# Checking Secure Interactions of Smart Card Applets: extended version*

P. Bieber[1], J. Cazin[1], P. Girard[2], J.-L. Lanet[2], V. Wiels[1], G. Zanon[1]

[1] ONERA-CERT/DTIM
BP 4025, 2 avenue E. Belin,
F-31055 Toulouse Cedex 4, France
{bieber,cazin,wiels,zanon}@cert.fr
[2] GEMPLUS
avenue du pic de Bertagne, 13881 Gémenos cedex, France
{Pierre.GIRARD,Jean-Louis.LANET}@gemplus.com

**Abstract.** This paper presents an approach enabling a smart card issuer to verify that a new applet securely interacts with already downloaded applets. A security policy has been defined that associates levels to applet attributes and methods and defines authorized flows between levels. We propose a technique based on model checking to verify that actual information flows between applets are authorized. We illustrate our approach on applets involved in an electronic purse running on Java enabled smart cards.

## Introduction

A new type of smart cards is getting more and more attractive: multiapplication smart cards. The main characteristics of such cards are that applications can be loaded after the card issuance and that several applications run on the same card. A few operating systems have been proposed to manage multiapplication smart cards, namely Java Card [1], Multos [2] and more recently Windows for Smart Cards [3]. In this paper, we will focus on Java Card. Following this standard, applications for multiapplication smart cards are implemented as interacting Java applets.

Multiapplication smart cards involve several participants: the card provider, the card issuer that proposes the card to the users, application providers and card holders (users). The card issuer is usually considered responsible for the security of the card. The card issuer does not trust application providers: applets could be malicious or simply faulty.

As in a classical mobile code setting, a malicious downloaded applet could try to observe, alter, use information or resources it is not authorized to. Of course, a set of JavaCard security functions were defined that severely restrict what an

* The Pacap project is partially funded by MENRT décision d'aide 98.B.0251
[1] http://java.sun.com/products/javacard
[2] http://www.multos.com
[3] http://www.microsoft.com/smartcard

applet can do. But these functions do not cover a class of threats we call illicit applet interactions that cause unauthorized information flows between applets.

Our goal is to provide techniques and tools enabling the card issuer to verify that new applets interact securely with already loaded applets.

The first section introduces security concerns related to multiapplication smart cards. The second section of the paper describes the electronic purse functionalities and defines the threats associated with this application. The third section presents the security policy and information flow property we selected to verify that applets interact securely. The fourth section shows how we verify secure interaction properties on the applets byte-code. The fifth section relates our approach to other existing work.

# 1 Security concerns

## 1.1 Java Card security mechanisms

Security is always a big concern for smart cards but it is all the more important with multiapplication smart cards and post issuance code downloading.

Opposed to monoapplicative smart cards where Operating System and application were mixed, multiapplication smart card have drawn a clear border between the operating system, the virtual machine and the applicative code. In this context, it is necessary to distinguish the security of the card (hardware, operating system and virtual machine) from the security of the application. The card issuer is responsible for the security of the card and the application provider is responsible for the applet security, which relies necessarily on the security of the card.

The physical security is obtained by the smart card media and its tamper resistance. The security properties that the OS guarantees are the quality of the cryptographic mechanisms (which should be leakage resistant, i.e. resistant against side channel attacks such Differential Power Analysis [9]), the correctness of memory and I/O management.

A Java Card virtual machine relies on the type safety of the Java language to guarantee the innocuousness of an applet with respect to the OS, the virtual machine [11], and the other applets. However this is guaranteed by a byte-code verifier which is not on board, so extra mechanisms have been added. A secure loader (like OP [18]) checks before loading an applet that it has been signed (and then verified) by an authorized entity (namely the card issuer). Even if an unverified applet is successfully loaded on the card, the card firewall [17], which is part of the virtual machine, will still deny to an aggressive applet the possibility to manipulate data outside its memory space.

To allow the development of multiapplication smart cards, the Java Card has introduced a new way for applets to interact directly. An applet can invoke another applet method through a shared interface. An applet can decide to share or not some data with a requesting applet based on its identifier.

### 1.2 Applets providers and end users security needs

Applet providers have numerous security requirements for their applications. Classical one are secret key confidentiality, protection from aggressive applets, integrity of highly sensitive data fields such as electronic purses balances, etc. These requirement are widely covered by the existing security mechanisms at various levels from silicon to secure loaders and are not covered in this paper. However, new security requirements appear with the growing complexity of applets and the ability for applets to interact directly with other applets.

Particularly, application providers do not want information to flow freely inside the card. They want to be sure that the commercial information they provide such as marketing information and other valuable data (especially short term ones such as stock quotes, weather forecast and so on) won't be retransmitted by applets without their consent. For example, marketing information will be collected by a loyalty applet when an end user buys some goods at a retailer. This information will be shared with partner applets but certainly not with competitor applets. However it will certainly be the case that some partners applets will interact with competitor applets. As in the real world trust between providers should not be transitive and so should be the authorized information flows.

It is also expected that pay services inside the card, as information, won't be able to be re-sold by rogue applets. In a very constrained environment, services, even as little as an array sorting method can be valuable as they can help other applets to save precious non volatile memory space and thus money.

So far we have just talked about confidentiality constraints, but integrity should also be considered. For example, corruption of data or of service outputs by an aggressive applet would be an extremely damaging attack for the brand image of an applet provider well-known for its information accuracy and quality of service.

