

# How to Disembed a Program? (Extended Abstract\*)

Benoît Chevallier-Mames<sup>1</sup>, David Naccache<sup>1</sup>, Pascal Paillier<sup>1</sup>, and  
David Pointcheval<sup>2</sup>

<sup>1</sup> Gemplus/Applied Research and Security Center  
{benoit.chevallier-mames,david.naccache,pascal.paillier}@gemplus.com  
<sup>2</sup> Ecole Normale Supérieure/CNRS  
david.pointcheval@ens.fr

**Abstract.** This paper presents the theoretical blueprint of a new secure token called the *Externalized Microprocessor* ( $X\mu P$ ). Unlike a smart-card, the  $X\mu P$  contains no ROM at all.

While exporting all the device's executable code to potentially untrustworthy terminals poses formidable security problems, the advantages of ROM-less secure tokens are numerous: chip masking time disappears, bug patching becomes a mere *terminal update* and hence does not imply any roll-out of cards in the field. Most importantly, code size ceases to be a limiting factor. This is particularly significant given the steady increase in on-board software complexity.

After describing the machine's instruction-set we introduce a public-key oriented architecture design which relies on a new RSA screening scheme and features a relatively low communication overhead. We propose two protocols that execute and dynamically authenticate arbitrary programs, provide a strong security model for these protocols and prove their security under appropriate complexity assumptions.

**Keywords:** Embedded cryptography, RSA screening schemes, ROM-less smart cards, program authentication, compilation theory, provable security, mobile code.

## 1 Introduction

The idea of inserting a chip into a plastic card is as old as public-key cryptography. The first patents are now 25 years old but mass applications emerged only a decade ago because of limitations in the storage and processing capacities of circuit technology. More recently new silicon geometries and cryptographic processing refinements led the industry to new generations of cards and more complex applications such as multi-applicative cards [7].

Over the last decade, there has been an increasing demand for more and more complex smart-cards from national administrations, telephone operators

---

\* The full version of this work can be found at [6].

and banks. Complexity grew to the point where current cards are nothing but miniature computers embarking a linker, a loader, a Java virtual machine, remote method invocation modules, a bytecode verifier, an applet firewall, a garbage collector, cryptographic libraries, a complex protocol stack plus numerous other clumsy OS components.

This paper ambitions to propose a disruptive secure-token model that tames this complexity explosion in a flexible and secure manner. From a theoretical standpoint, we look back to von Neumann's computing model wherein a processing unit operates on volatile and nonvolatile memories, generates random numbers, exchanges data via a communication tape and receives instructions from a program memory. We revisit this model by alleviating the integrity assumption on the executed program, explicitly allowing malevolent and arbitrary modifications of its contents. Assuming a cryptographic key is stored in non-volatile memory, the property we achieve is that no *chosen-program* attack can actually infer information on this key or modify its value: only authentic programs, the ones written by the genuine issuer of the architecture, may do so.

Quite customizable and generic in several ways, our execution protocols are directly applicable to the context of a ROM-less smart card (called the Externalized Microprocessor or  $X\mu P$ ) interacting with a powerful terminal (Externalized Terminal or XT). The  $X\mu P$  executes and dynamically authenticates external programs of *arbitrary size* without intricate code-caching mechanisms. This approach not only simplifies current smart-card-based applications but also presents immense advantages over state-of-the-art technologies on the security marketplace. Notable features of the  $X\mu P$  are further discussed in Section 7 and in the full version of this work [6]. We start by introducing the architecture and programming language of the  $X\mu P$  in the next section. After describing our execution protocols in Sections 4 and 5, Section 6 establishes a well-defined adversarial model and assesses their security under the RSA assumption and the collision-intractability of a hash function.

## 2 The $X\mu P$ 's Architecture and Instruction Set

XJVML. An executable program is modeled as a sequence of instructions  $P = (INS_1, \dots, INS_\ell)$  where  $INS_i$  is located at address  $i$  for  $i \in 1, \dots, \ell$  off-board. These instructions are in essence similar to instruction codes executed by any traditional microprocessor. Although the  $X\mu P$ 's instruction set could be similar to that of a 68HC05, MIPS32 or a MIX processor [10], we choose to model it as a JVMLO-like machine [13], extending this language into XJVML as follows. XJVML is a basic virtual processor operating on a volatile memory RAM, a non-volatile memory NVM, classical I/O ports denoted IO (for data) and XIO (for instructions), an internal random number generator denoted RNG and an operand stack ST, in which we distinguish

- **transfer instructions:** `load`  $x$  pushes the current value of  $RAM[x]$  (i.e. the memory cell at immediate address  $x$  in RAM) onto the operand stack. `store`

$x$  pops the top value off the operand stack and stores it at address  $x$  in RAM. Similarly, `load IO` captures the value presented at the I/O port and pushes it onto the operand stack whereas `store IO` pops the top value off the operand stack and sends it to the external world. `load RNG` generates a random number and pushes it onto the operand stack (the instruction `store RNG` does not exist). `getstatic` pushes `NVM[x]` onto the operand stack and `putstatic`  $x$  pops the top value off the operand stack and stores it into the nonvolatile memory at address  $x$ ;

