

StatVerif: verification of stateful processes

Mark D. Ryan

joint work with Myrto Arapinis, Joshua Phillips and Eike Ritter

FOSAD 2012

Mutable shared state

Agents that have **mutable state shared between runs**:

- Hardware tokens
 - Smart cards: capabilities, ...
 - RFID tags: pseudonyms, ...
 - TPMs: PCRs, ...
 - HSM: PIN codes, ...
- Trusted parties in contract signing
- Mobile phones: TMSIs, location
- Web servers, database servers, ...
- VANETS: pseudonyms
- ...

The Trusted Platform Module (TPM)



1. Secure storage
 2. Platform authentication
 3. Platform integrity reporting
- 200M in existence (laptops, desktops, servers)
 - specified by Trusted Computing Group (> 700 pages)

TPM functionality

Secure storage

- TPM stores keys and other sensitive data in **shielded memory**
- A user can store content that is encrypted by **keys only available to the TPM**

Platform authentication

- Each TPM chip has a **unique** and **secret** key
- A platform can obtain keys by which it can **authenticate** itself

Platform measurement and reporting

- TPM contains some **internal memory slots** called PCRs, and some keys can be **locked** to a particular PCR value
- PCR values can be modified, but only in a specific way

Digital rights management



unforgeable
configuration report

Secure environment






Richard Stallman
Creator of GNU, Emacs,
GCC, GPL, the Free
Software Foundation

“With a plan they call *trusted computing*, large media corporations, together with computer companies such as Microsoft and Intel, are planning to make your computer obey them instead of you.”

He calls it “treacherous computing”.

Attestation from cloud



Cloud server

Coding persistent state in pi calculus

let TPM	=	$\nu p.((!TPMsession) \mid TPMstate)$	
let TPMsession	=	$p(x).$	read & lock state
		$c(y).$	read command
		\dots	process command
		$\bar{c}\langle output \rangle.$	return output
		$\bar{p}\langle newstate \rangle.$	store & release state
let TPMstate	=	$\bar{p}\langle initstate \rangle.$	provide initial state
		$!(p(x).$	receive value into state
		$\bar{p}\langle x \rangle)$	return stored state

This is correct in applied pi calculus. However, because of the abstractions ProVerif makes, one cannot use this encoding to prove properties of the TPM in ProVerif.

Outline of the talk

- 1 StatVerif syntax
- 2 StatVerif semantics
- 3 Clauses
- 4 Abstractions for termination
- 5 Conclusion

StatVerif syntax: processes

$P, Q ::=$

$out (M, N); P$

$in (M, x); P$

$P \mid Q$

$!P$

$new a; P$

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

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

$[s \mapsto M]$

$read s_1, \dots, s_n \text{ as } x_1, \dots, x_n; P$

$s_1, \dots, s_n := M_1, \dots, M_n; P$

$lock s_1, \dots, s_n; P$

$unlock s_1, \dots, s_n; P$

processes

output

input

parallel composition

replication

restriction

destructor application

conditional

state cell

read

write

lock

unlock

Example

```
let TPMsession = lock s;  
    read s as xstate;  
    in (c, xcommand);  
    ... process command and compute new state. . .  
    s := newstate;  
    unlock s;  
    out (c, output)
```

Outline of the talk

- 1 StatVerif syntax
- 2 StatVerif semantics**
- 3 Clauses
- 4 Abstractions for termination
- 5 Conclusion

Semantic configurations are tuples $(\mathcal{E}, \mathcal{S}, \mathcal{P})$ where:

- \mathcal{E} is a set of names,
- \mathcal{S} maps cell names to their current values, and
- \mathcal{P} is a multiset of pairs (P, λ) with P a process and λ a set of cell names. $s \in \lambda$ iff P has exclusive access to cell s .

Initial configuration for a process P : $(\text{fn}(P), \emptyset, \{(P, \emptyset)\})$

StatVerif semantics: reduction relation

...