Finally, the end user security requirement will be related mainly with privacy. As soon as medical data or credit card record are handled by the card and transmitted between applets great care should be taken with the information flow ([4] details such a privacy threat).

### 1.3 Applet Certification

Applet providers and end users cannot control that their information flow requirements are enforced on the card because they do not manage it. Our goal is to provide techniques and tools enabling the card issuer to verify that new applets respect existing security properties defined as authorized information flows. The approach is described on figure 1.

If the applet provider wants to load a new applet on a card, she provides the bytecode for this applet. The card issuer has a security policy for the card and security properties that must be satisfied. This security policy should take into account the requirements of applet providers and end users. We provide techniques and tools to decide whether the properties are satisfied by the new

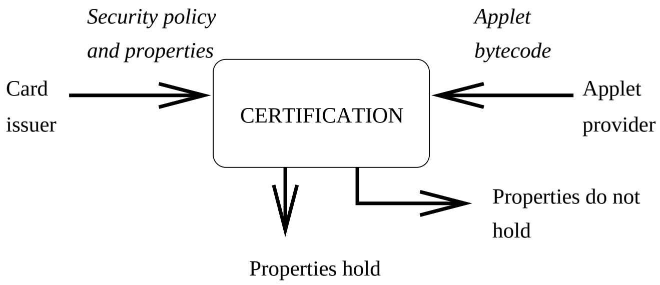


**Fig. 1.** Applet Certification

applet (these techniques are applied on the applet bytecode). If the properties hold, the applet can be loaded on the card; if they do not hold, it is rejected.

# 2 Electronic Purse

## 2.1 Electronic Purse Functionalities

A typical example of a multiapplication smart card is an electronic purse with one purse applet and two loyalty applets: a frequent flyer (Air France) application and a car rental (RentaCar) loyalty program. The purse applet manages debit and credit operations and keeps a log of all the transactions. As several currencies can be used (francs and euros for example) this applet also manages a conversion table. When the card owner wants to subscribe to a loyalty program, the corresponding loyalty applet is loaded on the card. This applet must be able to interact with the purse to get to know the transactions made by the purse in order to update loyalty points according to these transactions. For instance, the Air France applet will add miles to the account of the card owner whenever an Air France ticket is bought with the purse. The card owner can use these miles to buy a discounted Air France ticket. Agreements may also exist between loyalty applets to allow exchanges of points. For instance, loyalty points granted by RentaCar could be summed with Air France miles to buy a discounted ticket.

The electronic purse has been chosen as case study for our project. It has been implemented in Java by Gemplus.

## 2.2 Electronic Purse Threats

We suppose that all the relevant Java and JavaCard security functions are used in the electronic purse. But these functions do not cover all the threats. We are especially interested in threats particular to multiapplication smart cards like illicit applet interactions. An example of illicit interaction in the case of the electronic purse is described on figure 3.

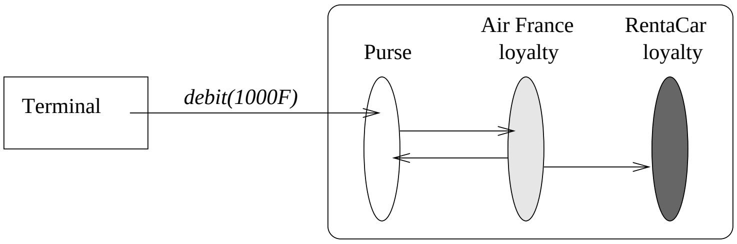


**Fig. 2.** An electronic purse

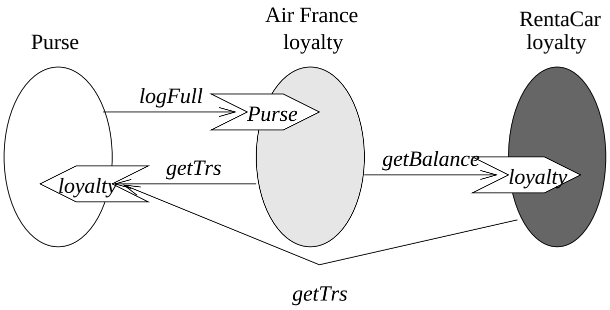


**Fig. 3.** Applet Interactions

The purse applet has a shared interface for loyalty applets to get their transactions and the loyalty applet has a shared interface for partner loyalty applets to get loyalty points.

A "logfull" service is proposed by the purse to the loyalty applets: when the transaction log is full, the purse calls the $logfull$ method of the loyalty applets that subscribed to the service to warn them that the log is full and they should get the transactions before some of them are erased and replaced by new ones. We suppose the Air France applet subscribed to the logfull service, but the RentaCar applet did not. When the log is full, the purse calls the $logfull$ method of the Air France applet. In this method, the Air France applet gets transactions from the purse but also wants to update its extended balance that contains its points plus all the points it can get from its loyalty partners. To update this extended balance, it calls the $getbalance$ method of the RentaCar loyalty applet. In that case, the car rental applet can guess that the log is full when the Air France applet calls its $getbalance$ method and thus get the transactions from the purse. There is a leak of information from the Air France applet to the RentaCar one and we want to be able to detect such illicit information flows. This illicit behaviour would not be countered by the applet firewall as all the invoked methods belong to shared interfaces.

## 3 Multiapplication Security Policy

To implement the applet certification approach, we first have to choose a security policy adapted to multiapplication cards and associated security properties.