- **arithmetic and logical operations:** `inc` increments the value on the top of the operand stack. `pop` pops the top of the operand stack. `push0` pushes the integer zero onto the operand stack. `xor` pops the two topmost values of the operand stack, exclusive-ors them and pushes the result onto the operand stack. `dec`'s effect on the topmost stack element is the exact opposite of `inc`. `mul` pops the two topmost values off the operand stack, multiplies them and pushes the result (two values representing the result's MSB and LSB parts) onto the operand stack;
- **control flow instructions:** letting  $1 \leq L \leq \ell$  be an instruction's index, `goto`  $L$  is a simple jump to program address  $L$ . Instruction `if`  $L$  pops the top value off the operand stack and either falls through when that value is the integer zero or jumps to  $L$  otherwise. The `halt` instruction halts execution.

Note that no program memory appears in our architecture: instructions are simply sent to the microprocessor which executes them in real time. To this end, a program counter  $i$  is maintained by the  $X\mu P$ :  $i$  is set to 1 upon reset and is updated by instructions themselves. Most of them simply increment  $i \leftarrow i + 1$  but control flow instructions may set  $i$  to arbitrary values in the range  $[1, \ell]$ . To request instruction `INS $i$` , the  $X\mu P$  simply sends  $i$  to the XT and receives `INS $i$`  via the specifically dedicated communication port XIO.

**SECURITY-CRITICAL INSTRUCTIONS.** While executing instructions, the device may be fed with misbehaving code crafted so as to read-out secrets from the NVM or even update the NVM at wish (for instance, illegally credit the balance of an e-Purse). It follows that the execution of instructions that have an irreversible effect on the device's NVM or on the external world must be authenticated in some way so as to validate their genuineness. For this reason we single-out the very few machine instructions that send signals out of the  $X\mu P^1$  and those instructions that modify the state of the  $X\mu P$ 's non-volatile memory<sup>2</sup>. These instructions will be called *security-critical* in the following sections and are defined as follows.

**Definition 1.** *A microprocessor instruction is security-critical if it might trigger the emission of an electrical signal to the external world or if it causes a*

<sup>1</sup> Typically the instruction allowing a data I/O port to toggle.

<sup>2</sup> Typically the latching of the control bit that triggers EEPROM/Flash update or erasure.

modification of the microprocessor’s internal nonvolatile memory. We denote by  $\mathcal{S}$  the set of security-critical instructions.

As we now see, posing  $\mathcal{S} = \{\text{putstatic } x, \text{ store IO}\}$  is not enough. Indeed, there exist subtle attacks that exploit  $i$  as a side channel. Consider the example below where  $k$  denotes the NVM address of a secret key byte  $u = \text{NVM}[k]$ :

$$P = (\text{getstatic } k, \text{ if } 1000, \text{ dec}, \text{ if } 1001, \text{ dec}, \text{ if } 1002, \dots).$$

The  $\text{X}\mu\text{P}$  will require from the  $\text{XT}$  a continuous sequence of instructions

$$\text{INS}_1, \text{INS}_2, \dots, \text{INS}_{u-1}, \text{INS}_u$$

followed by a sudden request of  $\text{INS}_{1000+u}$  and the value of  $u = \text{NVM}[k]$  has hence leaked-out.

Let us precisely formalize the problem: a microprocessor instruction is called *leaky* if it might cause a physically observable variable (e.g. the program counter) to take one of several possible values, depending on the data (RAM, NVM or ST element) handled by the instruction. The opposite notion is the one of *data indistinguishability* that characterizes those instructions for which the processed data have no influence whatsoever on environmental variables. Executing a `xor`, typically, does not reveal information (about the two topmost stack elements) which could be monitored from the outside of the  $\text{X}\mu\text{P}$ . As the execution of leaky instructions may reveal information about internal program variables, they fall under the definition of security-criticality and we therefore include them in  $\mathcal{S}$ . Following our instruction set, we have  $\mathcal{S} = \{\text{putstatic } x, \text{ store IO}, \text{ if } L\}$ .

### 3 Ensuring Program Authenticity

VERIFICATION PER INSTRUCTION. To ascertain that the instructions executed by the device are indeed those crafted by the code’s author, a naive approach consists in associating a signature to each instruction e.g. with RSA<sup>3</sup>. The program’s author generates a public and private RSA signature key-pair  $(N, e, d)$  and embeds  $(N, e)$  into the  $\text{X}\mu\text{P}$ . The code is enhanced with signatures  $P = ((\text{INS}_1, \sigma_1), \dots, (\text{INS}_\ell, \sigma_\ell))$  where  $\sigma_i = \mu(\text{ID}, i, \text{INS}_i)^d \bmod N$ ,  $\mu$  denotes a deterministic RSA padding function<sup>4</sup> and  $\text{ID}$  is a unique program identifier.

Note that the instruction address  $i$  appears in the padding function to avoid interchanging instructions in a program. The role of  $\text{ID}$  is to guard against code mixture attacks in which the  $i$ -th instructions of *two* programs are interchanged. The  $\text{X}\mu\text{P}$  keeps the  $\text{ID}$  of all authorized programs in nonvolatile memory. We consider the straightforward protocol shown on Figure 1.