$$(\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{([s \mapsto M], \emptyset)\}) \rightarrow (\mathcal{E}, \mathcal{S} \cup \{s \mapsto M\}, \mathcal{P})$$

if $s \in \text{dom}(\mathcal{E})$ **and** $s \notin \text{dom}(\mathcal{S})$

$$(\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{(\text{lock } s_1, \dots, s_n; P, \lambda)\}) \rightarrow (\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{(P, \lambda \cup \{s_1, \dots, s_n\})\})$$

if $\forall (Q, \lambda') \in \mathcal{P}. \lambda' \cap \{s_1, \dots, s_n\} = \emptyset$

$$(\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{(\text{unlock } s_1, \dots, s_n; P, \lambda)\}) \rightarrow (\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{(P, \lambda \setminus \{s_1, \dots, s_n\})\})$$

$$(\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{(\text{read } s_1, \dots, s_n$$

as $x_1, \dots, x_n; P, \lambda)\}) \rightarrow (\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{(P\{\mathcal{S}(s_1)/x_1, \dots, \mathcal{S}(s_n)/x_n\}, \lambda)\})$

if $\{s_1, \dots, s_n\} \subseteq \text{dom}(\mathcal{S})$ **and** $\forall (Q, \lambda') \in \mathcal{P}. \lambda' \cap \{s_1, \dots, s_n\} = \emptyset$

$$(\mathcal{E}, \mathcal{S}, \mathcal{P} \cup \{(s_1, \dots, s_n :=$$

 $M_1, \dots, M_n; P, \lambda)\}) \rightarrow (\mathcal{E}, \mathcal{S}[s_1 \mapsto M_1, \dots, s_n \mapsto M_n], \mathcal{P} \cup \{(P, \lambda)\})$

if $\{s_1, \dots, s_n\} \subseteq \text{dom}(\mathcal{S})$ **and** $\forall (Q, \lambda') \in \mathcal{P}. \lambda' \cap \{s_1, \dots, s_n\} = \emptyset$

Bounded number of cell names

Our translation only applies to StatVerif processes of the form:

$$\text{new } \tilde{m}; ([s_1 \mapsto M_1] \mid \cdots \mid [s_n \mapsto M_n] \mid P)$$

such that \mathbf{P} contains no $[s \mapsto M]$.

Outline of the talk

- 1 StatVerif syntax
- 2 StatVerif semantics
- 3 Clauses**
- 4 Abstractions for termination
- 5 Conclusion

The Horn clauses representation

The translation of a StatVerif process generates clauses built around the following two predicates

- $attacker(\tilde{M}, N)$ means that state \tilde{M} is reachable and in that state the attacker might know the value N ;
- $message(\tilde{M}, K, N)$ means that state \tilde{M} is reachable and in that state the value N is available on channel K .

Attacker clauses: constructors and destructors

- The attacker can build new messages by applying any **constructor** to messages he knows

Asymmetric encryption

$$\text{attacker}(xs, xk) \rightarrow \text{attacker}(xs, pk(xk))$$
$$\text{attacker}(xs, xk) \wedge \text{attacker}(xs, xm) \rightarrow \text{attacker}(xs, \text{aenc}(xk, xm))$$

- The attacker can analyse messages by applying any **destructor** to messages he knows

Asymmetric-key decryption

$$\text{attacker}(xs, xk) \wedge \text{attacker}(xs, \text{aenc}(pk(xk), xm)) \rightarrow \text{attacker}(xs, xm)$$

Attacker clauses: public channels

- The attacker can **send messages** on public channels

$$\text{attacker}(xs, xc) \wedge \text{attacker}(xs, xm) \rightarrow \text{message}(xs, xc, xm)$$

- The attacker can **eavesdrop** on public channels

$$\text{attacker}(xs, xc) \wedge \text{message}(xs, xc, xm) \rightarrow \text{attacker}(xs, xm)$$

Example of protocol clauses: the hardware token

```
let config = lock s;  
  in (c, x);  
  if y = init then  
    s := x;  
  unlock s
```

A tag not yet configured (*if y = init then*) can be configured with any value x sent to it:

$$\text{message}(\text{init}[], c[], x) \wedge \text{message}(\text{init}[], xc, xm) \\ \wedge \text{message}(\text{init}[], yc, ym) \rightarrow \text{message}(x, yc, ym)$$
$$\text{message}(\text{init}[], c[], x) \wedge \text{message}(\text{init}[], xc, xm) \\ \wedge \text{attacker}(\text{init}[], ym) \rightarrow \text{attacker}(x, ym)$$