### 3.1 Security policy

We propose to use a multilevel security policy [4] that was designed for multiapplication smart cards. Each applet provider is assigned a security level and we consider special levels for shared data. On the example of the electronic purse, we have a level for each applet: $AF$ for Air France, $P$ for purse and $RC$ for RentaCar and levels for shared data: $AF + RC$ for data shared by Air France and RentaCar, $AF + P$ for data shared by Air France and purse, etc. The relation between levels $\preceq$ is used to authorize or forbid information flows between applets. In the policy we consider, $AF + P \preceq AF$ and $AF + P \preceq P$, this means that information whose level is $AF+P$ is authorized to flow towards information whose level is $P$ or $AF$. So shared information from Air France and Purse may be received by Air France and Purse applets. To model that applets may only communicate through shared interfaces, direct flows between levels $AF$, $P$ and $RC$ are forbidden. So we have: $AF \not\preceq P$, $P \not\preceq AF$, $AF \not\preceq RC$, $RC \not\preceq AF$, $P \not\preceq RC$ and $RC \not\preceq P$.

The levels together with the $\preceq$ relation have a lattice structure, so there are a bottom level *public* and a top level *private*.

Figure 4 shows some of the authorized information flows.

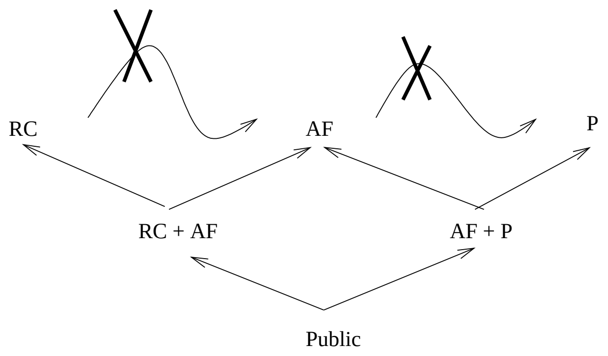


**Fig. 4.** Authorized information flows

### 3.2 Security properties

Now we have to define the security properties to be enforced. We have chosen the secure dependency model [1] that applies to systems where malicious applications might communicate confidential information to other applications. Like other information flow models such as non-interference [5], this model ensures that dependencies between system objects cannot be exploited to establish an indirect communication channel. We apply this model to the electronic purse: illicit interactions will be detected by controlling the dependencies between objects of the system.

A program is described by a set of evolutions that associate a value with each object at each date. We note $Ev \subseteq Objects \times Dates \rightarrow Values$ the set of evolutions of a program. The set $Objects \times Dates$ is made of three disjoint subsets: input objects that are not computed by the program, output objects that are computed by the program and are directly observable and internal objects that are not observable. We assume that function $lvl$ associates a security level with input and output objects.

The secure dependency property $SecDep$ requires that the value of output objects with security level $l$ only depends on the value of input objects whose security level is dominated by $l$:
$\forall o_t \in Output, \forall e \in Ev, \forall e' \in Ev, e \sim_{aut(o_t)} e' \Rightarrow e(o_t) = e'(o_t)$
where $aut(o_t) = \{o'_{t'} \in Input \mid t' < t, lvl(o'_{t'}) \preceq lvl(o_t)\}$ and $e \sim_{aut(o_t)} e'$ iff $\forall o'_{t'} \in aut(o_t), e(o'_{t'}) = e'(o'_{t'})$.

This property cannot be directly proved with a model checker such as SMV ([12]) because it is neither a safety or liveness property nor a refinement property. So we look for sufficient conditions of $SecDep$ that are better handled by SMV. By analysing the various instructions in a program, it is easy to compute for each object the set of objects it syntactically depends on. The set $dep(i, o_t)$ contains objects with date $t-1$ used by instruction at program location $i$ to compute the value of $o_t$. The program counter is an internal object such that $pc_{t-1}$ determines the current instruction used to compute the value of $o_t$. Whenever an object is modified (i.e. $o_{t-1}$ is different form $o_t$) then we consider that $pc_{t-1}$ belongs to $dep(i, o_t)$.

**Hypothesis 1** *The value of $o_t$ computed by the program is determined by the values of objects in $dep(e(pc_{t-1}), o_t)$:*
$\forall o_t \in Output, \forall e \in Ev, e' \in Ev, e \sim_{dep(e(pc_{t-1}), o_t)} e' \Rightarrow e(o_t) = e'(o_t)$

The latter formula looks like $SecDep$ so to prove $SecDep$ it could be sufficient to prove that the security level of any member of $dep(e(pc_{t-1}), o_t)$ is dominated by $o_t$ security level. But $dep(e(pc_{t-1}), o_t)$ may contain internal objects, as $lvl$ is not defined for these objects we might be unable to check this sufficient condition. To overcome this problem we define function $lvldep$ that associates, for each evolution, a computed level with each object. If $o_t$ is an input object then $lvldep(e, o_t) = lvl(o_t)$ otherwise $lvldep(e, o_t) = max\{lvldep(e, o'_{t-1}) \mid o'_{t-1} \in dep(e(pc_{t-1}), o_t)\}$ where $max$ denotes the least upper bound in the lattice of levels.