<sup>3</sup> Any other signature scheme featuring high-speed verification could be used here.

<sup>4</sup> Note that if a message-recovery enabling padding is used, the storage of  $P$  can be reduced.

0. The  $X_{\mu}P$  receives and checks ID and initializes  $i \leftarrow 1$
1. The  $X_{\mu}P$  queries from the XT instruction number  $i$
2. The XT sends  $(INS_i, \sigma_i)$  to the  $X_{\mu}P$
3. The  $X_{\mu}P$ 
  - (a) ascertains that  $\sigma_i^e = \mu(ID, i, INS_i) \bmod N$
  - (b) executes  $INS_i$
4. Goto step 1.

**Fig. 1.** The Authenticated  $X_{\mu}P$ (inefficient)

This protocol is quite inefficient because, although verifying RSA signatures can be relatively easy with the help of a cryptocoprocessor, verifying one RSA signature per instruction remains resource-consuming.

**RSA-BASED SCREENING SCHEMES.** We resort to the *screening* technique devised by Bellare, Garay and Rabin in [4]. Unlike verification, screening ascertains that a batch of messages has been signed instead of checking that each and every signature in the batch is individually correct. More technically, the RSA-screening algorithm proposed in [4] works as follows. Given a list of message-signature pairs  $\{m_i, \sigma_i = h(m_i)^d \bmod N\}$ , one screens this list by simply checking that

$$\left( \prod_{i=1}^t \sigma_i \right)^e = \prod_{i=1}^t h(m_i) \bmod N \quad \text{and} \quad i \neq j \Leftrightarrow m_i \neq m_j .$$

At a first glance, this primitive seems to perfectly suit our code externalization problem where one does not necessarily need to ascertain that all the signatures are individually correct, but rather control that all the code ( $\{INS_i, \sigma_i\}$ ) seen by the  $X_{\mu}P$  has indeed been signed by the program’s author at some point in time.

Unfortunately the restriction  $i \neq j \Leftrightarrow m_i \neq m_j$  has a very important drawback as loops are extremely frequent in executable code (in other words, the  $X_{\mu}P$  may repeatedly require the same  $\{INS_i, \sigma_i\}$  while executing a given program)<sup>5</sup>. To overcome this limitation, we introduce a new screening variant where, instead of checking that each message appears only once in the list, the screener controls that the number of elements in the list is strictly smaller than  $e$  (we assume throughout the paper that  $e$  is a prime number) *i.e.* :

$$\left( \prod_{i=1}^t \sigma_i \right)^e = \prod_{i=1}^t \mu(m_i) \bmod N \quad \text{and} \quad t < e .$$

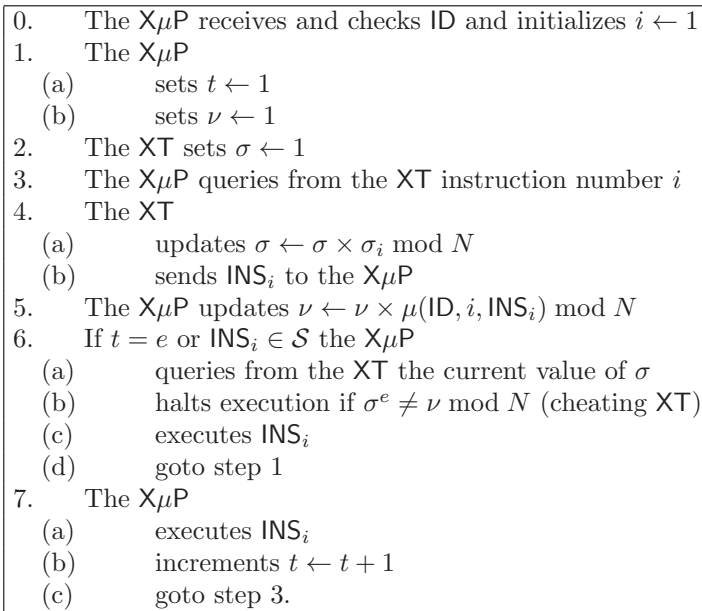
This screening scheme is referred to as  $\mu$ -RSA. The security of  $\mu$ -RSA for  $\mu = h$  where  $h$  is a full domain hash function, is guaranteed in the random oracle model [5] by the following theorem.

<sup>5</sup> Historically, [4] proposed only the criterion  $(\prod \sigma_i)^e = \prod \mu(m_i) \bmod N$ . This version was broken by Coron and Naccache in [9]. Bellare *et al.* subsequently repaired the scheme but the fix introduced the restriction that any message can appear at most once in the list.

