

Automated Backward Analysis of PKCS#11 v2.20

Robert Künnemann

Department of Computer Science – TU Darmstadt, Germany

Abstract. The PKCS#11 standard describes an API for cryptographic operations which is used in scenarios where cryptographic secrets need to be kept secret, even in case of server compromise. It is widely deployed and supported by many hardware security modules and smart cards. A variety of attacks in the literature illustrate the importance of a careful configuration, as API-level attacks may otherwise extract keys.

Formal verification of PKCS#11 configurations requires the analysis of a system that contains mutable state, a problem that existing methods solved by either artificially restricting the number of keys, introducing model-specific over-approximation or performing proofs by hand. At Security & Privacy 2014, Kremer and Künnemann presented a variant of the applied pi calculus that handles global state and, in conjunction with the tamarin prover for protocol verification, allows for the precise analysis of protocols with state. Using this tool chain, we show secrecy of keys for a PKCS#11 configuration that makes use of features introduced in version 2.20 of the standard, including wrap and unwrap templates in an extensible model.

This configuration supports the creation of so-called wrapping keys for import and export of sensitive keys (e. g., for backup or transfer), and it permits the co-existence of sensitive keys and non-sensitive keys on the same device.

1 Introduction

The more complex a system is, the more difficult is assuring its security. Given the complexity of the runtime environment provided by multi-purpose computers, it appears reasonable to compute security-critical operations outside the computer that actually runs the protocol (and usually more than one protocol), and instead on a device which is *a*) smaller, and thus more amenable to verification, and *b*) designed with security in mind. In security-critical contexts, e. g., the cash machine network or the public key infrastructure, the use of such devices, so-called *security tokens*, is common practice. In case the (complex) system running the protocol is under adversarial control, i. e., in case of server compromise, cryptographic secrets are protected by the (smaller, and thus more secure) security token, and the fact that these secrets were never directly accessible, even to the server.

PKCS#11 defines a platform-independent API to security tokens, for example smart cards, but also hardware security modules (HSMs). HSMs are physical computing devices that can be attached directly to a server via Ethernet, USB or other services, providing logical and physical protection for sensitive information and algorithms.

The fundamental security feature this standard provides is protecting the cryptographic values of keys marked as “sensitive” [22, Section 7]:

Additional protection can be given to private keys and secret keys by marking them as “sensitive” [...]. Sensitive keys cannot be revealed in plaintext off the token [...].

Clulow’s attack. It was discovered very soon that a faithful implementation of the standard violates this property. Clulow showed the following attack in 2003 [8], which serves as an introduction to the caveats in the design of PKCS#11. Keys are accessed via handles, hence indirectly. Some may be used for encryption and/or decryption, therefore an attacker that can access the token (e. g., in case of server compromise) can request the encryption of some message he supplies with a key of his choice, given that he knows the handle (which we will consider public) and that the key is configured to allow for encryption by setting an attribute ‘encrypt’ to true, analogously for decryption. Similarly, it is possible to *wrap* a key k_1 with another key k_2 , that is, encrypt k_1 under k_2 . This allows for backup, as well as for transfer. Given that k_2 is present on two security tokens A and B , one may wrap k_1 with k_2 on A (using two handles associated with k_1 and k_2 , respectively) and *unwrap* the resulting cypher-text on B , using another handle to k_2 . The scenario where k_2 is configured for wrapping *and* for decryption permits the attacker to request a wrapping of k_1 under k_2 and then decrypt the resulting cyphertext, thus obtaining the value of k_1 in the plain. The attack also applies when k_2 is wrapped under itself.

An ‘incomplete’ implementation of the standard, often called a *configuration* can thwart this attack by forbidding the same key to be used for wrapping and decryption. But there are other conflicts, e. g., between encryption and unwrapping. In the present paper, we propose a configuration and a method for verifying its security. We focus on attacks on the logical level, using only API calls that are (by themselves) perfectly harmless, as opposed to attacks on the implementation of cryptographic functionality [4].

Related Work. Building on the work of Longley and Rigby [20] and Bond and Anderson [5] on API attacks, several recent papers have investigated the security of APIs on the logical level adapting symbolic techniques for protocol analysis [6, 9, 11], finding many new attacks. There have also been academic proposals for new APIs [19, 10, 18]. While many attacks were found, a lot of effort was directed towards finding configurations that are secure, i. e., that preserve secrecy of keys.

In the analysis of PKCS#11 configurations, there are three major lines of work. The first one uses protocol verification techniques, regarding the security token as the (sole) participant in a protocol, with the adversary sending requests on the public network. Early results by Delaune, Kremer and Steel translated a given configuration into a satisfiability problem which is solved by model checking, providing secrecy of keys if the number of keys is known in advance [11]. This restriction was lifted in later work [14] by Fröschle and Steel, showing that a class of configurations can safely be abstracted by configurations that are *static* (i. e., a key’s attribute cannot be changed) and showing that the latter is soundly over-approximated in a bounded model. This method is used in further work by Bortolozzo, Centenaro, Focardi and Steel to find attacks on security tokens using automatic reverse engineering and to show a configuration very similar to ours secure.

The second line of work uses a program verification approach, modelling security tokens in a first-order linear time logic with past operators. In contrast to the protocol verification approach, proofs have been conducted by hand (using a tableau proof method that proceeds by backward analysis), but provide support for advanced data structures introduced by PKCS#11, version 2.20, in particular wrap/unwrap templates. An *attribute template* is a set of attributes, but contains the attributes ‘wrapping template’ and ‘unwrapping template’ which themselves are pointers to attribute templates, resulting in a recursive data structure. In early work, Fröschele and Sommer show security, but only for keys with ‘extractable’ set to false [15], i. e., keys that cannot be wrapped at all. In more recent work, the authors added support for wrap/unwrap templates [13]. Positives result only cover the secrecy of trusted wrapping keys, but not of keys wrapped using these keys (as opposed to the results in the present paper, and in the work by Bortolozzo et al.). The proof is done by hand and covers about six pages [16].

The third line of work uses static analysis on the implementation of the security token [7]. Using security type-checking, a software implementation (written in a subset of C) of the wrap, unwrap, encrypt and decrypt functions was shown to preserve secrecy of sensitive keys against a Dolev-Yao attacker for a configuration that is functionally equal to the one presented here, but uses a default type for wrapped keys rather than wrap/unwrap templates. A generalised version of this framework proposes an imperative programming language with cryptographic operations and a centralised store mapping values (possibly handles) to pairs of key values and their ground types [2]. Mapping sets of PKCS#11 attributes to types, the generalised framework again allows for the analysis of PKCS#11, more specifically, PKCS#11 v2.20 using the configuration analysed in the present paper.

Contribution. The present work shares common ground with the first two lines of work. It follows a protocol verification approach and provides machine support, but relies on backward analysis rather than finite model checking. Therefore, the resulting model can be extended without breaking the soundness of the decision procedure (or introducing limitations to the number of keys), but provides a largely automated proof procedure although decidability cannot be guaranteed. The model is formulated in a variant of the applied pi calculus augmented with operators for state manipulation. This high-level protocol description language can be translated to multiset rewrite rules using the Sapic tool which has been proven sound and complete in prior work [17]. The generated multiset rewrite system can be verified using the tamarin prover [23], which is sound and complete as well. The constraint solving algorithm employed by the tamarin prover uses backward analysis. While various methods in the first line of work achieve security for an unbounded number of keys indirectly (using bounded analysis and an over-approximation specific to the modelling of the PKCS#11 API) [14, 6], this method supports an unlimited number of fresh values by default. For one, our model is closer to an actual implementation, containing locking operations as well as database lookup, but most importantly, our result requires no additional over-approximation. Hence, the security proof is machine-verifiable (relying on the correctness of the translation procedure [17] and the solving algorithm [23]) and the model extensible. The same holds for the third line of work, too (except that parallelism and hence locking is not supported). The type-checking approach is much faster and requires less manual intervention than

our approach, but might not be able to type configurations that are actually secure. However, the configuration we show secure has been verified using type-checking [7, 2], hence the question of whether there actually is a more versatile policy that can be shown secure using backward analysis but is not amenable to security type-checking remains open.