**Theorem 1.** *A program satisfies SecDep if the computed level of an output object is always dominated by its security level:*
$\forall o \in Output, \forall e \in Ev, lvldep(e, o_t) \preceq lvl(o_t)$

**Proof:** We first prove by induction on $t$ that the program satisfies property $DepL$: $\forall o_t \in Output, \forall e \in Ev, e' \in Ev, e \sim_{depL(e,o_t)} e' \Rightarrow e(o_t) = e'(o_t)$ where $depL(e,o_t) = \{o'_{t'} \in Input \mid lvl(o'_{t'}) \preceq lvldep(e,o_t)\}$.

As there is no output initially,$DepL$ is trivially true for $t = 0$. We suppose that $DepL$ is true for output objects until time $t-1$. Let $e$ and $e'$ be two evolutions such that $e \sim_{depL(e,o_t)} e'$ where $o_t$ is an output object, we want to prove that $e(o_t) = e'(o_t)$. By definition of $lvldep$ if $o'_{t-1} \in dep(e(pc_{t-1}), o_t)$ then $lvldep(e, o'_{t-1}) \preceq lvl(o_t)$ so $\forall o''_{t''} \in Input, lvl(o''_{t''}) \preceq lvldep(e, o'_{t-1}) \Rightarrow lvl(o''_{t''}) \preceq lvl(o_t)$, hence $depL(e, o'_{t-1}) \subseteq depL(e, o_t)$. So, for all $o'_{t-1}$ in $dep(e(pc_{t-1}), o_t)$, $e \sim_{depL(e,o'_{t-1})} e'$ and by induction hypothesis $e(o'_{t-1})= e'(o'_{t-1})$. We have $e \sim_{dep(e(pc_{t-1}),o_t)} e'$ and by hypothesis 1, we can conclude that $e(o_t) = e'(o_t)$.

We now show that, under the theorem assumption, if a program satisfies $DepL$ it also satisfies $SecDep$. Let $e$ and $e'$ be two evolutions such that $e \sim_{aut(o_t)} e'$ where $o_t$ is an output object. We show that $depL(e,o_t)$ is included in $aut(o_t)$ so $e \sim_{depL(e,o_t)} e'$ and by $DepL$ we can conclude that $e(o_t) = e'(o_t)$.

Let $o'_{t'}$ be a member of $depL(e,o_t)$, by definition of $depL(e,o_t)$, $lvl(o'_{t'}) \preceq lvldep(e,o_t)$. According to the theorem assumption, $lvldep(e,o_t) \preceq lvl(o_t)$ because $o_t$ is an output object. Hence we have $lvl(o'_{t'}) \preceq lvl(o_t)$ so $o'_{t'}$ belongs to $aut(o_t)$.

As we want to use model checkers to verify the security properties, it is important to restrict the size of value domains in order to avoid state explosions during verifications. To check security, it is sufficient to prove that the previous property holds in any state of an abstracted program where object values are replaced with object computed levels. If $Ev$ is the set of evolution of the concrete program, then we note $Ev^a$ the set of evolutions of the corresponding abstract program.

**Hypothesis 2** *We suppose that the set of abstract evolution $Ev^a$ is such that the image of $Ev$ by $abs$ is included in $Ev^a$, where $abs(e)(o_t) = lvldep(e,o_t)$ if $o \neq pc$ and $abs(e)(pc_t) = e(pc_t)$.*

**Theorem 2.** *If $\forall o_t \in Output, \forall e^a \in Ev^a, e^a(o_t) \preceq lvl(o_t)$ then the concrete program guarantees $SecDep$.*

**Proof:** Let $o_t$ be an output object and $e$ be a concrete evolution in $Ev$, by Hypothesis 2 $abs(e)$ is an abstract evolution hence $abs(e)(o_t) \preceq lvl(o_t)$. By definition of $abs$, $abs(e)(o_t) = lvldep(e,o_t)$ so $lvldep(e,o_t) \preceq lvl(o_t)$ and by applying theorem 1, $SecDep$ is satisfied.

## 4 Applet certification

An application is composed of a finite number of interacting applets. Each applet contains several methods. For efficiency reasons, we want to limit the number of applets and methods analysed when a new applet is downloaded or when the security policy is modified.

We first present how we decompose the global analysis of interacting applets on a card into local verifications of a subset of the methods of one applet. Then we explain how the local verifications are implemented.

### 4.1 Global Analysis Technique

There are two related issues in order to be able to apply the approach in practice: we first have to determine what is the program we want to analyse (one method of one applet, all the methods in one applet, all the methods in all the applets on the card); then we have to identify inputs and outputs and to assign them a level.

We suppose the complete call graph of the application is given. Two kinds of methods will especially interest us because they are the basis of applet interactions: interface methods that can be called from other applets and methods that invoke external methods of other applets. We have decided to analyse subsets of the call graph that include such interaction methods. Furthermore an analysed subset only contains methods that belong to the same applet.

Let us consider for instance the Air France applet. Method $logfull$ is an interface method that calls two internal methods: $askfortransactions$ and $update$. $askfortransactions$ invokes method $gettransaction$ of Purse and credit the attribute $balance$ with the value of the transactions; $update$ invokes method $getbalance$ of RentaCar and updates the value of the $extendedbalance$ attribute. The program we are going to analyse is the set of 3 methods $logfull$, $askfortransactions$ and $update$.

For a given program, we consider as inputs results from external invocations and read attributes. We take as outputs parameters of external invocations and modified attributes. We thus associate security levels with applet attributes and with method invocations between applets. By default, we associate level $AF$ (resp. $P$ and $RC$) with all the attributes of Air France (resp. Purse and RentaCar) applet. As Air France can invoke the $getbalance$ or $debit$ methods of RentaCar, we assign the shared security level $RC + AF$ to these interactions. Similarly, as the Purse applet can invoke the $logfull$ method of Air France, we associate level $AF + P$ to this interaction. And we associate level $P + AF$ (resp. $P + RC$) to the invocation of $gettransaction$ method of the Purse applet by the Air France applet (resp. by RentaCar applet).

