Operational Semantics of the Java Card Virtual Machine

Igor A. Siveroni

Imperial College London. U.K.

Abstract

We present the operational semantics of Carmel, a language that models the Java Card Virtual Machine Language. We define a small-step relation between program configurations, including rules for exception handling, array objects and subroutines. We also include the basic structures needed to model object ownership and the Java Card firewall.

Authors:	Igor Siveroni
Date:	July, 2003
Number:	SECSAFE-ICSTM-011
Version:	1.0
Classification:	Public
Status:	To appear in JLAP.

Key words: Java Card, Semantics
PACS:

1 Introduction

Java Card was first introduced in 1996 as a new technology that enables Java applets to be loaded and executed in new generation smart cards and other devices with limited memory. Java Card is a simplified version of Java e.g. it does not include threads, floating point numbers and multi-dimensional arrays, and has a relatively small size API. Java Card applications are also small due to the memory limitations of the Smart Card processor. All these

Email address: `siveroni@doc.ic.ac.uk` (Igor A. Siveroni).

factors combined with the security-critical nature of its applications have made it an ideal case-study for the application of formal methods.

The Java Card platform [8] is defined by three (natural language) specifications: The *Java Card Virtual Machine* (JCVM) *Specification* [29], the *Java Card Application Programming Interface* (API) [27], and the *Java Card Runtime Environment* (JCRE) *Specification* [28]. The specifications of other important components of the Java Card system - such as the bytecode verifier, applet installer and CAP converter - can either be found scattered in the specification of the main components or deduced from their counterparts in the Java system (e.g. bytecode verification [18]). Unfortunately, the official Java Card specifications are sometimes ambiguous, contradictory and lack mathematical rigour.

The main goal of this work is to present a readable and complete specification of the JCVM that can be used as a component of the specification of the whole Java Card system which takes into account the JCRE and API. We define the operational semantics of Carmel [20], a language that models the Java Card Virtual Machine Language. We define a small-step relation between program configurations - including rules for exception handling and subroutines - in a style similar to the one used by Bertelsen [6] and use a type system based on the work of Freund and Mitchell [11].

The distinctive features of this article that are due to the fact that we deal with Java Card, and not Java, can be summarised as follows:

- The JCVM provides an object access mechanism, the JC firewall, based on object ownership. We define the basic structures needed to model object ownership and the Java Card firewall, and improve on the work of Éluard et al. [10] by providing a complete specification of the firewall which accounts for transient data.
- A Java Card is governed by the JCRE. Most of the interaction between a running applet and the JCRE is done via the invocation of API methods. A complete specification of the semantics of the JCRE and API belongs to another document [25]. In this paper we provide a framework for such specification by modelling the JCRE as part of the machine state and defining the basic interactions between the JCVM and the JCRE.

Several formal specifications of the JCVM have been developed in recent years [15]. Research in the program verification area have produced reasonably complete specifications [5] developed with the aid of automated proof assistants [9,22]. These specifications, though complete, are large and difficult to read. Other specifications focus in a reduced subset or a particular aspect of the machine e.g firewall [10]. Furthermore, these specifications do not take into account the other components of the Java Card system i.e. JCRE and

API. The main contribution of this article in relation to existing semantic descriptions of the JCVM is that it presents in a formal and clear way - avoiding notation specific to, for example, proof assistants - the most distinctive features of the JCVM and its instruction set and, as pointed out above, it models the basic interactions between the JVML and JCRE. In particular, special detail have been given to the specification of the firewall.

We start by presenting the syntax of Carmel and the basic program structures used in this paper (section 2). We then define the runtime structures and values used by the JCVM (section 3) and formalise the mechanism for throwing exceptions (section 4) and the firewall rules (section 5). Next, we explain in detail the configurations and rules that define the operational semantics of Carmel (section 6) and finish by summarising related work (section 7) and presenting conclusions and future work (section 8).

1.1 Notation

A Domain is defined as a set equipped with functions. For convenience, in some cases we define new domains as records. A record defined by

$$\mathsf{Dom} = (f_1 : \mathsf{Dom}_1) \times ... \times (f_n : \mathsf{Dom}_n)$$

is equivalent to the domain Dom equipped with the interface $f_i \in \mathsf{Dom} \to \mathsf{Dom}_i$, for $0 < i \leq n$. Given $e \in \mathsf{Dom}$, access to element f_i is written as $e.f_i$ and element update is performed by $e[f_i \mapsto v]$, where $v \in \mathsf{Dom}_i$. To simplify notation, we will usually write $o.f(args)$ instead of $f(o, args)$.

