

Formal analysis of protocols based on TPM state registers

Stéphanie Delaune¹, Steve Kremer¹, Mark D. Ryan²,
and Graham Steel¹

¹ LSV, ENS Cachan & CNRS & INRIA Saclay Île-de-France, France

² School of Computer Science, University of Birmingham, UK

Thursday, June 9th, 2011

TPM - What is it?

Trusted Platform Module

Hardware chip designed to enable commodity computers to achieve **greater levels of security** than is possible in software alone.



Trusted Platform Module

Hardware chip designed to enable commodity computers to achieve **greater levels of security** than is possible in software alone.



- more than **200 millions** currently in existence (mostly in laptops)
→ already used by some applications (*e.g.* Disk encryption)
- specified by the Trusted Computing Group
→ more than **700 pages** of specification

<http://www.trustedcomputinggroup.org>

Secure storage:

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

Secure storage:

- TPM stores keys and other sensitive data in its **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 reliably.

TPM functionality

Secure storage:

- TPM stores keys and other sensitive data in its **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 reliably.

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 using some specific command (*e.g.* command **Extend**).

TPM - How is it used?

Application programming interface:

- create new keys (e.g. `CreateWrapKey`), and load them into the device (e.g. `LoadKey2`);
- manipulate these keys, and the PCRs
 - e.g. `UnBind` allows one to decrypt a ciphertext using a key that is stored into the TPM and locked to the current PCR value
 - e.g. `Quote` allows one to obtain a certificate attesting that a key is locked to a particular PCR value
 - e.g. `Extend` allows one to extend the current value of a PCR with some data x , i.e. $p := \text{SHA1}(p||x)$.

TPM - How is it used?

Application programming interface:

- create new keys (e.g. `CreateWrapKey`), and load them into the device (e.g. `LoadKey2`);
- manipulate these keys, and the PCRs
 - e.g. `UnBind` allows one to decrypt a ciphertext using a key that is stored into the TPM and locked to the current PCR value
 - e.g. `Quote` allows one to obtain a certificate attesting that a key is locked to a particular PCR value
 - e.g. `Extend` allows one to extend the current value of a PCR with some data x , i.e. $p := \text{SHA1}(p||x)$.

The TPM provides a **root of trust** for a variety of protocols: e.g. Microsoft's hard drive encryption system "BitLocker", Direct Anonymous Attestation protocol, ...

Several attempts to formally analyse the TPM itself

- using a theorem prover [Lin, 2005];
- using ProVerif [Delaune *et al.*, 2010]; or
- in some specific models (with no tool support), e.g. [Gürgens *et al.*, 2007, Coker *et al.*, 2010]

Several attempts to formally analyse the TPM itself

- using a theorem prover [Lin, 2005];
- using ProVerif [Delaune *et al.*, 2010]; or
- in some specific models (with no tool support), e.g. [Gürgens *et al.*, 2007, Coker *et al.*, 2010]

→ These results do *not* consider TPM state registers.

Several attempts to formally analyse the TPM itself

- using a theorem prover [Lin, 2005];
- using ProVerif [Delaune *et al.*, 2010]; or
- in some specific models (with no tool support), e.g. [Gürgens *et al.*, 2007, Coker *et al.*, 2010]

→ These results do *not* consider TPM state registers.

Modelling state is challenging

[Herzog, 2006]

- extension of the strand space model to analyse optimistic fair exchange protocol [Guttman, 2011]
- extension of ProVerif to take global state into account [Modersheim, 2010, Arapinis *et al.*, 2011]

Several attempts to formally analyse the TPM itself

- using a theorem prover [Lin, 2005];
- using ProVerif [Delaune *et al.*, 2010]; or
- in some specific models (with no tool support), e.g. [Gürgens *et al.*, 2007, Coker *et al.*, 2010]

→ These results do *not* consider TPM state registers.

Modelling state is challenging

[Herzog, 2006]

- extension of the strand space model to analyse optimistic fair exchange protocol [Guttman, 2011]
- extension of ProVerif to take global state into account [Modersheim, 2010, Arapinis *et al.*, 2011]

→ These results are *not* suitable to analyse protocols based on TPM state registers.

Formal analysis of protocols based on TPM registers using an automatic tool

Formal analysis of protocols based on TPM registers using an automatic tool

Our approach:

- we use Horn clauses and rely on the **ProVerif tool**;
- we solve **non-termination issues** by providing a transformation that is sound and complete for the class of k -stable clauses; and
- we provide a **syntactic criterion** to conclude to k -stability.