The flexibility and precision of our modelling comes at the cost of losing some amount of automation. As the verification method is sound and complete, but is able to express the secrecy problem with unbounded nonces, which is undecidable [12], there is no guarantee that the tool terminates. In order to achieve termination, some intervention is necessary: Lemmas need to be used to cut branches in the proof attempt. The advantage is a result that applies to a version of PKCS#11 (2.20) with wrap/unwrap templates and an intuitive, extensible model. In the present proof, we have given 10 lemmas, 6 of which are at least partially guided when being used without any model-specific heuristics for the choice of goals in the proof. If we provide a (model-specific) heuristic, a proof is found without user intervention. The lemmas are model-specific, but not specific to the configuration we show secure (except for four trivial lemmas that speed up the proof, but could be left out). We note that the tool chain is sound and complete no matter which heuristic is used, i. e., even when the model is altered (e. g., to describe future versions of PKCS#11), if the heuristics are helpful in finding a proof, the result is correct.

The configuration of PKCS#11 we prove secure implements a policy with three kinds of keys: keys that are used to encrypt or decrypt payload and kept secret, keys that can wrap and unwrap the first kind of keys for backup and keys that can be read in the clear and are neither wrapped, nor used to wrap other keys but may encrypt/decrypt payload. Keys cannot change roles, which we consider a sensible best practice, thus this policy is static in the sense of [14].

This policy has been shown to provide secrecy of sensitive keys if implemented via restricting the wrap and unwrap commands [7] or via wrap/unwrap templates [2] using type-checking. As in these works, we show that non-sensitive keys (keys the attacker is allowed to read in the clear) do not impair the security of other keys.

2 The PKCS#11 Standard, v2.20

The PKCS#11 standard specifies a security API (this term was coined by Bond and Anderson [5]), that is, an API that separates trusted code operating on secret data from untrusted code. As in the case of PKCS#11, security APIs are often used to perform cryptographic computations. Separating the implementation of cryptographic operations from the rest of the system has the advantage that cryptographic secrets can be hidden *behind* the API, whilst the code only accesses these secrets indirectly, using handles. If the API gives access to an external piece of hardware, which is often the case, the hope is that malicious code is restricted to using the API. Smart cards and hardware security modules (HSMs) implementing PKCS#11 are much simpler than multi-purpose computers, and designed with security in mind, so the idea of “outsourcing” sensitive information and operations that depend on this information to a device that is easier to secure seems to be reasonable. However, a sound design of said API is essential in reaching this goal.

Table 1. Attributes relevant to key-management. Modifiable means that attributes may be modified after an object is created, or while it is copied, however, tokens may decide not to permit modification upon copying [22, Tab. 15, footnote 8, p. 66]. SO-only means the attribute can only be set by the SO.

name	modifiable	SO-only	comment
wrap (<i>wrap</i>)	yes	no	can be first argument for <i>wrap</i>
unwrap (<i>unwrap</i>)	yes	no	can be first argument for <i>unwrap</i>
encrypt (<i>enc</i>)	yes	no	can be used for encryption
decrypt (<i>dec</i>)	yes	no	can be used for decryption
sensitive (<i>sens</i>)	yes	no	value shall not be extracted (directly), but may be wrapped
extractable (<i>extr</i>)	yes	no	Value shall not be extracted (directly) or used as second argument to <i>wrap</i>
trusted (<i>trus</i>)	no	yes	has been generated by SO
wrap-with-trusted (<i>wwt</i>)	no	no	can only be wrapped by ‘trusted’ keys
wrap template (<i>wt</i>)	no	no	key that are wrapped need to match this template
unwrap template (<i>ut</i>)	no	no	template applied to keys after being imported by <i>unwrap</i>

PKCS#11 gives multiple applications access to several cryptographic devices via *slots* [22, p.14], abstracting away from their respective implementation technology, be it smart cards, PCMCIA cards, HSMs, etc. After a user, or an application, has established a *session* to a device (through a slot), they can identify either as a Security Officer (SO), or a normal user [22, p.15]. The role of the SO is to initialize a token and to set the normal user’s PIN. The normal user cannot log in until the SO has set the normal user’s PIN. Within a session, the user can manipulate *objects* stored on the token, such as keys and certificates. Objects are referred to by *handles*. The value of a handle does not reveal any information about the value of the key. Objects may be marked public and private. A normal user, if authenticated, can access public and private objects, otherwise only public objects. The SO can only access public objects, but perform operations the normal user cannot perform, such as setting the user’s PIN.

Attributes. Objects may have attributes (besides being public or private), some of them specific to their class and type (public keys of type *CKK_RSA* have a public exponent, for example), and some general. The latter pertain to the key-management functions of PKCS#11, and are listed in Table 1. Real devices offer only a subset of the functionality specified by PKCS#11, partly due to security considerations: The PKCS#11 standards permits modifying the attributes *wrap* and *dec*, immediately giving rise to Clulow’s attack, see p.220.

2.1 Modelling

Security APIs are a means to provide a higher level of assurance in case of server compromise. If the server is compromised, the user PIN can easily be intercepted and the attacker can establish a session with the device. Hence we model a network comprising

only of a single PKCS#11 token, which does not perform any kind of authentication for normal users. The Dolev-Yao adversary is in full control over the network and is thus able to issue arbitrary requests. The PKCS#11 standard discusses security against server compromise and states that “none of the [these attacks] can compromise keys marked ‘sensitive’, since a key that is sensitive will always remain sensitive” [22, p. 31]. This does not hold true for PKCS#11 itself (cf. p. 220), but we will show a configuration of PKCS#11, for which this holds true.

This paper describes a formal model of the core key-management functionality in PKCS#11, therefore we left out message digesting functions, signature and MAC functions, as well as dual-purpose cryptographic functions (like `C_DigestEncryptUpdate`) and random number generation, as they do not pertain to the key and object management part of PKCS#11. Key-derivation (`C_DeriveKey`) allows for creating a new key object from a base key and could be considered key-management, but there is no cryptographic mechanism in PKCS#11 that supports wrap/unwrap as well as key derivation [22, Section 12]. Thus we consider this function outside the key-management core of PKCS#11. Note that encryption and decryption functions may allow for producing, or decrypting wrappings (as Clulow’s attack shows), so they do pertain to the key-management part. We concentrate on encryption and decryption for single-part data, i. e., the functions for multi-part data `C_EncryptInit`, `C_EncryptUpdate`, `C_EncryptFinal` are left out in favour of `C_Encrypt`, as our focus lies on API-level attacks. The same holds for decryption functions. We furthermore ignored object management functions that reveal only information which the attacker, as the only user accessing the token, can compute anyway, such as `C_GetObjectSize`, `C_GetAttributeValue`, `C_FindObjectsInit`, `C_FindObjects` and `C_FindObjectsFinal`. The function `C_DestroyObject` is disabled, since our modelling gives the adversary an unlimited amount of space to store keys, thus any attack can be transformed into an attack that does not delete objects. Previous works

Table 2. Object- and Key-management operations in PKCS#11 and our modelling. Function marked ‘(n.p.)’ are not present.