We propose an assume-guarantee discipline that allows to verify a set of methods locally on each applet even if the methods call methods in other applets through shared interfaces (see figure 5). For instance, method $update$ of applet Air France calls method $getbalance$ of RentaCar, we will analyse both methods separately. We check that, in the $update$ method of the Air France applet, the level of parameters of the $getbalance$ method invocation is dominated by the level of this interaction (i.e. $RC + AF$). And we assume that $RC + AF$ is the level of the result of this method invocation. When we analyse the $getbalance$ method in the RentaCar applet, we will check that the level of the result of this method is dominated by the level of the interaction and we will assume that $RC + AF$ is the level of the parameters of the *getbalance* method.

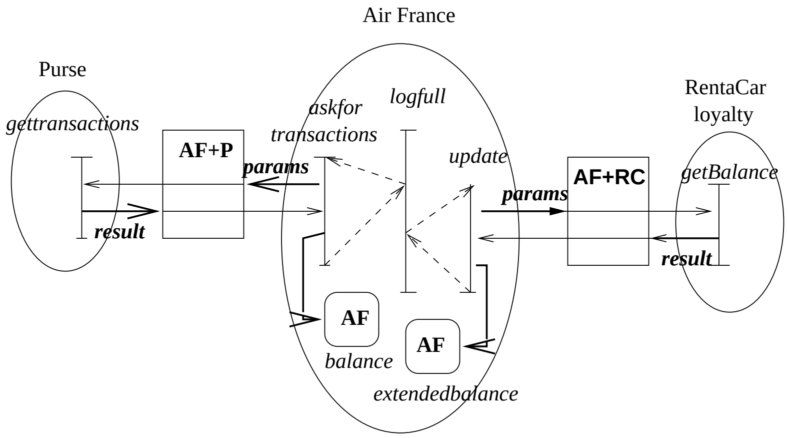


**Fig. 5.** Assume-Guarantee Verification

We adopt the same discipline for the attributes inside an applet. When an attribute is read, we assume that its level is the security level that was associated with it. When the attribute is modified, we check that the new level is dominated by the security level of this attribute. This assume-guarantee discipline inside an applet allows to verify only a subset of methods of an applet at a time (not the whole set of methods of this applet).

Thanks to this decomposition principle, it is possible to focus the analysis on the new applet that the card issuer wants to download on the cards. If the policy is unchanged there is no need to analyze again the already existing applets because levels associated with input and output objects will not change. If the security policy changes the security level associated with some input or output objects then only methods of already existing applets that use these objects should be checked again.

### 4.2 Local Analysis Technique

Our method to verify the security property on the application byte code is based on three elements:

- abstraction: we abstract all values of variables by computed levels;
- sufficient condition: we verify an invariant that is a sufficient condition of the security property;
- model checking: we verify this invariant by model checking.

We consider a set of methods at a time. Our approach uses the modularity capabilities of SMV: it first consists in building an SMV module that abstracts a method byte code for each method, then to build a `main` module that contains instances of each of these method abstraction modules and describes the interconnections between these modules. We begin by explaining how we build a module for each method, then we describe the `main` module and finally the properties. We illustrate the approach on the logfull example presented above that involves the analysis of methods $logfull$, *update* and $askfortransactions$.

**Method abstraction module.** We illustrate our technique on a simplified version of Air France *update*() method. This method directly invokes method *getbalance* of the RentaCar applet and updates the *extendedbalance* field.

```
Method void update()
   0 aload_0
   1 invokespecial 108 <Method int getbalance()>
   4 istore_1
   5 aload_0
   6 dup
   7 getfield 220 <Field int extendedbalance>
  10 iload_1
  11 iadd
  12 putfield 220 <Field int extendedbalance>
  15 return
```

*Abstraction* The *update*() byte code abstraction is modelled by an SMV [12] module. This module has got parameters (which are instantiated in the main module):

- *active* is an input of the module, it is a boolean that is true when the method is active (as we consider several methods, we have to say which one is effectively executing);
- *context* is a boolean representing the context level of the caller;
- *param* is an array containing the levels of the parameters of the method;
- $field$ is an array containing the levels of the attributes used by the method (only one here *extendedbalance*);
- *method* is an array containing the levels of the results of the external methods invoked by the *update* method (only one here *getbalance*).

The module also involves the following variables:

- *pc*: program counter;
- *lpc*: the level of the program counter, the context level of the method;
- *mem*[*i*]: an array modelling the memory locations;
- *stck*[*i*]: an array modelling the operand stack;
- $sP$: stack pointer;

– *ByteCode*: the name of the current instruction.