Formal analysis of protocols based on TPM registers using an automatic tool

Our approach:

- we use Horn clauses and rely on the **ProVerif tool**;
- we solve **non-termination issues** by providing a transformation that is sound and complete for the class of k -stable clauses; and
- we provide a **syntactic criterion** to conclude to k -stability.

Some case studies:

- a simplified version of the Microsoft BitLocker protocol
- a secure envelope protocol [Ables & Ryan, 2010]

→ both protocols crucially rely on the use of PCR

- 1 Overview of the TPM
- 2 Modelling using Horn clauses
- 3 Analysing with ProVerif
- 4 Case studies

- 1 Overview of the TPM
- 2 Modelling using Horn clauses
- 3 Analysing with ProVerif
- 4 Case studies

Cryptographic key

Keys are arranged in a **tree** structure and stored in the TPM memory
→ Storage Root Key created by a special command

Authdata, PCR

In particular, to each TPM key is associated an **authdata** value and also some **PCR** values

- **authdata** is a password shared between the user process and the TPM
- **PCR values** constrain the state of the TPM. The TPM will use the key only if certain PCRs currently have certain values.

CertifyKey command

Goal: allow a user to obtain a certificate on a key that is stored in the device.

CertifyKey command

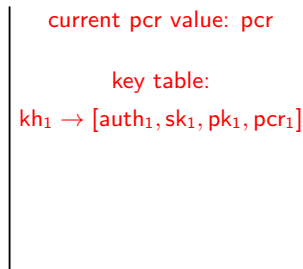
Goal: allow a user to obtain a certificate on a key that is stored in the device.

Description:

USER



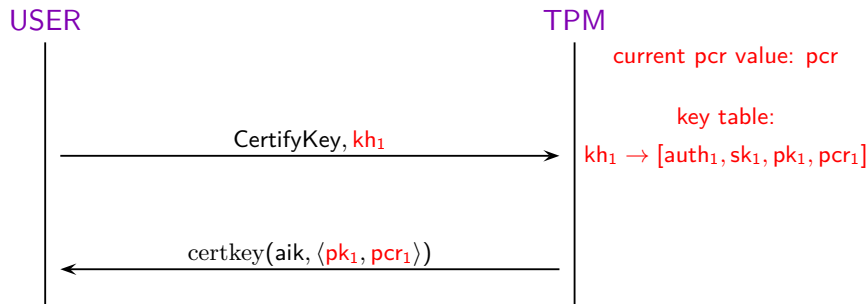
TPM



CertifyKey command

Goal: allow a user to obtain a certificate on a key that is stored in the device.

Description:



UnBind command

Goal: allow a user to retrieve the content of an encryption provided that the decryption key is stored in the key table of the TPM.

UnBind command

Goal: allow a user to retrieve the content of an encryption provided that the decryption key is stored in the key table of the TPM.

Description:

USER



TPM



current pcr value: **pcr₁**

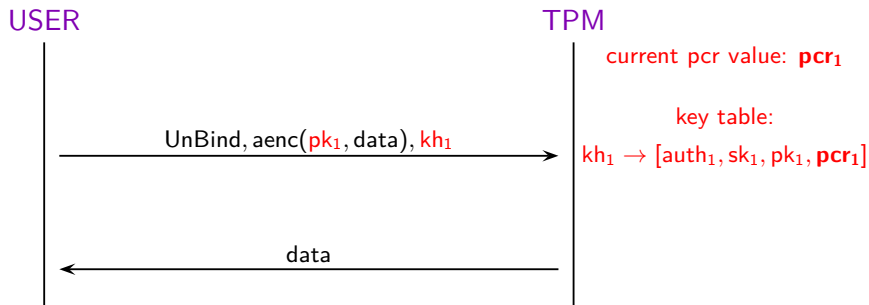
key table:

kh₁ → [auth₁, sk₁, pk₁, pcr₁]

UnBind command

Goal: allow a user to retrieve the content of an encryption provided that the decryption key is stored in the key table of the TPM.

Description:



Extend command

Goal: allow a user to update the value stored in one of the platform configuration register (PCR).

Extend command

Goal: allow a user to update the value stored in one of the platform configuration register (PCR).