We write $[\mathsf{X}]$ to denote the set of finite arrays of elements of domain X. The length of an array can be known using the function $\mathit{length} : [X] \to \mathsf{int}$. The indices of an array x range from 0 to $x.\mathit{length} - 1$. Elements of the array x can be accessed using the standard $x[i]$ notation. The i-th element of an array is updated by $x[i \to v]$. We use X^* to denote the set of sequences of elements of domain X, defined by $\mathsf{X}^* ::= \epsilon \mid x{::}\mathsf{X}^*$, where $x \in \mathsf{X}$ and ϵ is the empty sequence.

2 Carmel Syntax

The Java Card technology is a subset of the Java platform. Features such as characters and strings, threads, cloning, garbage collection, the security manager and dynamic class loading are not supported. Only basic object-oriented capabilities and features that are well-suited for writing programs for

$$
\begin{aligned}
\mathsf{Program} &= \mathcal{P}(\mathsf{Package})\\
\mathsf{Package} &= (aid : \mathsf{AID}) \times (classes : \mathcal{P}(\mathsf{Class})) \times (interfaces : \mathcal{P}(\mathsf{Interface}))\\
\mathsf{Class} &= (superClass : \mathsf{Class}) \times (implements : \mathcal{P}(\mathsf{Interface})) \times\\
&\quad (fields : \mathcal{P}(\mathsf{Field})) \times (methods : \mathcal{P}(\mathsf{Method}))\\
\mathsf{Interface} &= (superInterfaces : \mathsf{Interface}) \times (fields : \mathcal{P}(\mathsf{Field})) \times\\
&\quad (methods : \mathcal{P}(\mathsf{Method}))\\
\mathsf{Field} &= (type : \mathsf{Type}) \times (id : \mathsf{ID}) \times (isStatic : \mathsf{boolean})\\
\mathsf{Method} &= (type : \mathsf{MethodType}) \times (id : \mathsf{ID}) \times (isStatic : \mathsf{boolean}) \times\\
&\quad (classI : \mathsf{Class} \cup \mathsf{Interface}) \times (eHandlers : [\mathsf{EHandler}]) \times\\
&\quad (instructionAt : (\mathsf{Address} \rightarrow \mathsf{Instruction}))
\end{aligned}
$$

$$
aid \in \mathsf{AID} \quad cl \in \mathsf{Class} \quad iface \in \mathsf{Interface} \quad f \in \mathsf{Field} \quad m \in \mathsf{Method}
$$

$$
\begin{aligned}
\tau \in \mathsf{Type} &::= \mathsf{ReferenceType} \mid \mathsf{PrimitiveType}\\
\mathsf{PrimitiveType} &::= \texttt{boolean} \mid \texttt{byte} \mid \texttt{short} \mid \texttt{int}\\
\tau_{\mathrm{r}} \in \mathsf{ReferenceType} &::= \mathsf{ArrayType} \mid \mathsf{SimpleRef}\\
\mathsf{SimpleRef} &::= \mathsf{Class} \mid \mathsf{Interface}\\
\mathsf{ComponentType} &::= \mathsf{SimpleRef} \mid \mathsf{PrimType}\\
\mathsf{ArrayType} &::= \mathsf{ComponentType}\texttt{[ ]}\\
\tau_{\texttt{rm}} \in \mathsf{ResultType} &::= \mathsf{Type} \mid \texttt{void}\\
\mathsf{MethodType} &::= \mathsf{Type}^{*} \rightarrow \mathsf{ResultType}
\end{aligned}
$$

Fig. 1. Carmel Syntax

smart cards are supported. The JCVML is an optimised subset of the Java bytecode instruction set which, besides not including the instructions related to the unsupported Java features, does not support large primitive data types (`long`, `float` and `float`) and multi-dimensional arrays. Java Card programs, unlike Java, are not represented as `class` files. Instead, the compiled class files of a package are converted to a Java Card CAP (converted applet) file which is then ready to be installed on a card. The CAP file format is optimised for a small footprint by using compact data structures and limited indirection e.g. names are dropped, and tokens and offsets are used instead.