The values of the variables are abstracted into levels. Levels are defined in a module called Levels in such a way that a level is represented by a boolean. Hence the types of abstracted variables are boolean or array of boolean. We do not abstract the value of the program counter that gives the sequencing of instructions, we keep unchanged the value of the stack pointer that gives the index of the first empty slot.

```
module update(active, context, param, field, method){
 L: levels;
 pc : -1..9;
 lpc: boolean;
 mem : array 0..1 of boolean;
 stck : array 0..1 of boolean;
 sP : -1..1;
 ByteCode : {invoke_108, load_0, return, nop, store_1, dup,
            load_1, getfield_220,op, putfield_220};
```

The byte code execution starts at program location 0. Initially, the stack is empty, the level of the method parameter is stored in memory location 0. The level *lpc* is initialized to the context level of the caller *context*.

```
init(pc):= 0; init(sP):= 1; init(mem[0]):= param[0];
for (i=0; i< 2; i=i+1) {init(stck[i]) := L.public; }
init(lpc) := context;
```

The control loop defines the value of the program counter and of the current instruction. It is an almost direct translation of the Java byte code. When *pc* is equal to -1 then the execution is finished and the current instruction is nop that does nothing. As in [6], each instruction we consider models various instructions of the Java byte code. For instance, as we do not care about the type of memory and stack locations, instruction *load_0* represents Java instructions (aload_i, iload_i, lload_i,...). Similarly, the *op* instruction models all the binary operations as (iadd, ladd, iand, ior, ...).

Although methods in the Electronic Purse case study tend to be rather short, the value domain of the *pc* variable is generally the larger domain in the SMV model we generate. As a large value domain can be a source of state explosion for model checkers, it is interesting to reduce it. We consider that the next value of *pc* is $pc + 1$ (except of course in the case of conditional or jump instructions). In this example, the domain of *pc* is -1..9 that is smaller that the original domain 0..15.

```
if (active) {
  (next(pc), ByteCode) :=
   switch(pc) {
    -1: (-1, nop);
```

```
    0 : (pc+1, load_0 );
    1 : (pc+1, invoke_108 );
    2 : (pc+1, store_1 );
    3 : (pc+1, load_0 );
    4 : (pc+1, dup );
    5 : (pc+1, getfield_220 );
    6 : (pc+1, load_1 );
    7 : (pc+1, op );
    8 : (pc+1, putfield_220 );
    9 : (-1, return);
  };}
else {next(pc) := pc; next(ByteCode) := nop;}
```

The following section of the SMV model describes the effect of the instructions on the variables. The instructions compute levels for each variable. The *load* instruction pushes the level of a memory location on the stack, the *store* instruction pops the top of the stack and stores the least upper bound of this level and $lpc$ in a memory location. The least upper bound of levels $l1$ and $l2$ is modelled by the disjunction of two levels $l1 \vee l2$. The *dup* instruction duplicates on the stack the top of the stack. The *op* instruction computes the least upper bound of the levels of the two first locations of the stack. The *invoke* instruction pops from the stack the parameter and pushes onto the stack the result of this method invocation. Instruction $getfield$ pushes on the top of the stack the level of attribute $extendedbalance$. And, finally, instruction $putfield$ pops from the stack the level of attribute $extendedbalance$.

```
   switch(ByteCode) {
    nop :;
    load_0 : {next(stck[sP]):= mem[0];next(sP):=sP-1;}
    load_1 : {next(stck[sP]):= mem[1];next(sP):=sP-1;}
    store_1 : {next(mem[1]):=(stck[sP+1]|lpc) ;next(sP):=sP+1;}
    dup : {next(stck[sP]):= stck[sP+1]; next(sP):=sP-1;}
    op : {next(stck[sP+2]):=(stck[sP+1]|stck[sP+2]);
          next(sP):=sP+1};
    invoke_108 : {next(stck[sP]):=method[0];next(sP):= sP+1;}
    getfield_220 : {next(stck[sP+1]):=field[0];}
    putfield_220 : {next(sP):=sP+2;}
    return : ;
   }
```

*"Conditional" instructions.* No conditional instruction such as `ifne`, `tableswitch` ... occurs in the example presented above. The behaviour of these instructions is to jump at various locations of the program depending on the value of the top of the stack. As this value is replaced by a level, the abstract program cannot decide precisely which is the new value of the program counter. So an SMV non-deterministic assignment is used to state that there are several possible values for $pc$ (generally $pc + 1$ or the target location of the conditional instruction).

It is well known that conditional instructions introduce a dependency on the condition. This dependency is taken into account by means of the *lpc* variable. When a conditional instruction is encountered, *lpc* is modified and takes as new value the least upper bound of its current level and of the condition level. As each modified variable depends on *lpc*, we keep trace of the implicit dependency between variables modified in the scope of a conditional instruction and the condition.

**Main module.** We have an SMV module for the *update* method. We can build in a similar way a module for the *gettransactions* method and for the $logfull$ method. We also have a module Levels. The SMV `main` module is composed of two parts: the first part manages the connections between methods (definition of the active method, parameter passing); the second one assigns levels to attributes and interactions.

The `main` module includes an instance of each of the method abstraction modules and one instance of the level module. To define the active method and describe the activity transfer, we need a global variable *active* from which the activity parameter of the different methods is defined.

```
module main(){
active : {logfull, askfortransactions, update};
L: levels;
m_logfull: logfull((active=logfull));
m_aft: askfortransactions((active=askfortransactions),context_aft,
                          param_aft, field_aft,method_aft);
m_update: update((active=update),context_ud, param_ud,
                 field_ud,method_ud);
```

The `main` module contains a set of equations that define the value of *active* according to the call graph. Initially, the active method is $logfull$. When $askfortransactions$ is invoked (`m_logfull.ByteCode=invoke_235`), method $askfortransactions$ becomes active. When it terminates (`m_aft.pc=-1`), logfull becomes active again. When *update* is invoked (`m_logfull.ByteCode=invoke_192`) method *update* becomes active. When this method terminates (`m_update.pc=-1`), $logfull$ is active until the end.