function	description	process	comment
<i>Object management functions</i>			
<code>C_CreateObject</code>	creates an object	(n.p.)	forbidden by policy
<code>C_CopyObject</code>	creates a copy of an object	(n.p.)	not helpful to adversary (in this configuration)
<code>C_DestroyObject</code>	destroys an object	(n.p.)	not helpful to adversary
<code>C_SetAttributeValue</code>	modifies object’s attribute	(n.p.)	forbidden by policy
<i>Encryption/Decryption functions</i>			
<code>C_Encrypt</code>	encrypts single-part data	encrypt	
<code>C_Decrypt</code>	decrypts single-part data	decrypt	
<i>Key-management functions</i>			
<code>C_GenerateKey</code>	generates a secret key	create	
<code>C_GenerateKeyPair</code>	generates a public / private key pair	(n.p.)	asymm. wrapping keys permits ‘Trojan wrapped key attack’
<code>C_WrapKey</code>	wraps (encrypts) a key	wrap	
<code>C_UnwrapKey</code>	unwraps (decrypts) a key	unwrap	

[11, 14, 6, 15] have similar restrictions. The full list of functions is listed in Table 8 of the PKCS#11 standard [22, p. 27]. The rest of PKCS#11 is what we consider the core key-management part (Table 2). The configuration, while being more versatile than policies that were previously proposed, forbids certain operations that are potentially harmful. Another class of attacks Clulow discovered are “Trojan wrapped key attacks”, where a public unwrapping key is used to introduce a wrapping key that the attacker knows the value of, by producing the wrapping outside the device. We conclude that our policy should not allow for asymmetric wrapping keys at all, and simplify the model by only regarding symmetric keys. Extending the model to cover asymmetric keys for non-key-management operations is straight-forward, but unlikely to lead to new insights with respect to the security of policies.

2.2 Proposed Configuration

In this section, we will discuss policies that have been proposed in previous work and introduce requirements that guided the design of the policy we will prove secure. We will argue that, given these requirements, there are only two secure policies.

The first requirement shall be that there are ‘usage’ keys, i. e., keys that may encrypt or sign payload. Since only encryption and decryption are relevant to key-management, we propose:

Requirement A: There should be keys that can be used for encryption and decryption.

Delaune et al. have shown different policies secure, all of which have in common that wrap and unwrap are conflicting attributes [11]. Since wrapping keys are useful for backup and out-of-device-storage of keys, we could formulate the requirement that wrapping keys should be suitable for wrapping and unwrapping any key, however, we think that the backup, and possibly the distribution of ‘usage’ keys – in our model, keys used for encryption and decryption – is the main purpose of wrapping keys. Being able to wrap wrapping keys is useful, but only once it is established that usage keys can be wrapped.

Requirement B: Wrapping keys should be able to wrap and unwrap keys that can be used for encryption and decryption.

Policies in related work often support the setting and unsetting of some attributes [15, 13], requiring attributes to be unmodifiable where needed. Many configurations describe policies by declaring conflicting attributes (pairs of attributes that are forbidden to be set at once) and sticky (unmodifiable when set) attributes. We argue for a different approach: security policies become much easier to understand and to design if the key is assigned a role upon creation, and cannot be altered in subsequent steps. In most cases, it is clear from the beginning which keys needs to be protected, setting a key to sensitive after it was ready to be exposed is clearly bad practice. The attributes wrap, unwrap, enc, dec, extr, trust have in common that an attacker becomes only less powerful in unsetting them. The opposite holds for wwt. Being able to alter wrapping/unwrapping templates is much less useful as it seems, as wrappings in

PKCS#11 carry no information about which wrapping template was used. Altering the wrapping template could be useful to define the class of keys that can be wrapped as the union of permitted wrapping templates but, as we will see, this class can be defined using a single wrapping template.

Requirement C: Policies should disable `C_SetAttributeValue` altogether.

Bortolozzo et al. proposed a policy that is static in that attributes cannot be altered after the creation of keys [6]. Ignoring templates where `wrap`, `unwrap`, `enc` and `dec` are false, they propose three templates:

1. Freshly generated keys can have `wrap` and `unwrap` set.
2. Freshly generated keys can have `enc` and `dec` set.
3. Keys imported with `unwrap` may have `unwrap` and `encrypt` set, but `wrap` and `decrypt` unset.

Intuitively, the third template means that keys ‘degrade’ when being backed up, that is, they cannot resume to their full functionality. The policy we analyse here differs in this regard: ‘trusted’ keys can be used to wrap ‘usage’ keys, which can be adequately restored.

Table 3. The templates using in the proposed policy. (A dot • is present for each attribute that is set.)

name	wrap	unwrap	enc	dec	sens	extr	trus	wwt	wt	ut
trusted	•	•			•	•	•	•	usage	usage
usage			•	•	•	•		•	-	-
untrusted			•	•		•			-	-

These policies are incomparable; while the second policy supports lossless key-backup, the policy by Bortolozzo et al. allows for transferring wrapping keys between two PKCS#11 device A and B . Assume A and B have the same wrapping key set-up at the start (e. g., by the SO). A creates another wrapping key, which is wrapped using the common wrapping key, and unwrapped at B . Now A can use this second-level wrapping key to produce wrappings that B can import.

3 Preliminaries

Terms and equational theories. As usual in symbolic protocol analysis we model messages by abstract terms. Therefore we define an order-sorted term algebra with the sort msg and two incomparable subsorts pub and $fresh$. For each of these subsorts we assume a countably infinite set of names, FN for fresh names and PN for public names. Fresh names will be used to model cryptographic keys and nonces while public names model publicly known values. We furthermore assume a countably infinite set of variables for each sort s , \mathcal{V}_s and let \mathcal{V} be the union of the set of variables for all sorts. We write $u : s$ when the name or variable u is of sort s . Let Σ be a signature, i.e., a set of

function symbols, each with an arity. We write f/n when function symbol f is of arity n . We denote by \mathcal{T}_Σ the set of well-sorted terms built over Σ , PN , FN and \mathcal{V} . For a term t we denote by $names(t)$, respectively $vars(t)$ the set of names, respectively variables, appearing in t . The set of ground terms, i.e., terms without variables, is denoted by \mathcal{M}_Σ . When Σ is fixed or clear from the context we often omit it and simply write \mathcal{T} for \mathcal{T}_Σ and \mathcal{M} for \mathcal{M}_Σ .

We equip the term algebra with an equational theory E , that is, a finite set of equations of the form $M = N$ where $M, N \in \mathcal{T}$. From the equational theory we define the binary relation $=_E$ on terms, which is the smallest equivalence relation containing equations in E that is closed under application of function symbols, bijective renaming of names and substitution of variables by terms of the same sort. Furthermore, we require E to distinguish different fresh names, i.e., $\forall a, b \in FN : a \neq b \Rightarrow a \neq_E b$.

Example 1. Symmetric encryption can be modelled using a signature

$$\Sigma = \{ senc/2, sdec/2, true/0 \}$$

and an equational theory defined by

$$sdec(senc(m, k), k) = m.$$

For the rest of the paper we assume that E refers to some fixed equational theory and that the signature and equational theory always contain symbols and equations for pairing and projection, i.e., $\{\langle \cdot, \cdot \rangle, fst, snd\} \subseteq \Sigma$ and equations $fst(\langle x, y \rangle) = x$ and $snd(\langle x, y \rangle) = y$ are in E . We will sometimes use $\langle x_1, x_2, \dots, x_n \rangle$ as a shortcut for $\langle x_1, \langle x_2, \langle \dots, \langle x_{n-1}, x_n \rangle \dots \rangle \rangle$.