Protocol clauses

$$\begin{aligned}
 [0]\rho H\iota\phi\lambda &= \emptyset \\
 [Q_1 \mid Q_2]\rho H\iota\phi\emptyset &= [Q_1]\rho H\iota\phi\emptyset \cup [Q_2]\rho H\iota\phi\emptyset \\
 [!Q]\rho H\iota\phi\emptyset &= [Q]\rho H\iota\phi\emptyset \\
 [new\ a; Q]\rho H\iota\phi\lambda &= \begin{cases} [Q](\rho \cup \{a \mapsto a[u]\})H\iota\phi\lambda & \text{if } a \in \text{bn}(P) \\ [Q](\rho \cup \{a \mapsto \text{attn}\})H\iota\phi\lambda & \text{otherwise} \end{cases} \\
 [in\ (M, x); Q]\rho H\iota\phi\lambda &= [Q]\rho' H'(x :: \iota)\phi'\lambda \\
 \text{where } \phi' = \phi[k \mapsto \text{fresh} \mid k \notin \lambda] \text{ and } \rho' = \rho \cup \{x \mapsto x\} \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\} \text{ and } H' = H \wedge \text{message}(\phi', \rho(M), x) \\
 [out\ (M, N); Q]\rho H\iota\phi\lambda &= \{H \Rightarrow \text{message}(\phi, \rho(M), \rho(N))\} \cup [Q]\rho H\iota\phi\lambda \\
 [let\ x = g(M_1, \dots, M_n) \text{ in } Q_1 \text{ else } Q_2]\rho H\iota\phi\lambda &= [Q_2]\rho H\iota\phi\lambda \\
 &\cup \left\{ [Q_1]((\rho\sigma) \cup \rho')(H\sigma)(\iota\sigma)(\phi\sigma)\lambda \mid \right. \\
 &\quad \left. \begin{aligned} &g(p_1, \dots, p_n) \rightarrow p \in \text{def}(g) \\ &\text{and } \sigma' = \{z \mapsto \text{fresh} \mid z \in \text{fv}(g(p_1, \dots, p_n))\} \\ &\text{and } \sigma = \text{mgu}(g(M_1\rho, \dots, M_n\rho), g(p_1\sigma', \dots, p_n\sigma')) \\ &\text{and } \rho' = \{x \mapsto \rho\sigma'\} \cup \{z\sigma' \mapsto z\sigma' \mid z \in \text{fv}(g(p_1, \dots, p_n))\} \end{aligned} \right\} \\
 [if\ M = N \text{ then } Q_1 \text{ else } Q_2]\rho H\iota\phi\lambda &= [Q_1](\rho\sigma)(H\sigma)(\iota\sigma)(\phi\sigma)\lambda \cup [Q_2]\rho H\iota\phi\lambda \\
 &\quad \text{where } \sigma = \text{mgu}(\rho(M), \rho(N)) \\
 [lock\ s_{i_1}, \dots, s_{i_m}; Q]\rho H\iota\phi\lambda &= [Q](\rho \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\})H\iota\phi'(\lambda \cup \{i_1, \dots, i_m\}) \\
 &\quad \text{where } \phi' = \phi[k \mapsto \text{fresh} \mid k \notin \lambda] \\
 [unlock\ s_{i_1}, \dots, s_{i_m}; Q]\rho H\iota\phi\lambda &= [Q](\rho \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\})H\iota\phi'(\lambda \setminus \{i_1, \dots, i_m\}) \\
 &\quad \text{where } \phi' = \phi[k \mapsto \text{fresh} \mid k \notin \lambda] \\
 [read\ s_{i_1}, \dots, s_{i_m} \text{ as } x_1, \dots, x_m; Q]\rho H\iota\phi\lambda &= [Q]\rho' H'(x_1 :: \dots :: x_m :: \iota)\phi'\lambda \\
 &\quad \text{where } \rho' = \rho \cup \{x_j \mapsto \phi'_j \mid 1 \leq j \leq m\} \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\} \cup \{\text{vc} \mapsto \text{vc}, \text{vm} \mapsto \text{vm}\} \\
 &\quad \text{and } \phi' = \phi[k \mapsto \text{fresh} \mid k \notin \lambda] \\
 &\quad \text{and } H' = H \wedge \text{message}(\phi', \text{vc}, \text{vm}) \\
 &\quad \text{and } \text{vc}, \text{vm} \text{ fresh} \\
 [s_{i_1}, \dots, s_{i_m} := M_1, \dots, M_m; Q]\rho H\iota\phi\lambda &= [Q](\rho \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\})H\iota\phi''\lambda \\
 &\cup \{H \wedge \text{message}(\phi', \text{vc}, \text{vm}) \Rightarrow \text{message}(\phi'', \text{vc}, \text{vm})\} \\
 &\cup \{H \wedge \text{attacker}(\phi', \text{vm}) \Rightarrow \text{attacker}(\phi'', \text{vm})\} \\
 &\quad \text{where } \phi' = \phi[k \mapsto \text{fresh} \mid k \notin \lambda] \\
 &\quad \text{and } \phi'' = \phi'[i_j \mapsto \rho(M_j) \mid 1 \leq j \leq m] \\
 &\quad \text{and } \text{vc}, \text{vm} \text{ fresh}
 \end{aligned}$$