When $logfull$ invokes $askfortransactions$ or *update*, it involves parameter passing between methods. In the example, there is only one parameter in each case, so it is sufficient to copy the top of logfull stack into the parameter array of the invoked method. We also transfer the context level of the caller to the invoked method by copying *lpc* in the *context* parameter.

```
if(m_logfull.ByteCode = invoke_235) {
          next(param_aft[0]) := m_logfull.stck[m_logfull.sP+1];
          next(context_aft) := m_logfull.lpc; }

if(m_logfull.ByteCode = invoke_192) {
```

```
next(param_ud[0]) := m_logfull.stck[m_logfull.sP+1];
next(context_ud) := m_logfull.lpc; }
```

Remark: the methods in the example do not have result, but in the general case, we would also have to transfer the result (i.e. the top of the stack) from the invoked method to the caller.

It now remains to assign levels to attributes and interactions. In the example, we have two attributes with level $AF$: *balance* (146) which is a parameter of $askfortransactions$, and *extendedbalance* (220) which is a parameter of *update*; and two interactions: *getbalance* (108) between Air France and RentaCar (so its level is $AF+RC$), parameter of $askfortransactions$, and *gettransactions* (179) between Air France and Purse (so its level is $AF+P$), parameter of *update*. Remark: in our boolean encoding of levels, $l1+l2$ is expressed by $l1\&l2$.

```
field_aft[0] := L.AF;     method_aft[0] := L.AF & L.P;
field_ud[0] := L.AF;      method_ud[0] := L.AF & L.RC;
```

**Invariant.** We explained above how to compute a level for each variable. We also explained what security level we assigned to attributes and interactions. The invariant we verify is the sufficient condition we previously described: the computed level of each output is always dominated by its security level.

For the *update* method we should check two properties : one to verifiy that the interaction between *update* and *getbalance* is correct and the other one to check that *update* correctly uses attribute *extendedbalance*. Property *Smethod_108* means that, whenever the current instruction is the invocation of method *getbalance*, then the level of the transmitted parameters (the least upper bound of *lpc* and the top of the stack) is dominated by the level of the interaction $AF+RC$. In our boolean encoding of levels, $l1$ is dominated by $l2$ if `L.l1` implies `L.l2`. Property *Sfield_220* means that, whenever the current instruction is the modification of field *extendedbalance*, then the level of the new value (the least upper bound of *lpc* and the top of the stack) is dominated by the level of the attribute $AF$.

```
Smethod_108 :
  assert G (m_update.ByteCode=invoke_108 ->
           ((m_update.stck[sP+1]|m_update.lpc) -> L.AF & L.RC));
Sfield_220  :
  assert G (m_aft.ByteCode=putfield_220 ->
           ((m_aft.stck[sP+1]|m_aft.lpc) -> L.AF));
```

For the initial method (here $logfull$), we also have to verify the *Sresult* property which means that whenever the method is finished the level of the return value (the top of the stack) is dominated by the level of the interaction $AF+P$. As $logfull$ does not return any value, there is no need to verify this property.

```
Sresult     :
  assert G (m_logfull.ByteCode=return ->
          ((m_logfull.stck[sP+1]|m_logfull.lpc) -> L.AF & L.P));
```

### 4.3 Analysis.

Once we have the abstract model and the invariant, we model check the invariant properties on the model using SMV [12]. If the property does not hold the model checker produces a counter-model that represents an execution of the byte code leading to a state where the property is violated.

A security problem will be detected when checking property $Smethod\_108$ of method *update*. Indeed, the $logfull$ interaction between purse and Air France has $AF+P$ level. The $getbalance$ channel has $AF+RC$ level and we detect that the invocation of the $getbalance$ method depends on the invocation of the $logfull$ method. There is thus an illicit dependency from a variable of level $AF+P$ to an object of level $AF+RC$.

(cf figure 6).

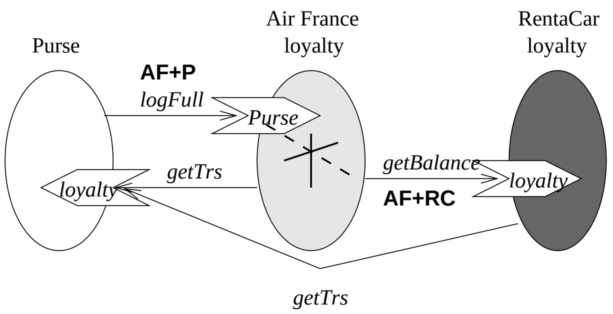


**Fig. 6.** Illicit interaction detection

To check all the possible interactions on the example of the electronic purse, we have to do 20 analyses such as the one we presented in this paper. These analyses involve about 100 methods and 60 properties to be verified. The complete (non simplified) version of the logfull interaction contains 5 methods and 8 properties, the verification of each property takes 1s on an Ultra 10 SUN station running Solaris 5.6 with 256 Mbytes of RAM. For the other interactions, 30 properties were checked individually in less than 5s, each of the 20 other properties were verified in less than one minute and the remaining properties are checked individually in less than 3 minutes. We observed that the verification time depends on the number of byte-code lines of methods but SMV was always able to verify the properties in a few minutes. Hence, we think that our technique could be applied to real-world applications because the electronic purse case-study is, by now, a "big" application with respect to smart-card memory size limitations.