**Theorem 1.** *Let  $(N, e)$  be an RSA public key where  $e$  is a prime number. If a forger  $\mathcal{F}$  can produce a list of  $t < e$  messages  $(m_1, \dots, m_t)$  and  $0 \leq \sigma < N$  such that  $\sigma^e = \prod_{i=1}^t h(m_i) \bmod N$  while the signature of at least one of  $m_1, \dots, m_t$  is not given to  $\mathcal{F}$ , then  $\mathcal{F}$  can be used to efficiently extract  $e$ -th roots modulo  $N$ .*

The theorem applies in both passive and active settings: in the former case,  $\mathcal{F}$  is given the list  $\{m_1, \dots, m_t\}$  as well as the signature of some of them. In the latter,  $\mathcal{F}$  is allowed to query a signing oracle and may choose the value of the  $m_i$ s. We refer the reader to [6, Appendix A.1] for a proof of Theorem 1 and detailed security reductions.

OPAQUE SCREENING. Signature screening is now used to verify instructions collectively as depicted on Figure 2. At any point in time,  $\nu$  is an accumulated product of  $t < e$  padded instructions  $\nu = \prod_i \mu(\text{ID}, i, \text{INS}_i)$ . Loosely speaking, both parties  $X\mu P$  and  $XT$  update their own security buffers  $\nu$  and  $\sigma$  which compatibility (in the sense of  $\sigma^e = \nu \bmod N$ ) is checked before executing any security-critical instruction. Note that a verification is also triggered when exactly  $e - 1$  instructions are aggregated in  $\nu$ .



**Fig. 2.** The Opaque  $X\mu P$  (secure but suboptimal)

As one can easily imagine, this protocol becomes rapidly inefficient when instructions of  $\mathcal{S}$  are frequently used. For instance, **ifs** constitute the basic ingredient of **while** and **for** assertions which are extremely common in executable code. Moreover, in many cases, **whiles** and **fors** are even nested or interwoven. It

follows that the Opaque  $X\mu P$  would incessantly trigger the relatively expensive<sup>6</sup> verification stage of steps 6a and 6b (we denote by `CheckOut` this verification stage throughout the rest of the paper). This is clearly an overkill: in many cases `ifs` can be safely performed on non secret data dependent<sup>7</sup> variables (for instance the variable that counts 16 rounds during a DES computation). We show in the next section how to optimize the number of `CheckOuts` while keeping the protocol secure.

## 4 Internal Security Policies

We now associate a *privacy bit* to each memory and stack cells, denoting by  $\varphi(\text{RAM}[j])$ ,  $\varphi(\text{NVM}[j])$  and  $\varphi(\text{ST}[j])$  the privacy bit associated to `RAM[j]`, `NVM[j]` and `ST[j]`. NVM privacy bits are nonvolatile. Informally speaking, the idea behind privacy bit is to prevent the external world from probing secret data handled by the  $X\mu P$ . RAM privacy bits are initialized to zero upon reset, NVM privacy bits are set to zero or one by the  $X\mu P$ 's issuer at the production or personalization stage,  $\varphi(\text{IO})$  and  $\varphi(\text{RNG})$  are always stuck to zero<sup>8</sup> and one by definition and privacy bits of released stack elements are automatically reset to zero.

We also introduce simple rules by which the privacy bits of new variables evolve as a function of prior  $\varphi$  values. Transfer instructions simply transfer the privacy bit of their variable (e.g. `getstatic 3` simultaneously sets  $\text{ST}[s] \leftarrow \text{NVM}[3]$  and  $\varphi(\text{ST}[s]) \leftarrow \varphi(\text{NVM}[3])$  where  $s$  denotes the stack pointer and `ST[s]` the topmost stack element). The rule we apply to arithmetical and logical instructions is *privacy-conservative* namely, the output privacy bits are all set to zero if and only if all input privacy bits were zero (otherwise they are all set to one). In other words, as soon as private data enter a computation all output data are tagged as private. This rule is easily hardwired as a simple boolean OR for non-unary operators.

This mechanism allows to process security-critical instructions in different ways depending on whether they run over private or non-private data. Typically, executing an `if L` does not provide critical information if the topmost stack element is non-private. A `CheckOut` may not be mandatorily invoked in this case. Accordingly, outputting a non-private value via a `store IO` instruction does not provide any sensitive information, and a `CheckOut` can be spared in this case as well. In fact, one can easily specify a *security policy* that contextually defines the conditions (over privacy bits) under which a security-critical instruction may or may not trigger a collective verification. To abstract away the security policy chosen by the issuer, we introduce the boolean predicate

$$\text{Alert} : \mathcal{S} \times \Phi \mapsto \{\text{True}, \text{False}\}$$

<sup>6</sup> While the execution of a regular instruction demands only one modular multiplication, the execution of an  $\text{INS}_i \in \mathcal{S}$  requires the transmission of an RSA signature (e.g. 1024 bits) and an exponentiation (e.g. to the power  $e = 2^{16} + 1$ ) in the  $X\mu P$ .

<sup>7</sup> Read: non-((secret-data)-dependent).

<sup>8</sup> i.e. any external data fed into the  $X\mu P$  is considered as publicly observable by opponents and hence non-private.

where  $\Phi$  denotes the set of all privacy bits  $\Phi = \varphi(\text{RAM}) \cup \varphi(\text{NVM}) \cup \varphi(\text{ST})$ .  $\text{Alert}(\text{INS}, \Phi)$  evaluates as True when a CheckOut is to be invoked. We hence twitch our protocol as now shown on Figure 3.

