number** of transactions.

Our **decision procedure** in short:

1. Execute a transaction.
2. Saturate the intruder knowledge by decrypting and comparing messages.
3. Verify (α, β)-privacy in the symbolic states reached.
4. Repeat until we reach the bound specified.

Computation

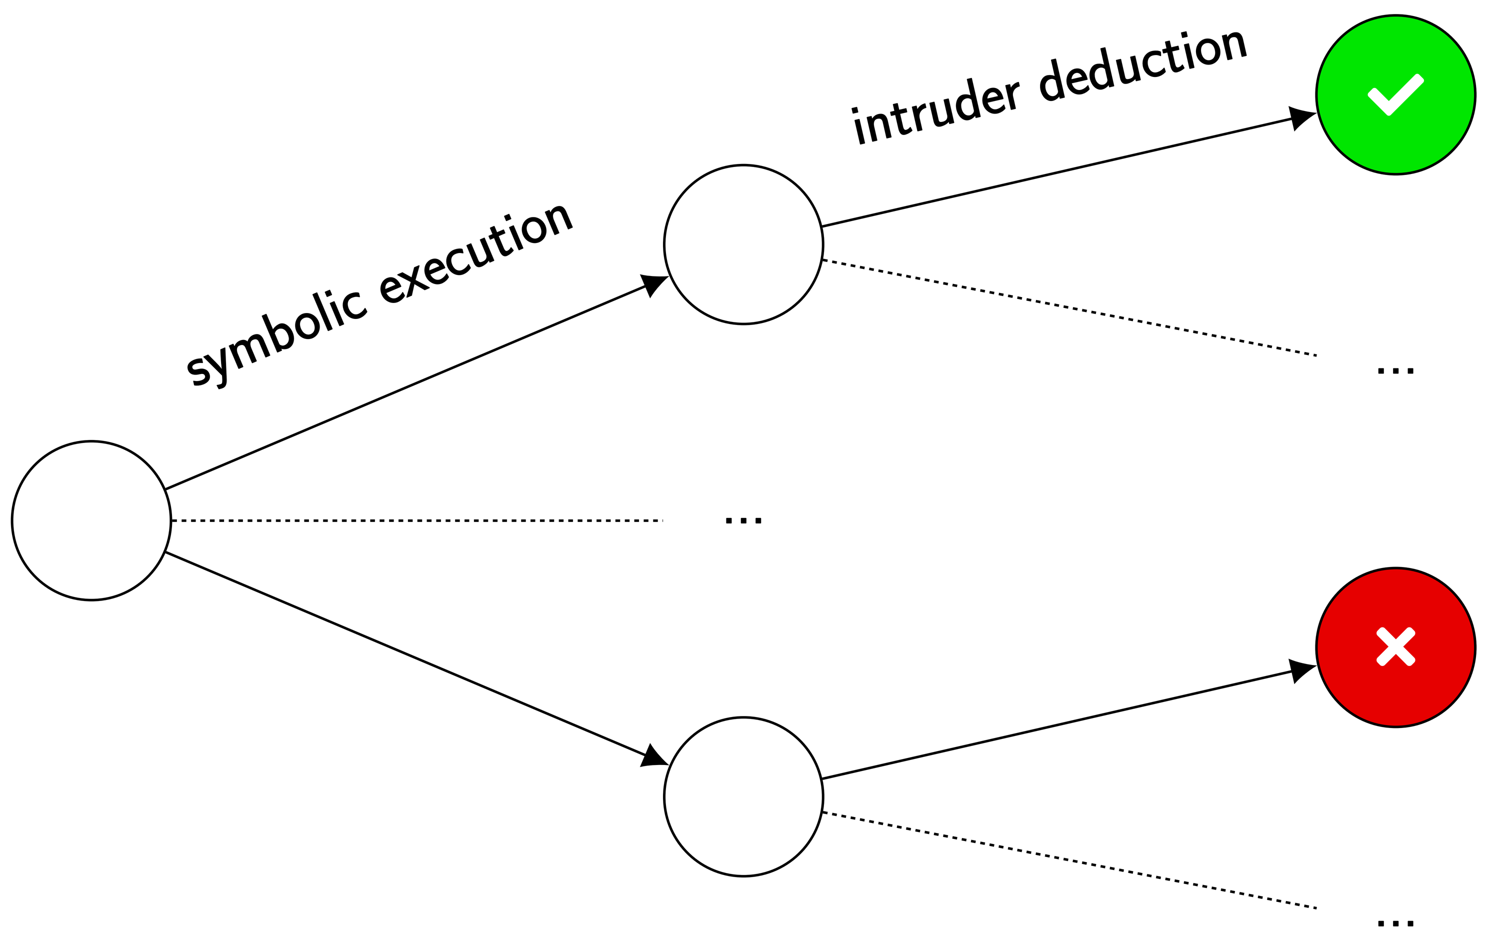

3. Theoretical results and tool support

Main outcomes:

1. decision procedure, with proofs of **correctness**, and **prototype tool**.
2. typing result: under certain conditions, we do not lose attacks if we restrict the intruder to sending only **well-typed messages**.
3. compositionality result: given a specification of components of a system and their **abstract interfaces**, if each component is secure then so is the entire system.

Input: specification of the protocol with a bound.
Output:

- either **attack trace**: reachable state with a violation of privacy.
- or confirmation that **the privacy goals are achieved**.

Case studies: Basic Hash, OSK, BAC, Private Authentication, NSL, simplified TLS.

Conclusion: (α, β)-privacy allows for **declarative and intuitive** specification of privacy and automated verification is **practical**.

Output

```
Privacy violation found after 2 transactions.
alpha: x in {a,b,i} and y in {yes,no}
beta implies: x = i and y = no
(alpha, beta)-privacy does not hold for the
state where the intruder has sent
crypt(pk(s),R1,R2) and has successfully
decrypted the reply from the server.
...
```