We also use the usual notation of positions for terms. A position p is a sequence of positive integers and $t|_p$ denotes the subterm of t at position p .

Facts. We also assume an unsorted signature Σ_{fact} , disjoint from Σ . The set of *facts* is defined as

$$\mathcal{F} := \{F(t_1, \dots, t_k) \mid t_i \in \mathcal{T}_\Sigma, F \in \Sigma_{fact} \text{ of arity } k\}.$$

Facts will be used both to annotate protocols, by the means of events, and for defining multiset rewrite rules. We partition the signature Σ_{fact} into *linear* and *persistent* fact symbols. We suppose that Σ_{fact} always contains a unary, persistent symbol !K and a linear, unary symbol Fr. Given a sequence or set of facts S we denote by $lfacts(S)$ the multiset of all linear facts in S and $pfacts(S)$ the set of all persistent facts in S . By notational convention facts whose identifier starts with ‘!’ will be persistent. \mathcal{G} denotes the set of ground facts, i.e., the set of facts that do not contain variables. For a fact f we denote by $ginsts(f)$ the set of ground instances of f . This notation is also lifted to sequences and sets of facts as expected.

Predicates. We assume an unsorted signature Σ_{pred} disjoint from Σ and Σ_{fact} . The set of *predicates* is defined as

$$\mathcal{P} := \{pr(t_1, \dots, t_k) \mid t_i \in \mathcal{T}_\Sigma, pr \in \Sigma_{pred} \text{ of arity } k\}.$$

Predicates will be used to describe branching conditions in protocols. Each predicate is defined via a first-order formula over ground atoms of the form $t_1 \approx t_2$, i.e., the grammar for such formulae is

$$\langle \phi \rangle ::= t_1 \approx t_2 \mid \neg \phi \mid \phi_1 \wedge \phi_2 \mid \exists x. \phi$$

where t_1, t_2 are terms and $x \in \mathcal{V}$. For an n -ary predicate pr , $pr(x_1, \dots, x_n)$ is defined by a formula ϕ_{pr} such that $fv(\phi_{pr}) \subseteq x_1, \dots, x_n$, where fv denotes the free variables in a formula, i.e., variables $v \in \mathcal{V}$ not bound by $\exists v$. The semantics of the first-order formulae is as usual where we interpret \approx as $=_E$. We use $\sigma_1 \vee \sigma_2$ as short-hand for $\neg(\neg\sigma_1 \wedge \neg\sigma_2)$.

Example 2. A predicate $pr = can_wrap$ is used to check whether the attributes associated to two handles (10 attributes each) allow for wrapping. For readability, we rename x_1 to $wrap_1$, x_{15} to $extr_2$, x_7 to $trus_1$ and x_{18} to wwt_2 :

$$\begin{aligned} \sigma_{can_wrap} := & (wrap_1 \approx 'on' \wedge extr_2 \approx 'on') \wedge \\ & ((wwt_2 \approx 'off') \vee (wwt_2 \approx 'on' \wedge trus_1 \approx 'on')) \end{aligned}$$

Substitutions. A substitution σ is a partial function from variables to terms. We suppose that substitutions are well-typed, i.e., they only map variables of sort s to terms of sort s , or of a subsort of s . We denote by $\sigma = \{t_1/x_1, \dots, t_n/x_n\}$ the substitution whose domain is $\mathbf{D}(\sigma) = \{x_1, \dots, x_n\}$ and which maps x_i to t_i . As usual we homomorphically extend σ to apply to terms and facts and use a postfix notation to denote its application, e.g., we write $t\sigma$ for the application of σ to the term t . A substitution σ is grounding for a term t if $t\sigma$ is ground. Given function g we let $g(x) = \perp$ when $x \notin \mathbf{D}(g)$. When $g(x) = \perp$ we say that g is undefined for x . We define the function $f := g[a \mapsto b]$ with $\mathbf{D}(f) = \mathbf{D}(g) \cup \{a\}$ as $f(a) := b$ and $f(x) := g(x)$ for $x \neq a$.

Sets, sequences and multisets. We write \mathbb{N}_n for the set $\{1, \dots, n\}$. Given a set S we denote by S^* the set of finite sequences of elements from S and by $S^\#$ the set of finite multisets of elements from S . We use the superscript $\#$ to annotate usual multiset operations, e.g., $S_1 \cup^\# S_2$ denotes the multiset union of multisets S_1, S_2 . Given a multiset S we denote by $set(S)$ the set of elements in S . The sequence consisting of elements e_1, \dots, e_n will be denoted by $[e_1, \dots, e_n]$ and the empty sequence is denoted

$$\begin{array}{ll} \langle P, Q \rangle ::= 0 & \langle P, Q \rangle ::= (continued) \\ \mid P \mid Q & \mid \text{insert } M, N; P \\ \mid !P & \mid \text{delete } M; P \\ \mid vn: fresh; P & \mid \text{lookup } M \text{ as } x \text{ in } P \text{ [else } Q] \\ \mid \text{out}(M, N); P & \mid \text{lock } M; P \\ \mid \text{in}(M, N); P & \mid \text{unlock } M; P \\ \mid \text{if } Pred \text{ then } P \text{ [else } Q] \quad Pred \in \mathcal{P} & \mid [L] \neg[A] \rightarrow [R]; P \quad (L, R, A \in \mathcal{F}^*) \\ \mid \text{event } F; P \quad (F \in \mathcal{F}) & \end{array}$$

Fig. 1. Syntax, where $M, N \in \mathcal{T}$

by $[\]$. Given a sequence S , we denote by $idx(S)$ the set of positions in S , i.e., \mathbb{N}_n when S has n elements, and for $i \in idx(S)$ S_i denotes the i th element of the sequence. Set membership modulo E is denoted by \in_E and defined as $e \in_E S$ if $\exists e' \in S. e' =_E e$. \subset_E and $=_E$ are defined for sets in a similar way. Application of substitutions is lifted to sets, sequences and multisets as expected. By abuse of notation we sometimes interpret sequences as sets or multisets; the applied operators should make the implicit cast clear.

3.1 A Cryptographic pi Calculus with Explicit State

Syntax and Informal Semantics. The Saptic calculus is a variant of the applied pi calculus [1]. In addition to the usual operators for concurrency, replication, communication and name creation, it offers several constructs for reading and updating an explicit global state. The grammar for processes is described in Fig. 1.

0 denotes the terminal process. $P \mid Q$ is the parallel execution of processes P and Q and $!P$ the replication of P , allowing an unbounded number of sessions in protocol executions. The construct $\nu n; P$ binds the name $n \in FN$ in P and models the generation of a fresh, random value. Processes $out(M, N); P$ and $in(M, N); P$ represent the output, respectively input, of message N on channel M . Readers familiar with the applied pi calculus [1] may note that we opted for the possibility of pattern matching in the input construct, rather than merely binding the input to a variable x . The process if $Pred$ then P else Q will execute P or Q , depending on whether $Pred$ holds. For example, if $Pred = equal(M, N)$, and $\phi_{equal} = x_1 \approx x_2$, then if $equal(M, N)$ then P else Q will execute P if $M =_E N$ and Q otherwise. (In the following, we will use $M = N$ as short-hand for $equal(M, N)$.) The event construct is merely used for annotating processes and will be useful for stating security properties. For readability we sometimes omit to write else Q when Q is 0 , as well as trailing 0 processes.