0. The  $X\mu P$  receives and checks ID and initializes  $i \leftarrow 1$
1. The  $X\mu P$ 
  - (a) sets  $t \leftarrow 1$
  - (b) sets  $\nu \leftarrow 1$
2. The XT sets  $\sigma \leftarrow 1$
3. The  $X\mu P$  queries from the XT instruction number  $i$
4. The XT
  - (a) updates  $\sigma \leftarrow \sigma \times \sigma_i \bmod N$
  - (b) sends  $\text{INS}_i$  to the  $X\mu P$
5. The  $X\mu P$  updates  $\nu \leftarrow \nu \times \mu(\text{ID}, i, \text{INS}_i) \bmod N$
6. If  $t = e$  or  $(\text{INS}_i \in \mathcal{S} \text{ and } \text{Alert}(\text{INS}_i, \Phi))$  the  $X\mu P$ 
  - (a) CheckOut
  - (b) executes  $\text{INS}_i$
  - (c) goto step 1
7. The  $X\mu P$ 
  - (a) executes  $\text{INS}_i$
  - (b) increments  $t \leftarrow t + 1$
  - (c) goto step 3.

**Fig. 3.** Enforcing a Security Policy: Protocol 1

## 5 Authenticating Code Sections Instead of Instructions

Following the classical definition of [1,11], we call a *basic block* a straight-line sequence of instructions that can be entered only at its beginning and exited only at its end. The set of basic blocks of a program  $P$  is usually given under the form of a graph  $\text{CFG}(P)$  and computed by the means of control flow analysis [12, 11]. In such a graph, vertices are basic blocks and edges symbolize control flow dependencies:  $B_0 \rightarrow B_1$  means that the last instruction of  $B_0$  may handover control to the first instruction of  $B_1$ . In our instruction set, basic blocks admit at most two sons with respect to control flow dependance; a block has two sons if and only if its last instruction is an *if*. When  $B_0 \rightarrow B_1$ ,  $B_0 \Rightarrow B_1$  means that  $B_0$  has no son but  $B_1$  (but  $B_1$  may have other fathers than  $B_0$ ). In this section we define a slightly different notion that we call *code sections*.

Informally, a code section is a maximal collection of basic blocks  $B_1 \Rightarrow B_2 \cdots \Rightarrow B_\ell$  such that no instruction of  $\mathcal{S} \cup \{\text{halt}\}$  appears in the blocks except, possibly, as the last instruction of  $B_\ell$ . The section is then denoted by  $S = \langle B_1, \dots, B_\ell \rangle$ . In a code section, the control flow is deterministic *i.e.* independent from program variables; thus a section may contain several cascading *goto* instructions. Code sections, unlike basic blocks, may share instructions; yet they have a natural graph structure induced by  $\text{CFG}(P)$  which we do not use in the sequel. It is known that computing a program's basic blocks can be



done in almost-linear time [12] and it is easily seen that the same holds for code sections. We refer to the full version of this work for an algorithm computing the set  $\text{Sec}(P)$  of code sections of a program  $P$ .

Given that instructions in a code section are executed sequentially, and that sections can be computed at compile time, signatures can certify sections rather than individual instructions. In other words, a single signature per code section suffices. The signature of a code section  $S$  starting at address  $i$  is:

$$\sigma_i = \mu(\text{ID}, i, h)^d \pmod{N},$$

with  $h = H(\text{INS}_1, \dots, \text{INS}_k)$  where  $\text{INS}_1, \dots, \text{INS}_k$  are the successive instructions in  $S$ . Here,  $H$  is an iterative hash function recursively defined by  $H(x_1, \dots, x_j) = F(x_j, H(x_1, \dots, x_{j-1}))$  and  $H(x_1) = F(x_1, IV)$  where  $F(x, y)$  is  $H$ 's compression function and  $IV$  an initialization constant. We summarize the new protocol on Figure 4.

- |    |  |
|----|--|
| 0. | The $X\mu P$ receives and checks ID and initializes $i \leftarrow 1$                                     |
| 1. | The $X\mu P$   |
|    | (a) sets $t \leftarrow 1$ ( <i>t now counts code sections</i> )  |
|    | (b) sets $\nu \leftarrow 1$  |
| 2. | The XT sets $\sigma \leftarrow 1$  |
| 3. | The $X\mu P$   |
|    | (a) sets $h \leftarrow IV$   |
|    | (b) queries the code section starting at address $i$   |
| 4. | The XT   |
|    | (a) updates $\sigma \leftarrow \sigma \times \sigma_i \pmod{N}$  |
|    | (b) sets $j = 1$   |
| 5. | The XT   |
|    | (a) sends $\text{INS}_j^i$ to the $X\mu P$   |
|    | (b) increments $j \leftarrow j + 1$  |
| 6. | The $X\mu P$   |
|    | (a) receives $\text{INS}_j^i$ ,  |
|    | (b) updates $h \leftarrow F(\text{INS}_j^i, h)$  |
| 7. | If $\text{INS}_j^i \in \mathcal{S}$ and ( $\text{Alert}(\text{INS}_j^i, \Phi)$ or $t = e$ ) the $X\mu P$ |
|    | (a) sets $\nu = \nu \times \mu(\text{ID}, i, h) \pmod{N}$  |
|    | (b) CheckOut   |
|    | (c) executes $\text{INS}_j^i$  |
|    | (d) goto step 1  |
| 8. | Else if $\text{INS}_j^i \in \mathcal{S}$ then the $X\mu P$   |
|    | (a) sets $\nu = \nu \times \mu(\text{ID}, i, h) \pmod{N}$  |
|    | (b) increments $t \leftarrow t + 1$  |
|    | (c) executes $\text{INS}_j^i$  |
|    | (d) goto step 3  |
| 9. | Else the $X\mu P$  |
|    | (a) executes $\text{INS}_j^i$  |
|    | (b) increments $j \leftarrow j + 1$  |
|    | (c) goto step 5.   |