Protocol clauses

$$\begin{aligned}
 \llbracket 0 \rrbracket \rho H \iota \phi \lambda &= \emptyset \\
 \llbracket Q_1 \mid Q_2 \rrbracket \rho H \iota \phi \emptyset &= \llbracket Q_1 \rrbracket \rho H \iota \phi \emptyset \cup \llbracket Q_2 \rrbracket \rho H \iota \phi \emptyset \\
 \llbracket !Q \rrbracket \rho H \iota \phi \emptyset &= \llbracket Q \rrbracket \rho H \iota \phi \emptyset \\
 \llbracket \text{new } a; Q \rrbracket \rho H \iota \phi \lambda &= \begin{cases} \llbracket Q \rrbracket (\rho \cup \{a \mapsto a[u]\}) H \iota \phi \lambda & \text{if } a \in \text{bn}(P) \\ \llbracket Q \rrbracket (\rho \cup \{a \mapsto \text{attn}[]\}) H \iota \phi \lambda & \text{otherwise} \end{cases} \\
 \llbracket \text{in } (M, x); Q \rrbracket \rho H \iota \phi \lambda &= \llbracket Q \rrbracket \rho' H'(x :: \iota) \phi' \lambda \\
 \text{where } \phi' = \phi [k \mapsto \text{fresh} \mid k \notin \lambda] \text{ and } \rho' = \rho \cup \{x \mapsto x\} \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\} \text{ and } H' = H \wedge \text{message}(\phi', \rho(M), x) \\
 \llbracket \text{out } (M, N); Q \rrbracket \rho H \iota \phi \lambda &= \{H \Rightarrow \text{message}(\phi, \rho(M), \rho(N))\} \cup \llbracket Q \rrbracket \rho H \iota \phi \lambda \\
 \llbracket \text{let } x = g(M_1, \dots, M_n) \text{ in } Q_1 \text{ else } Q_2 \rrbracket \rho H \iota \phi \lambda &= \llbracket Q_2 \rrbracket \rho H \iota \phi \lambda
 \end{aligned}$$

Assignments

$$\begin{aligned}
 \llbracket s_{i_1}, \dots, s_{i_m} := \\
 M_1, \dots, M_m; Q \rrbracket \rho H \iota \phi \lambda &= \llbracket Q \rrbracket (\rho \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\}) H \iota \phi'' \lambda \\
 &\cup \{H \wedge \text{message}(\phi', \text{vc}, \text{vm}) \Rightarrow \text{message}(\phi'', \text{vc}, \text{vm})\} \\
 &\cup \{H \wedge \text{attacker}(\phi', \text{vm}) \Rightarrow \text{attacker}(\phi'', \text{vm})\} \\
 &\text{where } \phi' = \phi [k \mapsto \text{fresh} \mid k \notin \lambda] \\
 &\text{and } \phi'' = \phi' [i_j \mapsto \rho(M_j) \mid 1 \leq j \leq m] \\
 &\text{and } \text{vc}, \text{vm} \text{ fresh}
 \end{aligned}$$

$$\begin{aligned}
 &\text{and } H' = H \wedge \text{message}(\phi', \text{vc}, \text{vm}) \\
 &\text{and } \text{vc}, \text{vm} \text{ fresh}
 \end{aligned}$$

$$\begin{aligned}
 \llbracket s_{i_1}, \dots, s_{i_m} := M_1, \dots, M_m; Q \rrbracket \rho H \iota \phi \lambda &= \llbracket Q \rrbracket (\rho \cup \{\phi'_k \mapsto \phi'_k \mid k \notin \lambda\}) H \iota \phi'' \lambda \\
 &\cup \{H \wedge \text{message}(\phi', \text{vc}, \text{vm}) \Rightarrow \text{message}(\phi'', \text{vc}, \text{vm})\} \\
 &\cup \{H \wedge \text{attacker}(\phi', \text{vm}) \Rightarrow \text{attacker}(\phi'', \text{vm})\} \\
 &\text{where } \phi' = \phi [k \mapsto \text{fresh} \mid k \notin \lambda] \\
 &\text{and } \phi'' = \phi' [i_j \mapsto \rho(M_j) \mid 1 \leq j \leq m] \\
 &\text{and } \text{vc}, \text{vm} \text{ fresh}
 \end{aligned}$$