Description:

USER



TPM




current pcr value: pcr

Extend command

Goal: allow a user to update the value stored in one of the platform configuration register (PCR).

Description:



- 1 Overview of the TPM
- 2 Modelling using Horn clauses
- 3 Analysing with ProVerif
- 4 Case studies

An introductory example

Goal: Alice has two secrets s_1 and s_2 . First, she interacts with Bob, and then Bob can learn one of the secrets (he chooses) but not both.

An introductory example

Goal: Alice has two secrets s_1 and s_2 . First, she interacts with Bob, and then Bob can learn one of the secrets (he chooses) but not both.

Description:

- 1 create and load a key pair $(k_1, pk(k_1))$ locked to $h(u_0, a_1)$ in Bob's TPM;
 - 2 create and load a key pair $(k_2, pk(k_2))$ locked to $h(u_0, a_2)$ in Bob's TPM;
- For sake of simplicity, we assume that the keys are already in Bob's TPM.

An introductory example

Goal: Alice has two secrets s_1 and s_2 . First, she interacts with Bob, and then Bob can learn one of the secrets (he chooses) but not both.

Description:

- 1 create and load a key pair $(k_1, pk(k_1))$ locked to $h(u_0, a_1)$ in Bob's TPM;
- 2 create and load a key pair $(k_2, pk(k_2))$ locked to $h(u_0, a_2)$ in Bob's TPM;
→ For sake of simplicity, we assume that the keys are already in Bob's TPM.
- 3 Bob provides some certificates to Alice (using **CertifyKey**);

An introductory example

Goal: Alice has two secrets s_1 and s_2 . First, she interacts with Bob, and then Bob can learn one of the secrets (he chooses) but not both.

Description:

- 1 create and load a key pair $(k_1, pk(k_1))$ locked to $h(u_0, a_1)$ in Bob's TPM;
- 2 create and load a key pair $(k_2, pk(k_2))$ locked to $h(u_0, a_2)$ in Bob's TPM;
→ For sake of simplicity, we assume that the keys are already in Bob's TPM.
- 3 Bob provides some certificates to Alice (using `CertifyKey`);
- 4 Alice sends $aenc(pk(k_1), s_1)$ and $aenc(pk(k_2), s_2)$ to Bob;

An introductory example

Goal: Alice has two secrets s_1 and s_2 . First, she interacts with Bob, and then Bob can learn one of the secrets (he chooses) but not both.

Description:

- 1 create and load a key pair $(k_1, pk(k_1))$ locked to $h(u_0, a_1)$ in Bob's TPM;
- 2 create and load a key pair $(k_2, pk(k_2))$ locked to $h(u_0, a_2)$ in Bob's TPM;
→ For sake of simplicity, we assume that the keys are already in Bob's TPM.
- 3 Bob provides some certificates to Alice (using **CertifyKey**);
- 4 Alice sends $aenc(pk(k_1), s_1)$ and $aenc(pk(k_2), s_2)$ to Bob;
- 5 Using **Extend** and **UnBind**, Bob can obtain either s_1 or s_2 , but not both.

Modelling the attacker

Predicate att

$\text{att}(u, v)$ means that there is a reachable state in which the PCR has value u and the attacker knows v .

Some rules:

$$\text{att}(x_p, x) \rightarrow \text{att}(x_p, \text{pk}(x))$$

$$\text{att}(x_p, x) \wedge \text{att}(x_p, y) \rightarrow \text{att}(x_p, \text{aenc}(x, y))$$

$$\text{att}(x_p, \text{aenc}(\text{pk}(x), y)) \wedge \text{att}(x_p, x) \rightarrow \text{att}(x_p, y)$$

Initial knowledge:

$$\text{att}(u_0, a_1)$$

$$\text{att}(u_0, a_2)$$

Predicate key

$\text{key}(u, sk, pk, v)$ means that there is a reachable state in which the PCR has value u , and the key table has an entry for the key pair (sk, pk) locked to the PCR value v .

Predicate key

$\text{key}(u, sk, pk, v)$ means that there is a reachable state in which the PCR has value u , and the key table has an entry for the key pair (sk, pk) locked to the PCR value v .

Some initial facts:

$$\text{key}(u_0, k_1, \text{pk}(k_1), h(u_0, a_1))$$
$$\text{key}(u_0, k_2, \text{pk}(k_2), h(u_0, a_2))$$

Predicate key

$\text{key}(u, sk, pk, v)$ means that there is a reachable state in which the PCR has value u , and the key table has an entry for the key pair (sk, pk) locked to the PCR value v .