**Fig. 4.** Authentication of Code Sections: Protocol 2

This protocol presents the advantage of being far less time consuming, because the number of `CheckOuts` (and updates of  $\nu$ ) is considerably reduced. The formats under which the code can be stored in the `XT` are diverse. The simplest of these consists in representing  $P$  as the list of all its signed code sections  $P = (\text{ID}, (1, \sigma_1, S_1), \dots, (k, \sigma_k, S_k))$ . Whatever the file format used in conjunction with our protocol is, the term *authenticated program* designates a program augmented with its signature material  $\Sigma(P) = \{\sigma_i\}_i$ . Thus, our protocols actually execute authenticated programs. A program is converted into an authenticated executable file via a specific compilation phase involving both code processing and signature generations.

## 6 Security Analysis

What we provide in this section is a formal proof that the protocols described above are secure. The security proof shall have two ingredients: a well-defined security model describing an adversary's goal and resources, and a reduction from some complexity-theoretic hard problem. Rather than rigorously introducing the numerous notions our security model is based upon (which the reader may find in [6], as well as the fully detailed reductions), we give here a high-level description of our security analysis.

**THE SECURITY MODEL.** We assume the existence of three parties in the game:

- a code issuer  $CI$  that compiles `XJVML` programs into authenticated executable files with the help of the signing key  $(N, d)$ ,
- an  $X\mu P$  that follows the communication protocol given in Section 4 and contains the verification key  $(N, e)$  matching  $(N, d)$ . The  $X\mu P$  also possesses some cryptographic private key material  $k$  stored in its NVM,
- an attacker  $\mathcal{A}$  willing to access  $k$  using means that are discussed below.

**ADVERSARIAL GOALS.** Depending on the role played by the  $X\mu P$ 's cryptographic key  $k$ , the adversary's goals might be of different nature. Of course, inferring information about  $k$  (worse, recovering  $k$  completely) comes immediately to one's mind, but there could also be weaker (somewhat easier) ways of having access to  $k$ . For instance if  $k$  is a symmetric encryption key,  $\mathcal{A}$  might try to decrypt ciphertexts encrypted under  $k$ . Similarly, if it is a public-key signature key,  $\mathcal{A}$  could attempt to rely on the protocol engaged with the  $X\mu P$  to help forging signatures in a way or an other. More exotically, the adversary could try to *hijack* the key  $k$  e.g. to use it (or a part of it thereof) as an AES key whereas  $k$  was intended to be employed some other way.  $\mathcal{A}$ 's goal in this case is a bit more intricate to capture, but we see no reason why we should prohibit that kind of scenario in our security model. Third, the adversary may attempt to modify  $k$ , thereby opening the door to fault attacks [2,3].

**THE ATTACK SCENARIO.** Parties behave as follows. The  $CI$  crafts polynomially many authenticated programs of polynomially bounded size and publishes them.

We assume no interaction between the  $CI$  and  $\mathcal{A}$ . Then  $\mathcal{A}$  and the  $X\mu P$  engage in the protocol and  $\mathcal{A}$  attempts to make the  $X\mu P$  execute a sequence of instructions  $\xi$  that was not originally issued by the  $CI$ . The attack succeeds when  $\xi$  contains a security-critical instruction that handles some part of  $k$  which the  $X\mu P$  nevertheless executes.

We say that  $\mathcal{A}$  is an  $(\ell, n, \tau, \varepsilon)$ -attacker if after seeing at most  $\ell$  authenticated programs  $P_1, \dots, P_\ell$  totalling at most  $n \geq \ell$  instructions and processing at most  $\tau$  steps,  $\Pr[\mathcal{A} \text{ succeeds}] \geq \varepsilon$ . In this definition, we include in  $\tau$  the execution time  $\text{Time}(\xi)$  of  $\xi$ , stipulating by convention that executing each instruction takes one step and that all transmissions (instruction addresses, instructions, signatures and IO data) are instantaneous.

SECURITY PROOF FOR PROTOCOL 1. We state:

**Theorem 2.** *If the screening scheme  $\mu$ -RSA is  $(q_k, \tau, \varepsilon)$ -secure against existential forgery under a known message attack, then Protocol 1 is  $(\ell, n, \tau, \varepsilon)$ -secure for  $n \leq q_k$ .*

Moreover, when  $\mu = \text{FDH}$ , outputting a valid forgery is equivalent to extracting  $e$ -th roots modulo  $N$  as shown in [6, Appendix A.1]. The following corollary is proved by invoking Theorem 1.