Carmel is a language that models the JCVML. It captures all the JCVML features while reducing the number of instructions (by grouping similar JCVML instructions into a single Carmel instruction) and abstracting away certain complexities of the CAP file format such as the use of tokens, offsets and the constant pool. The syntax of a Carmel program is shown in Figures 1 and 2. We have only included the elements of the program structure that are relevant for the presentation of this paper. The rest of the syntax e.g. parts needed for bytecode verification, can be found in [20]. A Carmel program is made of a set of packages - the packages that have been loaded to the card. A package contains classes and interfaces, and is uniquely identified by an AID (application identifier). A method structure m contains information about the method's ID - used for virtual method invocation (section 6.3.3) - its type and the class or interface ($m.\mathit{classI}$) where it was declared. We write $|m|$ to denote the method's arity. A non-abstract method has additional information concerning its body. The sequence of instructions that make up a method's body is accessed via elements of the domain $\mathsf{Address}$. We use $m.\mathit{instructionAt}(\mathit{addr})$ to access the instruction at address addr in method m. The first instruction of a method is located at address 0 and the instruction following the instruction at address addr can be found at $\mathit{addr} + 1$. A field is uniquely identified by a field ID. A field structure f contains information about the field's type and ID.

Classes and interfaces hold information about direct fields and methods, inherited elements have to be looked up in the superclass. A set of auxiliary functions is defined to help extract information from the class hierarchy. The functions $\mathit{superClasses}(\mathit{cl})$ and $\mathit{superInterfaces}^*(\mathit{iface})$ return the set of all the superclasses and superinterfaces of class cl and interface iface, respectively. The function $\mathit{implements}^*(\mathit{cl})$ returns the set of interfaces implemented by cl while $\mathit{instanceFields}(\mathit{cl})$ returns the set of non-static fields (direct and inherited) of class cl.

The Carmel instruction set is divided into six groups, each dealing with a different aspect of the language: imperative core (C), objects (O), method invocation (M), arrays (A), exceptions (E) and subroutines (S). The instruction set [1] is defined in Figure 2. As JCVM instructions, most Carmel instructions are typed and operate on values that have an expected runtime type - the operand type. One of the main differences between the JCVM language and Carmel resides in the way the operand type information is encoded. In Carmel, instead of having a specialised instruction for each operand type (e.g. `astore`, `istore`), instructions take the operand type as an argument (e.g. `store r 2`). This simplification reduces the number of instructions. Constants c are introduced by the `push` instruction and can be either integers or the special

[1] For this presentation, we have removed the `keyswitch`, `indexswitch` instructions, and those that take the `this` keyword as argument.

$$
\begin{array}{rcl}
I \in \mathsf{Instruction} & ::= & \texttt{nop} \mid \texttt{push}\ t\ c \mid \texttt{pop}\ n \mid \texttt{swap}\ n_1\ n_2 \mid \texttt{dup}\ n\ d \quad \text{(C)} \\
 & & \texttt{numop}\ t\ op\ t \mid \texttt{goto}\ addr \mid \texttt{if}\ t\ cmp\ \texttt{goto}\ addr \\
 & & \texttt{load}\ t\ i \mid \texttt{store}\ t\ i \mid \texttt{inc}\ t\ i\ c \\
 & & \texttt{new} \ \mid \texttt{checkcast}\ \tau \mid \texttt{instanceof}\ \tau \quad \text{(O)} \\
 & & \texttt{getstatic}\ f \mid \texttt{putstatic}\ f \mid \texttt{getfield}\ f \mid \texttt{putfield}\ f \\
 & & \texttt{invokevirtual}\ m \mid \texttt{invokedefinite}\ m \quad \text{(M)} \\
 & & \texttt{invokeinterface}\ m \mid \texttt{return}\ t \mid \texttt{return} \\
 & & \texttt{arraylength} \mid \texttt{arrayload}\ t \mid \texttt{arraystore}\ t \quad \text{(A)} \\
 & & \texttt{throw} \mid \texttt{jsr}\ addr \mid \texttt{ret}\ i \quad \text{(E) \& (S)} \\
t \in \mathsf{OpType} & = & \{\texttt{b}, \texttt{s}, \texttt{i}, \texttt{r}\} \\
op \in \mathsf{NumericOp} & = & \mathsf{UnaryNumOp} \cup \mathsf{BinaryNumOp} \\
\mathsf{UnaryNumOp} & = & \{\texttt{neg,to}\} \\
\mathsf{BinaryNumOp} & = & \{\texttt{add,sub,mul,div,rem,and,or,xor,shl,shr,ushr}\} \\
cmp \in \mathsf{ComparisonOp} & = & \{\texttt{eq,ge,gt,le,lt,ne}\ \}
\end{array}
$$