## 5 Related work

A lot of work has been going on about the analysis of security properties of Java byte code. The major part of this work is concerned with properties verified by SUN byte code verifier like correct typing, no stack overflow, etc. Among this work, two kinds of approaches can be distinguished depending on the technique used for the verification. Most of the approaches are based on static analysis techniques, particularly type systems [2, 16]. One approach has used model checking (with SMV) to specify the correct typing of Java byte code [14]. We also based our approach on model-checking tools because they tend to be more generic and expressive than type-checking algorithms. This allowed us to obtain results faster because we did not have to implement a particular type-checking algorithm. This should also enable us to perform experiments with other security policies and properties.

Recently, several researchers investigated the static analysis of information flow properties quite similar to our secure dependency property but, to our knowledge, none of them applied their work on Java byte-code. Girard et alter [7] defined a type system to check that a programme written in a subset of the C language does not transfer High level information in Low level variables. In [3], the typing relation was related to the secure dependency property. Volpano and Smith [15] proposed a type system with a similar objective for a language that includes threads. They relate their typing relation to the non-interference property. The secure dependency property was compared with non-interference in [1]. Myers and Liskov [13] propose a technique to analyze information flows of imperative programs annotated with labels that aggregate the sensitivity levels associated with various information providers. One of the interesting feature is the declassify operation that allows providers to modify labels. They propose a linear algorithm to verify that labels satisfy all the constraints.

A few pieces of work deal with other kind of security properties. In [8] the authors propose an automatic method for verifying that an implementation using local security checks satisfies a global security property. However, their approach is limited to control flow properties such as Java Virtual Machine stack inspection. The work described in [10] focusses on integrity property by controlling exclusively write operations to locations of references of sensitive objects such as files or network connections.

## 6 Conclusion and future work

In this paper, we have presented an approach for the certification of applets that are to be loaded on a Javacard. The security checks we propose are complementary to the security functions already present on the card. The applet firewall controls the interaction between two applets, while our analysis has a more global view and is able to detect illicit information flow between several applets.

As stated in section 4.3, the complete analysis of the application would concern 20 sets of methods, involving globally about 100 methods. Consequently a lot of SMV models need to be built. Automatization of the production of models is thus mandatory for the approach to be practicable. Such an automatization is relatively straightforward provided that preliminary treatments are made to prepare the model construction, such as construction of the call graph, method name resolution, etc. However, a complete automatization is hardly possible: an interaction with the user will be needed for the definition of security policy and level assignments.

Another interesting issue is the analysis of results. When SMV produces a counter-example for a security property, we have to study how to interpret this counter-example as an execution of the concrete byte code program at the application level.

## References


1. P. Bieber and F. Cuppens. A Logical View of Secure Dependencies. *Journal of Computer Security*, 1(1):99–129, 1992.
2. Stephen N. Freund and John C. Mitchell. A type system for object initialization in the java byte code language. In *Proceedings of OOPSLA 98*, 1998.
3. P. Girard. *Formalisation et mise en oeuvre d'une analyse statique de code en vue de la verification d'applications securisees*. PhD thesis, ENSAE, 1996.
4. Pierre Girard. Which security policy for multiapplication smart cards? In *USENIX workshop on smartcard technology*, 1999.
5. J. Goguen and J. Meseguer. Unwinding and Inference Control. In *IEEE Symposium on Security and Privacy*, Oakland, 1984.
6. Pieter H. Hartel, Michael J. Butler, and Moshe Levy. The operational semantics of a java secure processor. Technical Report DSSE-TR-98-1, Declarative systems and Software Engineering group, University of Southampton,Highfield, Southampton SO17 1BJ, UK, 1998.
7. C. O'Halloran J. Cazin, P. Girard and C. T. Sennett. Formal Validation of Software for Secure Systems. In *Anglo-french workshop on formal methods, modelling and simulation for system engineering*, 1995.
8. T. Jensen, D. Le Métayer, and T. Thorn. Verification of control flow based security policies. In *Proceedings of the 20th IEEE Security and Privacy Symposium*, 1999.
9. Paul Kocher, Joshua Jaff, and Benjamin Jun. Differential power analysis: Leaking secrets. In *Advances in Cryptology – CRYPTO'99 Proceedings*. Springer-Verlag, 1999.
10. X. Leroy and F. Rouaix. Security properties of typed applets. In *Proceedings of POPL*, 1998.
11. T. Lindholm and F. Yellin. *The Java Virtual Machine Specification*. Addison Wesley, 1997.
12. K.L. McMillan. *The SMV language*. Cadence Berkeley Labs, 1999.
13. A.C. Myers and B. Liskov. A decentralized model for information flow control. In *Proceedings of the 16th ACM symposium on operating systems principles*, 1997.
14. Joachim Posegga and Harald Vogt. Offline verification for java byte code using a model checker. In *Proceedings of ESORICS*, number 1485 in LNCS. Springer, 1998.
15. G. Smith and D.M. Volpano. Secure information flow in a multi-threaded imperative language. In *Proceedings of POPL*, 1998.

16. Raymie Stata and Martin Abadi. A type system for java bytecode subroutines. In *Proc. 25th Symposium on Principles of Programming Languages*, 1998.
17. Sun Microsystems. *Java Card 2.1 Realtime Environment (JCRE) Specification*, February 1999.
18. Visa. *Open Platform, Card Specification*, April 1999. Version 2.0.