**Corollary 1.** *If  $\mu$  is a full domain hash function, then Protocol 1 is secure under the RSA assumption in the random oracle model.*

SECURITY PROOF FOR PROTOCOL 2. We now move on to the (more efficient) Protocol 2 defined in Section 5.  $(\mu, H)$ -RSA is defined as being the RSA screening scheme with padding function  $(x, y, z) \mapsto \mu(x, y, H(z))$ . We slightly redefine  $(\ell, n, \tau, \varepsilon)$ -security as the resistance against adversaries that have access to at most  $\ell$  authenticated programs totalling at most  $n$  code sections. We state:

**Theorem 3.** *If the screening scheme  $(\mu, H)$ -RSA is  $(q_k, \tau, \varepsilon)$ -secure against existential forgery under a known message attack, then Protocol 2 is  $(\ell, n, \tau, \varepsilon)$ -secure for  $n \leq q_k$ .*

When  $\mu(a, b, c) = h(a||b||H(c))$  and  $h$  is seen as a random oracle, a security result similar to Corollary 1 can be obtained for Protocol 2. However, a bad choice for  $H$  could allow the adversary  $\mathcal{A}$  to easily find collisions over  $\mu$  via collisions over  $H$ . Nevertheless, unforgeability can be formally proved under the assumption that  $H$  is collision-intractable. We refer the reader to the corresponding theorem given in [6, Appendix B]. Associating this result with Theorem 3, we conclude:

**Corollary 2.** *Assume  $\mu(a, b, c) = h(a||b||H(c))$  where  $h$  is a full-domain hash function seen as a random oracle. Then Protocol 2 is secure under the RSA assumption and the collision-intractability of  $H$ .*

WHAT ABOUT ACTIVE ATTACKS? Although RSA-based screening schemes may feature strong unforgeability under chosen-message attacks (see [6, Appendix A.2] for such a proof for FDH-RSA), it is easy to see that our protocols cannot resist chosen-message attackers whatever the security level of the underlying screening scheme happens to be. Indeed, assuming that the adversary is allowed to query the code issuer  $CI$  with messages of her choosing, a trivial attack consists in obtaining the signature

$$\sigma = \mu(\text{ID}, 1, H(\text{INS}_1, \text{INS}_2, \text{INS}_3))^d \pmod N$$

of a program  $P$  where ID is known to be accepted by the  $X\mu P$  and the single-section program  $P$  is

$$P = (\text{getstatic } 17, \text{store } \text{I0}, \text{halt})$$

wherein NVM[17] is known to contain a fraction of the cryptographic key  $k$ , the value 17 being purely illustrative here<sup>9</sup>. Similarly, the attacker may query the signature of some trivial key-modifying code sequence. Obviously, nothing can be done to resist chosen-message attacks.

## 7 Deployment Considerations and Engineering Options

From a practical engineering perspective, our new architecture is likely to deeply impact the smart card industry. We briefly discuss some advantages of our technology.

CODE PATCHING. A bug in a program does not imply the roll-out of devices in the field but a simple terminal update. Patching a future smart card can hence become as easy as patching a PC. A possible bug patching mechanism consists in encoding in ID a backward compatibility policy signed by the  $CI$  that either instructs the  $X\mu P$  to replace its old ID by a new one and stop accepting older version programs or allow the execution of new or old code (each at a time, i.e. no blending possible). The description of this mechanism is straightforward and omitted here.

CODE SECRECY. Given that the XT contains the application's code, our architecture assumes that the algorithm's specifications are public. It is possible to reach *some* level of secrecy by encrypting the XT's program under a key (common to all  $X\mu P$ s). Obviously, morphologic information about the algorithm will leak out to some extent (loop structure *etc.*) but important elements such as S-box contents or the actual type of boolean operators used by the code could remain confidential if programmed appropriately.

SIMPLIFIED PRODUCT MANAGEMENT. Given that a GSM  $X\mu P$  and an electronic-purse  $X\mu P$  differ only by a few NVM bytes (essentially ID), by opposition to smart-cards,  $X\mu P$ s are real commodity products (such as capacitors,

<sup>9</sup> The `halt` instruction is even superfluous as the attacker can power off the device right after the second instruction is executed.

resistors or Pentium processors) which stock management is greatly simplified and straightforward. Given the very small NVM room needed to store an ID and a public-key, a single  $X\mu P$  can very easily support several applications provided that the sum of the NVM spaces used by these applications does not exceed the  $X\mu P$ 's total NVM capacity and that these NVM spaces are properly firewalled. From the user's perspective the  $X\mu P$  is tantamount to a key ring carrying all the secrets (credentials) used by the applications that the user interacts with but *not* these applications themselves.

A wide range of trade-offs and variants is possible when implementing the architecture described in this paper. Referring to the extended version of this work [6] for more, a few engineering options are considered here.