Main result

Theorem (The StatVerif compiler is correct)

Let M be a message. Let P be a protocol of the form

$$\text{new } \tilde{m}; ([s_1 \mapsto M_1] \mid \cdots \mid [s_n \mapsto M_n] \mid Q)$$

such that P contains no $[s \mapsto M]$.

$$\text{Clauses}(P) \vdash \text{secrecy}(M) \Rightarrow P \models \text{secrecy}(M)$$

If $\forall \tilde{K}$ the fact *attacker*(\tilde{K}, M) is not derivable from $\text{Clauses}(P)$, then P preserves the secrecy of M .

Outline of the talk

- 1 StatVerif syntax
- 2 StatVerif semantics
- 3 Clauses
- 4 Abstractions for termination**
- 5 Conclusion

Making ProVerif work on the Horn clauses



Making ProVerif work on the Horn clauses



Theorem: There is no selection function such that the saturation of

$$\left\{ \begin{array}{l} \mathit{attacker}(u, x) \rightarrow \mathit{attacker}(h(u, v), x), \\ \mathit{attacker}(u, \mathit{enc}(u, m)) \rightarrow \mathit{attacker}(u, m) \end{array} \right\}$$

will always terminate.

Making ProVerif work on the Horn clauses

Safe abstraction:

Replace

$$attacker(x_p, x_v) \wedge attacker(x_p, x) \rightarrow attacker(h(x_p, x_v), x)$$

with n instances, in which x_p is

- $zero[]$
- $h(zero[], x_1)$
- $h(h(zero[], x_1), x_2)$
- $h(h(h(zero[], x_1), x_2), x_3)$
-

Notion of k -stability

Definition: k -stable

A rule R is **k -stable** if for any substitution θ grounding for R , for any PCR value $u = h(u_1, u_2)$ such that $\text{length}_{\text{pcr}}(u) > k$ we have that:

- either $(R\theta)[h(u_1, u_2) \rightarrow u_1] = R(\theta[h(u_1, u_2) \rightarrow u_1])$,
- or $(R\theta)[h(u_1, u_2) \rightarrow u_1]$ is a tautology.

Examples

- $\text{attacker}(x_p, \text{certkey}(\text{aik}[], \text{pair}(x_{pk}, h(\text{zero}[], a_1)))) \rightarrow \text{attacker}(x_p, \text{aenc}(x_{pk}, s_1))$
- $\text{attacker}(x_p, x_v) \wedge \text{attacker}(x_p, x) \rightarrow \text{attacker}(h(x_p, x_v), x)$

Proposition

Let \mathcal{R} be a finite set of rules and Q be a query such that \mathcal{R} and Q are k -stable. If Q is satisfiable then there exists a **k -bounded derivation** witnessing this fact.

Syntactic criterion to check k -stability

Lemma

Let $k \geq 0$ be an integer and $R = H \rightarrow C$ be a rule such that:

- 1 for all $h(v_1, v_2) \in st(R)$, $length_{pcr}(v_1, v_2) \leq k$;
- 2 for all $h(v_1, v_2) \in st(H)$, we have that $v_1 \notin \mathcal{X}$;
- 3 for all $h(v_1, v_2) \in st(C)$ such that $v_1 \in \mathcal{X}$, we have that $C[h(v_1, v_2) \rightarrow v_1] \in H$.

Then, we have that **the rule is k -stable**.

Examples

- $attacker(x_p, certkey(aik[], pair(x_{pk}, h(zero[], a_1)))) \rightarrow attacker(x_p, aenc(x_{pk}, s_1))$
- $attacker(x_p, x_v) \wedge attacker(x_p, x) \rightarrow attacker(h(x_p, x_v), x)$

→ Going back to our running example, it is sufficient to consider **1-bounded derivation** when checking satisfiability of a query.

Outline of the talk

- 1 StatVerif syntax
- 2 StatVerif semantics
- 3 Clauses
- 4 Abstractions for termination
- 5 Conclusion**

Conclusion

- Defined the **StatVerif calculus**
- Designed **StatVerif compiler**
- **Proved correctness** of the StatVerif compiler
- Proved security of:
 - our **hardware token device**
 - the Mackenzie *et al.* **optimistic contract signing protocol**
 - **TPM authentication protocols**

- **Prototype implementation**
- **More case studies**
 - UMTS protocols,
 - RFID protocols,
 - ...
- Extension to **authentication properties**
- Extension to **observational equivalence**