The remaining constructs are used for manipulating state and are new compared to the applied pi calculus. We offer two different mechanisms for state. The first construct is *functional* and allows to associate a value to a key. The construct $insert\ M, N$ binds the value N to a key M . Successive inserts allow to change this binding. The delete M operation simply “undefines” the mapping for the key M . The lookup M as x in P else Q allows to retrieve the value associated to M , binding it to the variable x in P . If the mapping is undefined for M the process behaves as Q . The lock and unlock constructs allow to gain exclusive access to a resource M . This is essential for writing protocols where parallel processes may read and update a common memory. We additionally offer another kind of global state in form of a multiset of ground facts, as opposed to the previously introduced functional store. The purpose of this construct is to provide access to the underlying notion of state in tamarin, but we stress that it is distinct from the previously introduced functional state, and its use is only advised to expert users. It is not used in the present modelling.

The bound names of a process are those that are bound by νn . We suppose that all names of sort *fresh* appearing in the process are under the scope of such a binder. Free names must be of sort *pub*. A variable x can be bound in three ways: (i) by the construct $lookup\ M\ as\ x$, or (ii) $x \in vars(N)$ in the construct $in(M, N)$ and x is not under the scope of a previous binder, (iii) $x \in vars(L)$ in the construct $[L] \text{ } \dashv\!\!\dashv\!\!\! \rightarrow [R]$ and x is not under the scope of a previous binder. While the construct $lookup\ M\ as\ x$ always

acts as a binder, the input and $[L] \text{---} [A] \text{---} [R]$ constructs do not rebind an already bound variable but perform pattern matching.

A process is ground if it does not contain any free variables. We denote by $P\sigma$ the application of the homomorphic extension of the substitution σ to P . As usual we suppose that the substitution only applies to free variables. We sometimes interpret the syntax tree of a process as a term and write $P|_p$ to refer to the subprocess of P at position p (where $|$, if and lookup are interpreted as binary symbols, all other constructs as unary).

Semantics

Frames and deduction. Before giving the formal semantics of our calculus we introduce the notions of frame and deduction. A *frame* consists of a set of fresh names \tilde{n} and a substitution σ and is written $\nu\tilde{n}.\sigma$. Intuitively a frame represents the sequence of messages that have been observed by an adversary during a protocol execution and secrets \tilde{n} generated by the protocol, a priori unknown to the adversary. Deduction models the capacity of the adversary to compute new messages from the observed ones.

Definition 1 (Deduction). *We define the deduction relation $\nu\tilde{n}.\sigma \vdash t$ as the smallest relation between frames and terms defined by the deduction rules in Figure 2.*

$$\begin{array}{c}
 \frac{a \in FN \cup PN \quad a \notin \tilde{n}}{\nu\tilde{n}.\sigma \vdash a} \text{ DNAME} \qquad \frac{\nu\tilde{n}.\sigma \vdash t \quad t =_E t'}{\nu\tilde{n}.\sigma \vdash t'} \text{ DEQ} \\
 \\
 \frac{x \in \mathbf{D}(\sigma)}{\nu\tilde{n}.\sigma \vdash x\sigma} \text{ DFRAME} \qquad \frac{\nu\tilde{n}.\sigma \vdash t_1 \cdots \nu\tilde{n}.\sigma \vdash t_n \quad f \in \Sigma^k}{\nu\tilde{n}.\sigma \vdash f(t_1, \dots, t_n)} \text{ DAPPL}
 \end{array}$$

Fig. 2. Deduction rules

Operational semantics. We can now define the operational semantics of our calculus. The semantics is defined by a labelled transition relation between process configurations. A *process configuration* is a 6-tuple $(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L})$ where

- $\mathcal{E} \subseteq FN$ is the set of fresh names generated by the processes;
- $\mathcal{S} : \mathcal{M}_\Sigma \rightarrow \mathcal{M}_\Sigma$ is a partial function modeling the functional store;
- $\mathcal{S}^{\text{MS}} \subseteq \mathcal{G}^\#$ is a multiset of ground facts and models the multiset of stored facts;
- \mathcal{P} is a multiset of ground processes representing the processes executed in parallel;
- σ is a ground substitution modeling the messages output to the environment;
- $\mathcal{L} \subseteq \mathcal{M}_\Sigma$ is the set of currently acquired locks.

The transition relation is defined by the rules described in Fig. 3. Transitions are labelled by sets of ground facts. For readability we omit empty sets and brackets around singletons, i.e., we write \rightarrow for $\xrightarrow{\emptyset}$ and \xrightarrow{f} for $\xrightarrow{\{f\}}$. We write \rightarrow^* for the reflexive, transitive closure of \rightarrow (the transitions that are labelled by the empty sets) and write \xRightarrow{f} for $\rightarrow^* \xrightarrow{f} \rightarrow^*$. We can now define the set of traces, i.e., possible executions, that a process admits.

Definition 2 (Traces of P). Given a ground process P we define the set of traces of P as

$$\text{traces}^{pi}(P) = \left\{ [F_1, \dots, F_n] \mid (\emptyset, \emptyset, \emptyset, \{P\}, \emptyset, \emptyset) \xrightarrow{F_1} (\mathcal{E}_1, \mathcal{S}_1, \mathcal{S}_1^{\text{MS}}, \mathcal{P}_1, \sigma_1, \mathcal{L}_1) \xrightarrow{F_2} \dots \xrightarrow{F_n} (\mathcal{E}_n, \mathcal{S}_n, \mathcal{S}_n^{\text{MS}}, \mathcal{P}_n, \sigma_n, \mathcal{L}_n) \right\}$$

3.2 Security Properties

The formalism used for defining security properties in the Saptic tool, which is used to define key secrecy as well as helping lemmas was introduced with the tamarin tool [23]. Security properties are described in an expressive two-sorted first-order logic. The sort $temp$ is used for time points, \mathcal{V}_{temp} are temporal variables.

Definition 3 (Trace formulas). A trace atom is either false \perp , a term equality $t_1 \approx t_2$, a timepoint ordering $i \leq j$, a timepoint equality $i \doteq j$, or an action $F@i$ for a fact $F \in \mathcal{F}$ and a timepoint i . A trace formula is a first-order formula over trace atoms.

As we will see in our case studies this logic is expressive enough to analyze a variety of security properties, including complex injective correspondence properties.

To define the semantics, let each sort s have a domain $\mathbf{D}(s)$. $\mathbf{D}(temp) = \mathcal{Q}$, $\mathbf{D}(msg) = \mathcal{M}$, $\mathbf{D}(fresh) = FN$, and $\mathbf{D}(pub) = PN$. A function $\theta : \mathcal{V} \rightarrow \mathcal{M} \cup \mathcal{Q}$ is a valuation if it respects sorts, that is, $\theta(\mathcal{V}_s) \subset \mathbf{D}(s)$ for all sorts s . If t is a term, $t\theta$ is the application of the homomorphic extension of θ to t .

Definition 4 (Satisfaction relation). The satisfaction relation $(tr, \theta) \models \varphi$ between trace tr , valuation θ and trace formula φ is defined as follows:

$$\begin{aligned} (tr, \theta) \models \perp & \quad \text{never} \\ (tr, \theta) \models F@i & \quad \text{iff } \theta(i) \in \text{idx}(tr) \text{ and } F\theta \in_E tr_{\theta(i)} \\ (tr, \theta) \models i \leq j & \quad \text{iff } \theta(i) < \theta(j) \\ (tr, \theta) \models i \doteq j & \quad \text{iff } \theta(i) = \theta(j) \\ (tr, \theta) \models t_1 \approx t_2 & \quad \text{iff } t_1\theta =_E t_2\theta \\ (tr, \theta) \models \neg\varphi & \quad \text{iff not } (tr, \theta) \models \varphi \\ (tr, \theta) \models \varphi_1 \wedge \varphi_2 & \quad \text{iff } (tr, \theta) \models \varphi_1 \text{ and } (tr, \theta) \models \varphi_2 \\ (tr, \theta) \models \exists x : s.\varphi & \quad \text{iff there is } u \in \mathbf{D}(s) \text{ such that} \\ & \quad (tr, \theta[x \mapsto u]) \models \varphi \end{aligned}$$

When φ is a ground formula we sometimes simply write $tr \models \varphi$ as the satisfaction of φ is independent of the valuation.