**SPEEDING UP MODULAR OPERATIONS.** While the multiplication of two  $\kappa$ -bit integers theoretically requires  $\kappa^2$  operations, multiplying a random  $\nu$  by  $\mu(x)$  may require only  $\kappa^2/4$  operations when  $\mu$  is adequately chosen. Independently, an adequate usage of RAM counters allows to decrease the value of  $e$  without sensibly increasing the expected number of CheckOut on the average.

**REPLACING RSA.** Clearly, any signature scheme that admits a screening variant (*i.e.* a homomorphic property) can be used in our protocols. RSA features a low (and customizable) verification time, but replacing it by EC-based schemes for instance, could present some advantages.

**CODE SIZE VERSUS EXECUTION SPEED.** The access to a virtually unlimited ROM renders vacuous the classical dilemma between optimizing code size or speed. Here, for instance, one can cheaply unwind (inline) loops or implement algorithms using pre-computed space-consuming look-up tables instead of performing on-line calculations *etc.*

**SMART USAGE OF SECURITY HARDWARE FEATURES.** Using the Alert predicate, the  $X\mu P$  could selectively activate hardware-level protections against physical attacks whenever a private variable is handled or forecasted to be used a few cycles later.

**HIGH SPEED XIO.** A high-speed communication interface is paramount for servicing the extensive information exchange between the  $X\mu P$  and the XT. Evaluating transmission performances for a popular standard, the Universal Serial Bus (USB)<sup>10</sup>, we found that transfers of 32 bits can be done at 25 Mb/s in USB High Speed mode which corresponds to 780K 32-bit words per second. When servicing Protocol 1, this corresponds approximately to a 32-bit  $X\mu P$  working at 390 KHz; when parallel execution and look-ahead transmission take place, one gets a 32-bit machine running at 780 KHz. An 8-bit USB interface leads to 830 KHz. There is no doubt that these figures can be greatly improved.

---

<sup>10</sup> Note that USB is unadapted to our application as this standard was designed for good bandwidth rather than for good latency.

## 8 Further Work

The authors believe that the concept introduced in this paper raises a number of practical and theoretical questions. Amongst these is the safe externalization of Java's *entire* bytecode set, the safe co-operative development of code by competing parties (*i.e.* mechanisms for the secure handover of execution from program  $ID_1$  to program  $ID_2$ ), or the devising of faster execution protocols.

Interestingly, the paradigm of signature screening on which Protocols 1 and 2 are based also exists in the symmetric setting, where RSA signatures are replaced by MACs and a few hash functions. Security can also be assessed formally in this case under adequate assumptions. We refer the reader to [6] for details.

This paper showed how to provably securely externalize programs from the processor that runs them. Apart from answering a theoretical question, we believe that our technique provides the framework of novel practical solutions for real-life applications in the world of mobile code and cryptography-enabled embedded software.

## References

1. A. Aho, R. Sethi, J. Ullman, *Compilers: Principles, Techniques, and Tools*, Addison-Wesley, 1986.
2. E. Biham and A. Shamir, *Differential Fault Analysis of Secret Key Cryptosystems*, In *Advances in Cryptography, Crypto'97*, LNCS 1294, pages 513–525, 1997.
3. I. Biehl, B. Meyer and V. Müller, *Differential Fault Attacks on Elliptic Curve Cryptosystems*, In M. Bellare (Ed.), *Proceedings of Advances in Cryptology, Crypto 2000*, LNCS 1880, pages 131–146, Springer Verlag, 2000.
4. M. Bellare, J. Garay and T. Rabin, *Fast Batch Verification for Modular Exponentiation and Digital Signatures*, Eurocrypt'98, LNCS 1403, pages 236–250. Springer-Verlag, Berlin, 1998.
5. M. Bellare and P. Rogaway, *Random Oracles Are Practical: a Paradigm for Designing Efficient Protocols*, *Proceedings of the first CCS*, pages 62–73. ACM Press, New York, 1993.
6. B. Chevallier-Mames, D. Naccache, P. Paillier and D. Pointcheval, *How to Disembed a Program?*, IACR ePrint Archive, <http://eprint.iacr.org>, 2004.
7. Z. Chen, *Java Card Technology for Smart Cards: Architecture and Programmer's Guide*, The Java Series, Addison-Wesley, 2000.
8. J.-S. Coron, *On the Exact Security of Full-Domain-Hash*, *Crypto'2000*, LNCS 1880, Springer-Verlag, Berlin, 2000.
9. J.-S. Coron and D. Naccache, *On the Security of RSA Screening*, *Proceedings of the Fifth CCS*, pages 197–203, ACM Press, New York, 1998.
10. D.E. Knuth, *The Art of Computer Programming, vol. 1, Seminumerical Algorithms*, Addison-Wesley, Third edition, pages 124–185, 1997.
11. S. Muchnick, *Advanced Compiler Design and Implementation*, Morgan Kaufmann, 1997.
12. G. Ramalingam, *Identifying Loops in Almost Linear Time*, *ACM Transactions on Programming Languages and Systems*, 21(2):175–188, March 1999.
13. R. Stata and M. Abadi, *A Type System for Java Bytecode Subroutines*, SRC Research Report 158, June 11, 1998, <http://www.research.digital.com/SRC/>.