Fig. 2. Carmel Instruction Set

constant `null`. Number of words are denoted by n and d, local variable indices by i, while *addr* is used to denote an instruction address.

All references to classes, interfaces, methods and fields have been resolved to their respective representation structure and can be accessed directly from wherever they are referenced. Thus, the field structure f can be extracted directly from the `getstatic` f instruction. However, in order to make the presentation more readable, names - concrete syntax - are introduced when necessary. Thus, when we write `Object` we are actually referring to the structure associated to the `java.lang.Object`[2] class in the program.

[2] By default we will assume that classes without prefix belong to the `java.lang` package.

3 Runtime Structures and Values

3.1 Objects and the Heap

Objects are stored in the heap. The heap is defined as a mapping from locations to objects, which can be class instances or array objects. Java Card objects are defined in Figure 3. A class instance $\langle \tau, fv, own, ep \rangle^{\mathrm{c}}$ contains information about the object's class (τ), its instance fields runtime values (fv), owner (own) and JCRE entry point type (ep). An array object $\langle \tau, n, vs, own, g, t \rangle^{\mathrm{a}}$ contains the type of its elements (τ) and information about the array's size (n), values (vs), owner (own), global access (g) and memory persistency (tr). Object ownership, JCRE entry points, and global and transient (i.e. not persistent) arrays are special features of Java Card. The domains associated with these features are:

$$own \in \mathsf{Owner} = (context : \mathsf{Context}) \times (aid : \mathsf{AID} \cup \{JCRE\})$$

$$ctxt \in \mathsf{Context} = \mathsf{AID} \cup \{JCRE\} \qquad ep \in \mathsf{EntryPoint} = \{\mathtt{perm}, \mathtt{temp}, \mathtt{no}\}$$

$$tr \in \mathsf{Transient} = \{\mathtt{CLEAR\_ON\_RESET}, \mathtt{CLEAR\_ON\_DESELECT}, \mathtt{NOT\_TRANSIENT}\}$$

An object's Owner is made of two components: the AID of the owning applet and the AID of the package where the owning applet is defined, called the object's (group) *Context*. In general, the Java Card firewall only allows access among objects that belong to the same context. Special JCRE-owned objects i.e. JCRE entry points and global arrays, can be accessed from any context. Object ownership and access checks are discussed in detail in the section dedicated to the Firewall (section 5).

For security reasons, sometimes it is not desirable to preserve data stored in objects across sessions. An object is transient if the data stored in its fields is cleared whenever the card is reset or the applet is deselected. In the current specification of the JCVM [29] only arrays with primitive components or references to `Object` can be designated as transient objects. There are two types of transient objects, namely, `CLEAR_ON_RESET` and `CLEAR_ON_DESELECT`. The interaction between transient objects and the applet firewall is explained in section 5.

New objects are created and allocated by the JCVM according to the *newObject* and *newArray* relations. A new object $\langle \tau, fv, own, ep \rangle^{\mathrm{c}}$ - or a new array $\langle \tau, n, vs, own, g, tr \rangle^{\mathrm{a}}$ - is correctly allocated in heap H - generating new heap H' - and assigned location loc if the following conditions are satisfied :

$$
\begin{aligned}
H \in \mathsf{Heap} &= \mathsf{Location} \rightarrow \mathsf{Object} \\
\mathsf{Object} &= \mathsf{ClassInst} \cup \mathsf{ArrayObject} \\
o \in \mathsf{ClassInst} &= (class : \mathsf{Class}) \times (fieldValue : \mathsf{FieldID} \rightarrow \mathsf{Value}) \times \\
&\quad (entryPoint : \mathsf{EntryPoint}) \times (owner : \mathsf{Owner}) \\
a \in \mathsf{ArrayObject} &= (eleType : \mathsf{Type}) \times (length : \mathsf{integer}) \times \\
&\quad (values : [\mathsf{Value}]) \times (owner : \mathsf{Owner}) \times \\
&\quad (global : \mathsf{boolean}) \times (transient : \mathsf{Transient}) \\
\langle \tau, fv, own, ep\rangle^{\mathrm{c}} \in \mathsf{ClassInst} &\quad \langle \tau, n, vs, own, g, tr\rangle^{\mathrm{a}} \in \mathsf{ArrayObject}
\end{aligned}
$$