Definition 5 (Validity, satisfiability). Let $Tr \subseteq (\mathcal{P}(\mathcal{G}))^*$ be a set of traces. A trace formula φ is said to be valid for Tr , written $Tr \models^\forall \varphi$, if for any trace $tr \in Tr$ and any valuation θ we have that $(tr, \theta) \models \varphi$.

A trace formula φ is said to be satisfiable for Tr , written $Tr \models^\exists \varphi$, if there exist a trace $tr \in Tr$ and a valuation θ such that $(tr, \theta) \models \varphi$.

Note that $Tr \models^\forall \varphi$ iff $Tr \not\models^\exists \neg\varphi$. Given a ground process P we say that φ is valid, written $P \models^\forall \varphi$, if $\text{traces}^{pi}(P) \models^\forall \varphi$, and that φ is satisfied in P , written $P \models^\exists \varphi$, if $\text{traces}^{pi}(P) \models^\exists \varphi$.

Standard operations:

$$\begin{aligned}
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{0\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P|Q\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P, Q\}, \sigma, \mathcal{L}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{!P\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{!P, P\}, \sigma, \mathcal{L}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{\nu a; P\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E} \cup \{a'\}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\{a'/a\}\}, \sigma, \mathcal{L}) \\
&\quad \text{if } a' \text{ is fresh} \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L}) &\xrightarrow{K(M)} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L}) \quad \text{if } \nu \mathcal{E}. \sigma \vdash M \\
(\dots, \mathcal{P} \cup^{\#} \{\text{out}(M, N); P\}, \sigma, \mathcal{L}) &\xrightarrow{K(M)} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\}, \sigma \cup \{^N/x\}, \mathcal{L}) \\
&\quad \text{if } x \text{ is fresh and } \nu \mathcal{E}. \sigma \vdash M \\
(\dots, \mathcal{P} \cup^{\#} \{\text{in}(M, N); P\}, \sigma, \mathcal{L}) &\xrightarrow{K((M, N\tau))} (\dots, \mathcal{P} \cup^{\#} \{P\tau\}, \sigma, \mathcal{L}) \\
&\quad \text{if } \exists \tau. \tau \text{ is grounding for } N, \nu \mathcal{E}. \sigma \vdash M, \nu \mathcal{E}. \sigma \vdash N\tau \\
(\dots, \mathcal{P} \cup^{\#} \{\text{out}(M, N); P, \\
&\quad \text{in}(M', N'); Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P, Q\tau\}, \sigma, \mathcal{L}) \\
&\quad \text{if } M =_E M' \text{ and } \exists \tau. N =_E N'\tau \text{ and } \tau \text{ grounding for } N' \\
(\dots, \mathcal{P} \cup^{\#} \{\text{if } pr(M_1, \dots, M_n) \\
&\quad \text{then } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\}, \sigma, \mathcal{L}) \\
&\quad \text{if } \phi_{pr}\{^{M_1}/x_1, \dots, ^{M_n}/x_n\} \text{ is satisfied} \\
(\dots, \mathcal{P} \cup^{\#} \{\text{if } pr(M_1, \dots, M_n) \\
&\quad \text{then } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{Q\}, \sigma, \mathcal{L}) \\
&\quad \text{if } \phi_{pr}\{^{M_1}/x_1, \dots, ^{M_n}/x_n\} \text{ is not satisfied} \\
(\dots, \mathcal{P} \cup^{\#} \{\text{event}(F); P\}, \sigma, \mathcal{L}) &\xrightarrow{F} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup \{P\}, \sigma, \mathcal{L})
\end{aligned}$$

Operations on global state:

$$\begin{aligned}
(\dots, \mathcal{P} \cup^{\#} \{\text{insert } M, N; P\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E}, \mathcal{S}[M \mapsto N], \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\}, \sigma, \mathcal{L}) \\
(\dots, \mathcal{P} \cup^{\#} \{\text{delete } M; P\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E}, \mathcal{S}[M \mapsto \perp], \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\}, \sigma, \mathcal{L}) \\
(\dots, \mathcal{P} \cup^{\#} \{\text{lookup } M \text{ as } x \\
&\quad \text{in } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\{V/x\}\}, \sigma, \mathcal{L}) \\
&\quad \text{if } \mathcal{S}(N) =_E V \text{ is defined and } N =_E M \\
(\dots, \mathcal{P} \cup^{\#} \{\text{lookup } M \text{ as } x \\
&\quad \text{in } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{Q\}, \sigma, \mathcal{L}) \\
&\quad \text{if } \mathcal{S}(N) \text{ is undefined for all } N =_E M \\
(\dots, \mathcal{P} \cup^{\#} \{\text{lock } M; P\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\}, \sigma, \mathcal{L} \cup \{M\}) \\
&\quad \text{if } M \notin_E \mathcal{L} \\
(\dots, \mathcal{P} \cup^{\#} \{\text{unlock } M; P\}, \sigma, \mathcal{L}) &\longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup^{\#} \{P\}, \sigma, \mathcal{L} \setminus \{M' : M' =_E M\}) \\
(\dots, \mathcal{P} \cup^{\#} \{[l \text{ } \neg a] \rightarrow r\}; P\}, \sigma, \mathcal{L}) &\xrightarrow{a'} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}} \setminus \text{lfacts}(l') \cup^{\#} r', \mathcal{P} \cup^{\#} \{P\tau\}, \sigma, \mathcal{L}) \\
&\quad \text{if } \exists \tau, l', a', r'. \tau \text{ grounding for } l \text{ } \neg a \rightarrow r, l' \text{ } \neg a' \rightarrow r' =_E (l \text{ } \neg a] \rightarrow r)\tau, \\
&\quad \text{lfacts}(l') \subseteq^{\#} \mathcal{S}^{\text{MS}}, \text{pfacts}(l') \subset \mathcal{S}^{\text{MS}}
\end{aligned}$$

Fig. 3. Operational semantics

4 Model

In this section, we will introduce our model of a PKCS#11 token. The complete code is available at <http://sapic.gforge.inria.fr/pkcs11templates.zip>. We consider a security device that allows the creation of keys in its secure memory. The user can access the device via an API. If he creates a key, he obtains a handle, which he can use to let the device perform operations on his behalf. For each handle the device also stores a list of attributes, which define what operations are permitted for this handle. The goal is that the user can never gain knowledge of the key, as the user's machine might be compromised. We model the device by the following process (we use $\text{out}(m)$ as a shortcut for $\text{out}(c, m)$ for a public channel c):

$P_{init}; !(P_{create} \mid P_{dec} \mid P_{enc} \mid P_{wrap} \mid P_{unwrap} \mid P_{get_keyval})$, where

```

Pinit :=
/* wrap, unwrap, enc, dec, sens, extr, trus, wwt, wt, ut */
insert ⟨'template', 'trusted'⟩,
  ⟨'on', 'on', 'off', 'off', 'on', 'on', 'on', 'on', 'usage', 'usage'⟩;
insert ⟨'template', 'usage'⟩,
  ⟨'off', 'off', 'on', 'on', 'on', 'on', 'off', 'on', 'undef', 'undef'⟩;
insert ⟨'template', 'untrusted'⟩,
  ⟨'off', 'off', 'on', 'on', 'off', 'on', 'off', 'off', 'undef', 'undef'⟩;