Some initial facts:

$$\text{key}(u_0, k_1, \text{pk}(k_1), h(u_0, a_1))$$
$$\text{key}(u_0, k_2, \text{pk}(k_2), h(u_0, a_2))$$

Remarks:

- we do **not allow keys to be deleted** from the memory of the TPM;
→ we allow an unbounded number of keys to be loaded
- the attacker is not allowed to modify the key table (only through the API).

Modelling the TPM commands (1/2)

CertifyKey

$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \rightarrow \text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, x_{pcr} \rangle))$$

UnBind

$$\text{att}(x_p, \text{aenc}(x_{pk}, x_{data})) \wedge \text{key}(x_p, x_{sk}, x_{pk}, x_p) \rightarrow \text{att}(x_p, x_{data})$$

Modelling the TPM commands (2/2)

The TPM rule for **extending** and **rebooting the PCR** is treated in a particular way. We have a dedicated set of **inheritance rules**.

Modelling the TPM commands (2/2)

The TPM rule for **extending** and **rebooting the PCR** is treated in a particular way. We have a dedicated set of **inheritance rules**.

Key table:

$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \wedge \text{att}(x_p, x_v) \rightarrow \text{key}(h(x_p, x_v), x_{sk}, x_{pk}, x_{pcr})$$

$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \rightarrow \text{key}(u_0, x_{sk}, x_{pk}, x_{pcr}) \quad (\text{optional})$$

Modelling the TPM commands (2/2)

The TPM rule for **extending** and **rebooting the PCR** is treated in a particular way. We have a dedicated set of **inheritance rules**.

Key table:

$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \wedge \text{att}(x_p, x_v) \rightarrow \text{key}(h(x_p, x_v), x_{sk}, x_{pk}, x_{pcr})$$

$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \rightarrow \text{key}(u_0, x_{sk}, x_{pk}, x_{pcr}) \quad (\text{optional})$$

Knowledge of the attacker:

$$\text{att}(x_p, x_v) \wedge \text{att}(x_p, x) \rightarrow \text{att}(h(x_p, x_v), x)$$

$$\text{att}(x_p, x) \rightarrow \text{att}(u_0, x)$$

Modelling the protocol

Protocol rules:

Considering our introductory example, the **role of Alice** can be described by the following two rules:

$$\text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, \text{h}(u_0, a_1) \rangle)) \rightarrow \text{att}(x_p, \text{aenc}(x_{pk}, s_1))$$

$$\text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, \text{h}(u_0, a_2) \rangle)) \rightarrow \text{att}(x_p, \text{aenc}(x_{pk}, s_2))$$

Modelling the protocol

Protocol rules:

Considering our introductory example, the **role of Alice** can be described by the following two rules:

$$\text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, h(u_0, a_1) \rangle)) \rightarrow \text{att}(x_p, \text{aenc}(x_{pk}, s_1))$$

$$\text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, h(u_0, a_2) \rangle)) \rightarrow \text{att}(x_p, \text{aenc}(x_{pk}, s_2))$$

Query

Is Bob able to learn both secrets?

$$Q = \{\text{att}(x, s_1), \text{att}(x, s_2)\}$$

Outline

- 1 Overview of the TPM
- 2 Modelling using Horn clauses
- 3 Analysing with ProVerif**
- 4 Case studies

The ProVerif tool (B. Blanchet)

Available on line:

<http://www.proverif.ens.fr/>

Input: protocols written in Horn clauses

Characteristics

- **unbounded** number of sessions
- primitives given by an **equational theory**
- **security properties:** (strong) **secrecy**, correspondence properties, equivalence properties
- sound but not complete, **termination is not guaranteed**
→ the tool works well in practice

Termination problem

The termination problem seems due to the way PCR is modeled:

$$\text{att}(x_p, x_v) \wedge \text{att}(x_p, x) \rightarrow \text{att}(h(x_p, x_v), x)$$

$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \wedge \text{att}(x_p, x_v) \rightarrow \text{key}(h(x_p, x_v), x_{sk}, x_{pk}, x_{pcr})$$

Termination problem

The termination problem seems due to the way PCR is modeled:

$$\text{att}(x_p, x_v) \wedge \text{att}(x_p, x) \rightarrow \text{att}(h(x_p, x_v), x)$$

$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \wedge \text{att}(x_p, x_v) \rightarrow \text{key}(h(x_p, x_v), x_{sk}, x_{pk}, x_{pcr})$$

Main idea

- 1 Could we **bound** the length of the PCR, *i.e.* the number of times a PCR may be extended between two resets?
- 2 If the answer is 'yes', can we **compute** such a bound?