Fig. 3. Java Card Objects and the Heap

$$
\begin{aligned}
&\langle H', loc\rangle \in newObject(own, \tau, ep, H) \Leftrightarrow && (1) \\
&\quad loc \notin dom(H) \wedge H' = H[loc \mapsto o] \wedge o = \langle \tau, fv, own, ep\rangle^{\mathrm{c}} \wedge \\
&\quad fv = \{(f.id, def(f.type)) \mid f \in instanceFields(\tau)\} \\
&\langle H', loc\rangle \in newArray(own, \tau, n, g, tr, H) \Leftrightarrow && (2) \\
&\quad loc \notin dom(H) \wedge H' = H[loc \mapsto a] \wedge a = \langle \tau, n, vs, own, g, tr\rangle^{\mathrm{a}} \wedge \\
&\quad vs.length = n \wedge \forall i.\ 0 \leq i < n.\ vs[i] = def(\tau)
\end{aligned}
$$

where *instanceFields* is the set of non-static fields of a class, including the fields defined in its superclasses, and $def(\tau)$ returns the default value for a particular type τ (0 for numeric types and `null` for reference types). For this presentation we assume that fields and array values are initialised to the default values indicated by their types.

Note that *newObject* and *newArray* are non-deterministic. No constraints are placed on *loc* being unique, we only make sure that *loc* is a fresh location i.e. $loc \notin dom(H)$.

3.1.1 Heap Update

Class instances and array objects are identified by their location in the heap. Similarly an element location, defined by

$$eloc \in \mathsf{ElementLocation} = (\mathsf{Location} \times \mathsf{FieldID}) + (\mathsf{Location} \times \mathsf{ShortValue}),$$

is used to identify instance fields and array entries of class instances and array objects, respectively. Thus, we define a more succinct definition of field and array update with the *updateEle* function:

$$
\begin{aligned}
&\mathit{updateEle} : \mathsf{Heap} \times \mathsf{ElementLocation} \times \mathsf{Value} \rightarrow \mathsf{Heap}\\
&\qquad \mathit{updateEle}(H, (\mathit{loc}, \mathit{id}), v) = H[\mathit{loc} \mapsto o'],\\
&\qquad\qquad o' = H(\mathit{loc})[\mathit{fieldValues} \mapsto H(\mathit{loc}).\mathit{fieldValues}[\mathit{id} \mapsto v]]\\
&\qquad \mathit{updateEle}(H, (\mathit{loc}, i), v) = H[\mathit{loc} \mapsto a'],\\
&\qquad\qquad a' = H(\mathit{loc})[\mathit{values} \mapsto H(\mathit{loc}).\mathit{values}[i \mapsto v]]
\end{aligned} \tag{3}
$$

3.2 Values and Runtime Data Types

The JCVM supports two kind of values: primitive values and reference values. Primitive values consist on numeric values (`byte`, `short` and `int`) and return addresses. Reference values consist on references to class instances and one-dimensional array objects (locations, section 3.1) and the special constant `null`. Java Card virtual machine values can be:

$$
\begin{aligned}
v \in \mathsf{Value} &= \mathsf{PrimitiveValue} \ \cup \ \mathsf{ReferenceValue}\\
\mathsf{PrimitiveValue} &= \mathsf{NumericValue} \ \cup \ \mathsf{ReturnAddressValue}\\
\mathsf{NumericValue} &= \mathsf{ByteValue} \ \cup \ \mathsf{ShortValue} \ \cup \ \mathsf{IntegerValue}\\
\mathsf{ReferenceValue} &= \mathsf{Location} \cup \{\mathtt{null}\} \quad \mathsf{ReturnAddressValue} = \mathsf{Address}
\end{aligned}
$$

ReturnAddressValue values are elements of the Address domain (instruction addresses) and are included for the execution of subroutine instructions (section 6.3.4). Storage is managed by the JCVM in terms of an abstract storage unit called word. A word is large enough to hold reference and numeric values, with the exception of values of type `int`. Two words are large enough to hold a value of type `int`. The function *nbWords* returns the number of words of a value or a sequence of values.