```

This sets up the templates before starting operation in the replicated process. Key-generation is handled by P_{create} :

```

Pcreate := in(⟨'create', atts, ptr⟩);
lock 'device';
ν h; ν k;
lookup ⟨'template', ptr⟩ as templ in
if permits( attwrap(templ), [...], attut(templ),
           attwrap(atts), [...], attut(atts)) then
  event NewKey(h, k, attsens(atts));
  insert ⟨'obj', h⟩, ⟨k, atts⟩;
  event WrapKey(h, k, attwrap(atts)); event DecKey(h, k, attdec(atts));
  event EncKey(h, k, attenc(atts)); event UnwrapKey(h, k, attunwrap(atts));
  out(h) ;
unlock 'device'

```

Upon reception of a key generation request with a list of attributes and a pointer to a template, the device is locked. Then, the device creates a new handle h and a key k . The pointer is retrieved from the database P_{init} has written to. The functions $attwrap$, to $attut$ are simple deconstructors, $attwrap$, for example, extracts the first element from a list of 10 attributes (see Tab. 1). The predicate $permits$ compares the attributes given by the adversary with the attributes stored with the template. In subsection 2.2, we have argued that templates should determine the key's attributes, hence $permits$ is true if and only if the attributes including the pointers to the templates match exactly. To obtain

a more permissive modelling, this predicate can be altered, e. g., to allow for certain attributes to be changed.

The creation of keys is logged as an event $\text{NewKey}(h, k, \text{attsens}(atts))$. If the third argument is ‘on’, this key is sensitive, i. e., secrecy needs to be preserved. Events are used to state security properties and helping lemmas. Next, the device stores the key that belongs to the handle by associating the pair $\langle \text{‘key’}, h \rangle$ to the value of the key k and the attributes. The events WrapKey to UnwrapKey are used to refer to the attributes of keys in helping lemmas and otherwise irrelevant. Finally, the handle is output and the device unlocked.

Remark 1. The predicate *permits* compares the attributes, including the pointers wt and ut , literally. Since our policy only accepts pointers to templates created by P_{mit} , and since those are distinct, this is without loss of generality. Furthermore, the adversary has to provide the pointer to the template, which is without loss of generality, too, since the pointers are of sort *pub*.

If a handle has the ‘dec’ attribute set, it can be used for decryption:

```

 $P_{\text{dec}} := \text{in}(\langle h, \text{senc}(m, k) \rangle);$ 
lock ‘device’;
lookup  $\langle \text{‘obj’}, h \rangle$  as  $v$  in
  if  $\text{can\_decrypt}(\text{attwrap}(\text{tem}(v)), [\dots], \text{attut}(\text{tem}(v)))$  then
    if  $\text{key}(v) = k$  then
      event  $\text{DecUsing}(k, m); \text{out}(m); \text{unlock ‘device’}$ 

```

The lookup stores the value associated to $\langle \text{‘obj’}, h \rangle$ in v . The predicate *can_decrypt* is satisfied, iff. the fourth argument, $\text{attdec}(\text{tem}(v))$ equals ‘on’. The function symbol *tem* extracts the second element of a pair, it is defined by the equation $\text{tem}(\langle k, t \rangle) = t$. Similarly, $\text{key}(\langle k, t \rangle) = k$. If the key stored with this handle matches the key used to generate the encryption, the plain-text is output and the device unlocked.

If a key has the ‘wrap’ attribute set, it can be used to encrypt the value of a second key:

```

 $P_{\text{wrap}} := \text{in}(\langle h1, h2 \rangle);$ 
lock ‘device’;
lookup  $\langle \text{‘obj’}, h1 \rangle$  as  $v1$  in
  lookup  $\langle \text{‘obj’}, h2 \rangle$  as  $v2$  in
    if  $\text{can\_wrap}(\text{attwrap}(\text{tem}(v1)), [\dots], \text{attut}(\text{tem}(v1)),$ 
       $\text{attwrap}(\text{tem}(v2)), [\dots], \text{attut}(\text{tem}(v2)))$  then
      lookup  $\langle \text{‘template’}, \text{attwt}(\text{tem}(v1)) \rangle$  as  $wt$  in
        if  $\text{permits}(\text{attwrap}(wt), [\dots], \text{attut}(wt),$ 
           $\text{attwrap}(\text{tem}(v2)), [\dots], \text{attut}(\text{tem}(v2)))$  then
          event  $\text{Wrap}(\text{key}(v1), \text{key}(v2));$ 
          out( $\text{senc}(\text{key}(v2), \text{key}(v1))$ );
          unlock ‘device’

```

Given the two handles h_1 and h_2 , the corresponding keys and attributes are retrieved. The predicate *can_wrap* is defined in Example 2 on page 228. If the template referred

to by the entry ‘wt’ in the first key’s attribute list permits wrapping the second key, i. e., it is equal to the second handle’s attributes, then the encryption is output. The complete model including P_{enc} , P_{unwrap} and P_{get_keyval} can be found at <http://sapic.gforge.inria.fr/pkcs11templates.zip>.

Limitations. While being more detailed than previous works in terms of the attributes and commands supported, our model does have the following limitations: *Integrity of wrappings.* In our model, a wrapping can only be imported if it is an encryption of some message with a key that is on the device. In reality, this integrity property cannot be given. There are techniques that allow to account for malleability of cypher-texts in the symbolic model [3]. The current draft for version 2.40 indicates that authenticated encryption might be part of the future version of this standard [21]. *Multiple tokens.* We currently model a single token. By encapsulating the current process P in $\nu device: pub.P$ and prepending $device$ to database keys, this situation could be modelled. *Copying keys, asymmetric keys, key derivation.* $C_CopyObject$, support for asymmetric keys and key-derivation are not modelled for various reasons explained in Section 2.1. While we conjecture copying objects could be enabled, and asymmetric keys as well as key-derivation keys be allowed as ‘usage’ keys, this is missing in the current model (as opposed to other related work [11, 7]).

5 Security Results for the Proposed Policy

The main security result is the secrecy of keys generated on the device that have been marked ‘sensitive’ upon creation, in our case, keys created with the templates ‘trusted’ or ‘usage’. This is expressed by the following trace formula:

$$\neg(\exists h, k: msg, i, j: temp. \text{NewKey}(h, k, 'on')@i \wedge K(k)@j)$$

The action NewKey refers to the event in the key-generation process P_{create} , which we introduced in Section 4. If the third argument is ‘on’, k has been created ‘sensitive’. To derive the result, we have defined 9 helping lemmas, four of which are rather trivial, but help speeding up the proof. The first, dec_limits , establishes typing invariants, most importantly, it states that decrypt is not useful to the adversary: Any message obtained by decryption was either known to the adversary in advance, or a key that was created ‘sensitive’ or imported was leaked, or there was some key that had the attributes ‘wrap’ and ‘dec’ set at different points in time. The following four lemmas state that given the templates it is, e. g., not possible to create a key with ‘trusted’ as wrapping template. The lemma bad_keys states that a key that was created by unwrapping must earlier have been created on the device, unless something bad happened, i. e., either a sensitive key leaked, or a key had wrap and dec, or unwrap and enc set before. The latter two conflicts are known to cause attacks [8]. The following lemma $\text{no_key_is_wrap_and_dec}$ says that the first conflict can only occur if the second occurred before, or a sensitive key was leaked. Subsequently, it is shown the second conflict cannot occur, unless a key was leaked. The lemma $\text{cannot_obtain_key_ind}$ is an inductive version of the security property cannot_obtain_key .