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.

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

- $\text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, h(u_0, a_1) \rangle)) \rightarrow \text{att}(x_p, \text{aenc}(x_{pk}, s_1))$
- $\text{att}(x_p, x_v) \wedge \text{att}(x_p, x) \rightarrow \text{att}(h(x_p, x_v), x)$

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

- $\text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, h(u_0, a_1) \rangle)) \rightarrow \text{att}(x_p, \text{aenc}(x_{pk}, s_1))$
- $\text{att}(x_p, x_v) \wedge \text{att}(x_p, x) \rightarrow \text{att}(h(x_p, x_v), x)$

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 $h(v_1, v_2) \in st(R)$, $\text{length}_{\text{pcr}}(v_1, v_2) \leq k$;
- 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 R is k -stable**.

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 $h(v_1, v_2) \in st(R)$, $\text{length}_{\text{pcr}}(v_1, v_2) \leq k$;
- 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 R is k -stable**.

Examples

- $\text{att}(x_p, \text{certkey}(\text{aik}, \langle x_{pk}, h(u_0, a_1) \rangle)) \rightarrow \text{att}(x_p, \text{aenc}(x_{pk}, s_1))$
- $\text{att}(x_p, x_v) \wedge \text{att}(x_p, x) \rightarrow \text{att}(h(x_p, x_v), x)$

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 R is k -stable**.

Examples

- $att(x_p, certkey(aik, \langle x_{pk}, h(u_0, a_1) \rangle)) \rightarrow att(x_p, aenc(x_{pk}, s_1))$
- $att(x_p, x_v) \wedge att(x_p, x) \rightarrow att(h(x_p, x_v), x)$

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

Our transformation

Goal: A set of *k-stable* rules can be transformed into another “equivalent” set of rules that is *more suitable* for analysis with ProVerif.

Our transformation

Goal: A set of *k-stable* rules can be transformed into another “equivalent” set of rules that is *more suitable* for analysis with ProVerif.

Transformation: we replace each rule R by the set of rules:

$$\{R[x \mapsto u] \mid x \in \mathcal{X}, p(x, t_1, \dots, t_\ell) \in R, u \in U_k\}$$

$$\text{where } U_k = \{ \begin{array}{l} u_0, \\ h(u_0, x_1), \\ \dots, \\ h(\dots h(u_0, x_1), \dots, x_k) \end{array} \}.$$

Our transformation

Goal: A set of *k-stable* rules can be transformed into another “equivalent” set of rules that is *more suitable* for analysis with ProVerif.

Transformation: we replace each rule R by the set of rules:

$$\{R[x \mapsto u] \mid x \in \mathcal{X}, p(x, t_1, \dots, t_\ell) \in R, u \in U_k\}$$

$$\text{where } U_k = \{ \begin{array}{l} u_0, \\ h(u_0, x_1), \\ \dots, \\ h(\dots h(u_0, x_1), \dots, x_k) \end{array} \}.$$

This transformation effectively bounds the PCR length of possible PCR values that may appear as the first argument of a predicate.

Theorem

If the initial set of rules is *k-stable*, then the initial and transformed set of rules are equivalent w.r.t. satisfiability of queries.

Outline

- 1 Overview of the TPM
- 2 Modelling using Horn clauses
- 3 Analysing with ProVerif
- 4 Case studies

TPM commands

TPM's commands – We consider the following commands.

- Read
- Quote
- CreateWrapKey
- LoadKey2
- CertifyKey
- UnBind
- Seal
- UnSeal

TPM commands

TPM's commands – We consider the following commands.

- Read
- Quote
- CreateWrapKey
- LoadKey2
- CertifyKey
- UnBind
- Seal
- UnSeal

Simplifications and/or abstractions

TPM commands

TPM's commands – We consider the following commands.

- Read
- Quote
- CreateWrapKey
- LoadKey2
- CertifyKey
- UnBind
- Seal
- UnSeal

Simplifications and/or abstractions

- 1 we **do not** consider **authdata**;
→ this is equivalent to giving all the authdata to the attacker

TPM commands

TPM's commands – We consider the following commands.

- Read
- Quote
- CreateWrapKey
- LoadKey2
- CertifyKey
- UnBind
- Seal
- UnSeal

Simplifications and/or abstractions

- 1 we **do not** consider **authdata**;
→ this is equivalent to giving all the authdata to the attacker
- 2 the key **AIK** (attestation identity key) is initially and permanently loaded in the TPM;
→ In reality, we have to create it (**Makeldentity**) and to load it (**Activateldentity**)

TPM commands

TPM's commands – We consider the following commands.

- Read
- Quote
- CreateWrapKey
- LoadKey2
- CertifyKey
- UnBind
- Seal
- UnSeal

Simplifications and/or abstractions

- 1 we **do not** consider **authdata**;
→ this is equivalent to giving all the authdata to the attacker
- 2 the key **AIK** (attestation identity key) is initially and permanently loaded in the TPM;
→ In reality, we have to create it (**Makeldentity**) and to load it (**Activateldentity**)
- 3 we only consider **one** PCR, instead of **24**.

A simplified version of the Bitlocker protocol (1/2)

Goal: protect the data that are stored on your disk.

→ your data are encrypted using **VEK**, which is in turn encrypted with **VMK**.

A simplified version of the Bitlocker protocol (1/2)

Goal: protect the data that are stored on your disk.

→ your data are encrypted using **VEK**, which is in turn encrypted with **VMK**.

Description of the set-up phase:

- A new key pair (sk, pk) is generated and loaded in Alice's TPM
→ using **CreateWrapKey** and **LoadKey2**;
- **VMK** is encrypted under the key **pk** locked to $h(h(u_0, bios), loader)$
→ using **Seal**

$seal(pk, vmk, tpmproof, h(h(u_0, bios), loader))$

A simplified version of the Bitlocker protocol (1/2)

Goal: protect the data that are stored on your disk.

→ your data are encrypted using **VEK**, which is in turn encrypted with **VMK**.

Description of the set-up phase:

- A new key pair (sk, pk) is generated and loaded in Alice's TPM
→ using **CreateWrapKey** and **LoadKey2**;
- **VMK** is encrypted under the key **pk** locked to $h(h(u_0, bios), loader)$
→ using **Seal**

$seal(pk, vmk, tpmproof, h(h(u_0, bios), loader))$

Description of the retrieval phase:

- a **trust chain** is built: Pre-BIOS → BIOS → loader
- retrieve **VMK** using **Unseal**
- prevent unauthorised retrievals, by extending “**deny**” into the PCR

Modelling - Bitlocker protocol (2/2)

Alice's role setting up the drive encryption in a trusted state:

$$\text{key}(x_p, x_{sk}, \text{pk}(x_{sk}), \text{nil}) \rightarrow \text{att}(x_p, \text{seal}(\text{pk}(x_{sk}), \text{vmk}[x_p], \text{tpmproof}, \text{h}(\text{h}(u_0, \text{bios}), \text{loader})))$$

PCR reboot rules:

$$\text{att}(x_p, x) \rightarrow \text{att}(\text{h}(\text{h}(\text{h}(u_0, \text{bios}), \text{loader}), \text{deny}), x)$$

$$\text{att}(x_p, x) \rightarrow \text{att}(\text{h}(\text{h}(u_0, \text{bios}), \text{loader_rogue}), x)$$

$$\text{att}(x_p, x) \rightarrow \text{att}(\text{h}(u_0, \text{bios_rogue}), x)$$

Results of our analysis: $\text{att}(x_p, \text{vmk}[x])$

- the rules are **3-stable**
- ProVerif quickly concludes that the protocol is **safe** (using the set of rules obtained by applying our transformation).

Goal: provide some data (**secret**) to Bob in such a way that Bob can either access the data or revoke his right to access the data.

→ Now, we consider the fact that the TPM can be **rebooted**.