Table 4. Evaluation of our proof method with and without the model-specific heuristics, broken down into lemmas. Interaction is measured in terms of mouse clicks. The user choses the next proof goal out of a list in a web interface. The value was determined by counting occurrences of the keyword ‘solve’ in the file. The proof size is the number of case distinctions.

name	interaction		proof size	
	(w/o heur.)	(w/ heur.)	(w/o heur.)	(w/ heur.)
dec_limits	11	0	3394	2235
trusted_as_ut_impossible	0	0	4	4
untrusted_as_ut_impossible	0	0	4	4
untrusted_as_wt_impossible	0	0	4	4
trusted_as_wt_impossible	0	0	4	4
bad_keys	0	0	2988	683
no_key_is_wrap_and_dec[..]	15	0	1177	2396
no_key_is_enc_and_unwrap	29	0	2669	352
cannot_obtain_key_ind	6	0	7306	14598
cannot_obtain_key	0	0	0	0

User intervention is necessary to find the proof when the tamarin prover is used with heuristics adapted to general Sapic output. They differ from the standard ‘smart’ heuristic only in that the actions corresponding to unlock operations and premises corresponding to the previous state in the execution are prioritized. Using a model-specific heuristic, it is possible to find the proof automatically. This heuristic prioritizes the resolution of insert operations with ‘template’ in the first argument (thereby moving the case distinction about which templates is used upwards in the proof tree) and deprioritizes the adversary’s deduction of handles (as they are public anyway). With this heuristic, the complete proof is found within half an hour on a computation server with 24 Intel Xeon 2.67GHz cores and 50GB RAM. Experiments on desktop machines are planned. See Table 4 for more details. The manual part of the proof, transcripts of the complete proof and the tools used are available at <http://sapic.gforge.inria.fr/pkcs11templates.zip>.

6 Conclusion and Future Work

We have investigated a new method of verifying key secrecy in PKCS#11 configurations following the protocol verification approach. It requires manual effort in defining helping lemmas, but overcomes the need for a model-specific approximation techniques and models PKCS#11 in a precise and intuitive manner. In particular, our model supports features that have been added in version 2.20 of the standard. The upcoming [21] version 2.40 is available for public review and will support encryption with authenticated data, a mechanism that has long been requested and has the potential to remove the restriction to three-level policies outlined in Section 2.2 and allow for new policies. Future work on version 2.40 will benefit from the fact that our verification technique applies without model-specific approximation. The second contribution is a secure configuration of PKCS#11 that permits wrapping and unwrapping for backup purposes.

Besides adapting our model to the case of multiple tokens in the network, we plan to increase the degree of automation by refining the model-specific heuristic for the tamarin prover's constraint solving algorithm. The current heuristic is rather simple, but may be generalized. We stress that completeness and soundness of the tool chain are independent of the heuristic used, suggesting a tool chain without approximation but adapted heuristics as an alternative to ad-hoc approximations that does not necessarily provide decidability, but is more flexible with regard to extensions and less reliant on hand-written proofs.

Acknowledgements. This work was supported by the German Federal Ministry of Education and Research (BMBF) within EC SPRIDE.

References

- [1] Abadi, M., Fournet, C.: Mobile Values, New Names, and Secure Communication. In: POPL 2001. ACM Press (2001)
- [2] Adão, P., Focardi, R., Luccio, F.L.: Type-Based Analysis of Generic Key Management APIs. In: CSF, pp. 97–111. IEEE (2013)
- [3] Ahmed, N., Jensen, C.D., Zenner, E.: Towards Symbolic Encryption Schemes. In: Foresti, S., Yung, M., Martinelli, F. (eds.) ESORICS 2012. LNCS, vol. 7459, pp. 557–572. Springer, Heidelberg (2012)
- [4] Bardou, R., Focardi, R., Kawamoto, Y., Simionato, L., Steel, G., Tsay, J.-K.: Efficient Padding Oracle Attacks on Cryptographic Hardware. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 608–625. Springer, Heidelberg (2012)
- [5] Bond, M., Anderson, R.: API level attacks on embedded systems. *IEEE Computer Magazine* 34(10) (2001)
- [6] Bortolozzo, M., et al.: Attacking and Fixing PKCS#11 Security Tokens. In: CCS 2010. ACM Press (2010)
- [7] Centenaro, M., Focardi, R., Luccio, F.L.: Type-based analysis of key management in PKCS#11 cryptographic devices. *Journal of Computer Security* 21(6) (2013)
- [8] Clulow, J.: On the security of PKCS #11. In: Walter, C.D., Koç, Ç.K., Paar, C. (eds.) CHES 2003. LNCS, vol. 2779, pp. 411–425. Springer, Heidelberg (2003)
- [9] Cortier, V., Keighren, G., Steel, G.: Automatic Analysis of the Security of XOR-Based Key Management Schemes. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 538–552. Springer, Heidelberg (2007)
- [10] Cortier, V., Steel, G., Wiedling, C.: Revoke and let live: a secure key revocation API for cryptographic devices. In: CCS 2012. ACM (2012)
- [11] Delaune, S., Kremer, S., Steel, G.: Formal Analysis of PKCS#11 and Proprietary Extensions. *Journal of Computer Security* 18(6) (2010)
- [12] Durgin, N., et al.: Undecidability of Bounded Security Protocols. In: Workshop on Formal Methods and Security Protocols. IEEE (1999)
- [13] Fröschle, S., Sommer, N.: Concepts and Proofs for Configuring PKCS#11. In: Barthe, G., Datta, A., Etalle, S. (eds.) FAST 2011. LNCS, vol. 7140, pp. 131–147. Springer, Heidelberg (2012)
- [14] Fröschle, S., Steel, G.: Analysing PKCS#11 key management APIs with unbounded fresh data. In: Degano, P., Viganò, L. (eds.) ARSPA-WITS 2009. LNCS, vol. 5511, pp. 92–106. Springer, Heidelberg (2009)

- [15] Fröschle, S., Sommer, N.: Reasoning with past to prove PKCS#11 keys secure. In: Degano, P., Etalle, S., Guttman, J. (eds.) FAST 2010. LNCS, vol. 6561, pp. 96–110. Springer, Heidelberg (2011)
- [16] Fröschle, S.B., Sommer, N.: When is a PKCS#11 configuration secure? Tech. rep. Reports of SFB/TR 14 AVACS 82, SFB/TR 14 AVACS (2011), <https://vhome.offis.de/sibyllef/cryptokireport.pdf>
- [17] Kremer, S., Künnemann, R.: Automated analysis of security protocols with global state. In: Security and Privacy. IEEE Computer Society (2014)
- [18] Kremer, S., Künnemann, R., Steel, G.: Universally Composable Key-Management. In: Crampton, J., Jajodia, S., Mayes, K. (eds.) ESORICS 2013. LNCS, vol. 8134, pp. 327–344. Springer, Heidelberg (2013)
- [19] Kremer, S., Steel, G., Warinschi, B.: Security for Key Management Interfaces. In: CSF 2011, pp. 66–82. IEEE Computer Society (2011)
- [20] Longley, D., Rigby, S.: An Automatic Search for Security Flaws in Key Management Schemes. *Computers and Security* 11(1) (March 1992)
- [21] PKCS #11 Cryptographic Token Interface Base Specification Version 2.40, Committee Specification 01. OASIS Open (September 2014), <http://docs.oasis-open.org/pkcs11/pkcs11-base/v2.40/cs01/pkcs11-base-v2.40-cs01.html>
- [22] PKCS #11: Cryptographic Token Interface Standard. RSA Security Inc. v2.20 (June 2004)
- [23] Schmidt, B., et al.: Automated Analysis of Diffie-Hellman Protocols and Advanced Security Properties. In: CSF 2012. IEEE (2012)