Description

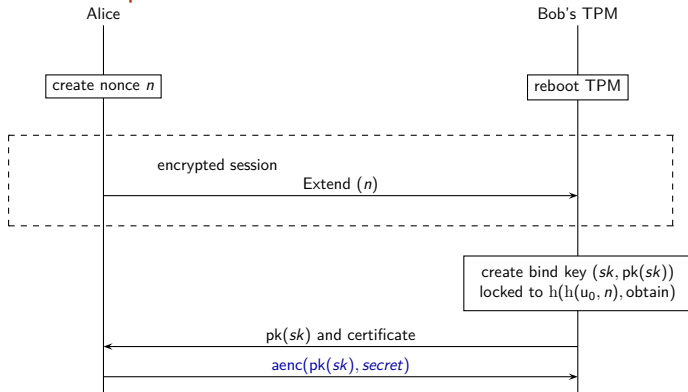
- 1 Sealing the envelope:

Goal: provide some data (**secret**) to Bob in such a way that Bob can either access the data or revoke his right to access the data.

→ Now, we consider the fact that the TPM can be **rebooted**.

Description

1 Sealing the envelope:



Goal: provide some data (*secret*) to Bob in such a way that Bob can either access the data or revoke his right to access the data.

→ Now, we consider the fact that the TPM can be *rebooted*.

Description

① *Sealing the envelope:*

② *Opening the envelope:*

→ use *Extend* to extend *obtain* into the PCR,

→ use *UnBind* to decrypt the ciphertext $\text{aenc}(\text{pk}(sk), \text{secret})$;

Goal: provide some data (**secret**) to Bob in such a way that Bob can either access the data or revoke his right to access the data.

→ Now, we consider the fact that the TPM can be **rebooted**.

Description

① **Sealing the envelope:**

② **Opening the envelope:**

→ use **Extend** to extend **obtain** into the PCR,

→ use **UnBind** to decrypt the ciphertext $\text{aenc}(\text{pk}(sk), \text{secret})$;

③ **Returning the envelope:**

→ use **Extend** to extend **deny** into the PCR,

→ use **Quote** to obtain a signature attesting that the current value of the PCR is $h(h(u_0, n), \text{deny})$. This certificate can be used as a proof that Bob will never have access to **secret**.

Envelope protocol (2/3)

Alice's role

$\text{att}(x_p, x) \rightarrow \text{att}(\text{h}(x_p, n[x_p]), x)$

$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \rightarrow \text{key}(\text{h}(x_p, n[x_p]), x_{sk}, x_{pk}, x_{pcr})$

$\text{att}(x_p, \text{certkey}(\text{aik}, \text{pk}(\text{sk}), \text{h}(\text{h}(u_0, n[y]), \text{obtain})))$
 $\rightarrow \text{att}(x_p, \text{aenc}(\text{pk}(\text{sk}), \text{secret}[y]))$

Envelope protocol (2/3)

Alice's role

$$\text{att}(x_p, x) \rightarrow \text{att}(\mathbf{h}(x_p, n[x_p]), x)$$
$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \rightarrow \text{key}(\mathbf{h}(x_p, n[x_p]), x_{sk}, x_{pk}, x_{pcr})$$
$$\begin{aligned} \text{att}(x_p, \text{certkey}(\text{aik}, \text{pk}(\text{sk}), \mathbf{h}(\mathbf{h}(u_0, n[y]), \text{obtain}))) \\ \rightarrow \text{att}(x_p, \text{aenc}(\text{pk}(\text{sk}), \text{secret}[y])) \end{aligned}$$

Query

- $\text{att}(x_p, \text{secret}[y])$, and
- $\text{att}(x_p, \text{certpcr}(\text{aik}, \mathbf{h}(\mathbf{h}(u_0, n[y]), \text{deny}), x))$.

Envelope protocol (2/3)

Alice's role

$$\text{att}(x_p, x) \rightarrow \text{att}(\mathbf{h}(x_p, n[x_p]), x)$$
$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \rightarrow \text{key}(\mathbf{h}(x_p, n[x_p]), x_{sk}, x_{pk}, x_{pcr})$$
$$\begin{aligned} \text{att}(x_p, \text{certkey}(\text{aik}, \text{pk}(\text{sk}), \mathbf{h}(\mathbf{h}(u_0, n[y]), \text{obtain}))) \\ \rightarrow \text{att}(x_p, \text{aenc}(\text{pk}(\text{sk}), \text{secret}[y])) \end{aligned}$$

Query

- $\text{att}(x_p, \text{secret}[y])$, and
- $\text{att}(x_p, \text{certpcr}(\text{aik}, \mathbf{h}(\mathbf{h}(u_0, n[y]), \text{deny}), x))$.

All the rules are **2-stable** and **ProVerif terminates** on the set of rules obtained after applying our transformation.

Envelope protocol (2/3)

Alice's role

$$\text{att}(x_p, x) \rightarrow \text{att}(\mathbf{h}(x_p, n[x_p]), x)$$
$$\text{key}(x_p, x_{sk}, x_{pk}, x_{pcr}) \rightarrow \text{key}(\mathbf{h}(x_p, n[x_p]), x_{sk}, x_{pk}, x_{pcr})$$
$$\begin{aligned} \text{att}(x_p, \text{certkey}(\text{aik}, \text{pk}(\text{sk}), \mathbf{h}(\mathbf{h}(u_0, n[y]), \text{obtain}))) \\ \rightarrow \text{att}(x_p, \text{aenc}(\text{pk}(\text{sk}), \text{secret}[y])) \end{aligned}$$

Query

- $\text{att}(x_p, \text{secret}[y])$, and
- $\text{att}(x_p, \text{certpcr}(\text{aik}, \mathbf{h}(\mathbf{h}(u_0, n[y]), \text{deny}), x))$.

All the rules are **2-stable** and **ProVerif terminates** on the set of rules obtained after applying our transformation.

→ **false attack** due to the nonce abstraction.

Envelope protocol (3/3)

→ **Add freshness** by adding an additional boot parameter to the att and key predicates.

$$\text{att}(x_b, x_p, x) \rightarrow \text{att}(x_b, h(x_p, n[x_b]), x)$$

...

Envelope protocol (3/3)

→ **Add freshness** by adding an additional boot parameter to the att and key predicates.

$$\text{att}(x_b, x_p, x) \rightarrow \text{att}(x_b, h(x_p, n[x_b]), x)$$

...

PCR reboot rules:

$$\begin{aligned} \text{att}(x_b, x_p, x) &\rightarrow \text{att}(b(x_b, x_p), u_0, x) \\ \text{key}(x_b, x_p, \text{srk}, \text{pk}(\text{srk}), \text{nil}) &\rightarrow \text{key}(b(x_b, x_p), u_0, \text{srk}, \text{pk}(\text{srk}), \text{nil}) \\ \text{key}(x_b, x_p, \text{aik}, \text{pk}(\text{aik}), \text{nil}) &\rightarrow \text{key}(b(x_b, x_p), u_0, \text{aik}, \text{pk}(\text{aik}), \text{nil}) \end{aligned}$$

Envelope protocol (3/3)

→ **Add freshness** by adding an additional boot parameter to the att and key predicates.

$$\text{att}(x_b, x_p, x) \rightarrow \text{att}(x_b, h(x_p, n[x_b]), x)$$

...

PCR reboot rules:

$$\begin{aligned} \text{att}(x_b, x_p, x) &\rightarrow \text{att}(b(x_b, x_p), u_0, x) \\ \text{key}(x_b, x_p, \text{srk}, \text{pk}(\text{srk}), \text{nil}) &\rightarrow \text{key}(b(x_b, x_p), u_0, \text{srk}, \text{pk}(\text{srk}), \text{nil}) \\ \text{key}(x_b, x_p, \text{aik}, \text{pk}(\text{aik}), \text{nil}) &\rightarrow \text{key}(b(x_b, x_p), u_0, \text{aik}, \text{pk}(\text{aik}), \text{nil}) \end{aligned}$$

Result of our analysis:

→ Due to the boot parameter, ProVerif encounters **termination problems**.

→ ProVerif confirms that the protocol is secure (around 30 min) for 1 reboot.

Formal Horn clauses-based framework for modelling protocols of the TPM that use PCRs.

Our method:

- 1 model everything using Horn clauses;
- 2 show that the set of clauses needed are k -stable, and apply our attack-preserving transformation;
- 3 launch ProVerif (or another tool) on the resulting set of clauses.

Conclusion

Formal Horn clauses-based framework for modelling protocols of the TPM that use PCRs.

Our method:

- 1 model everything using Horn clauses;
- 2 show that the set of clauses needed are k -stable, and apply our attack-preserving transformation;
- 3 launch ProVerif (or another tool) on the resulting set of clauses.

Case studies:

- 1 A simplified version of the **Microsoft Bitlocker protocol**
→ rules are 3-stable, ProVerif quickly concludes on the resulting set of rules, the VMK remains secret for unbounded reboots and PCR extends
- 2 The envelope protocol [Ables & Ryan, 10]
→ add freshness through the boot parameter to avoid a false attack, rules are 2-stable, ProVerif concludes for one reboot.