The Java types `byte` and `short` - and in some implementations, the `int` type - are supported by the JCVM. Values of the Java type `boolean` are implemented as byte values, where 1 represents `true` and 0 `false`.

Operations between numeric values are performed using the *applyUnary* and *applyBinary* functions:

$$
\begin{aligned}
\mathit{applyUnary} &: \mathsf{UnaryNumOp} \times \mathsf{Value} \times \mathsf{OpType} \times \mathsf{OpType} \rightarrow \mathsf{Value}\\
\mathit{applyBinary} &: \mathsf{BinaryNumOp} \times \mathsf{Value} \times \mathsf{Value} \times \mathsf{OpType} \rightarrow \mathsf{Value}
\end{aligned}
$$

$$\frac{\tau \in \mathsf{ComponentType}}{\tau \preceq \tau} \qquad \frac{cl' \in \mathit{superClasses}(cl)}{cl \preceq cl'} \qquad \frac{\mathit{iface} \in \mathit{implements}^*(cl)}{cl \preceq \mathit{iface}}$$

$$\frac{}{\mathit{iface} \preceq \texttt{Object}} \qquad \frac{\mathit{iface}' \in \mathit{superInterfaces}^*(\mathit{iface})}{\mathit{iface} \preceq \mathit{iface}'}$$

$$\frac{}{\tau\texttt{[ ]} \preceq \texttt{Object}} \qquad \frac{\tau \preceq \tau'}{\tau\texttt{[ ]} \preceq \tau'\texttt{[ ]}}$$

$$\frac{\begin{array}{c} H(loc) \in \mathsf{ArrayObject} \\ \tau = H(loc).\mathit{eleType} \end{array}}{H \vdash loc : \tau\texttt{[ ]}} \qquad \frac{\begin{array}{c} H(loc) \in \mathsf{ClassInst} \\ \tau.\mathit{class} = H(loc).\mathit{class} \end{array}}{H \vdash loc : \tau}$$

$$\frac{\tau \in \mathsf{ReferenceType}}{H \vdash \texttt{null} :\preceq \tau} \qquad \frac{(H \vdash loc : \tau') \wedge (\tau' \preceq \tau)}{H \vdash loc :\preceq \tau}$$

Fig. 4. Typing Rules

The *applyUnary* function requires as arguments the operand and result types. The result type is needed by the type conversion operator `to`. Binary operations do not perform type conversions and, therefore, operand and result types are the same.

`boolean` and `byte` values have to be sign-extended before they are pushed to the operand stack. Similarly, `short` values placed on the stack have to be truncated before they are stored into a field or array element of type `boolean` or `byte`. Value conversions from byte to short, and from short to byte are performed by the *toShort* and *toByte* functions, respectively. We write c^t to denote the runtime representation of numeric constant c in type t.

Figure 4 lists the typing rules for Carmel. The subtype relation $\tau \preceq \tau'$ is defined following the "can be assigned to" criteria as indicated by the description of `checkcast` and `instanceof` in [28]. We only show the runtime typing rules that correspond to runtime locations and reference types. The judgement $H \vdash loc : \tau$ assigns the Java reference type τ to the location loc. The judgement $H \vdash loc :\preceq \tau$ is used to indicate that location loc can be assigned to an element (a field or variable) of Java type τ i.e. if the runtime type of loc is a subtype of τ or if loc is `null`.

3.3 The JCRE domain

The JCRE domain ($J \in \mathsf{JCRE}$) models the data maintained and actions performed by the JCRE. This paper only deals with two aspects of the JCRE: the management of pre-allocated exception objects ($J.getException$, section 4.1) and current applet information ($J.appletContext$, section 5). A complete description of the JCRE domain is presented as part of the semantics of the Java Card system [25].

3.4 Frames

A frame contains information about the state of execution of a method. It is defined as follows:

$$
\begin{aligned}
F \in \mathsf{Frame} &= \mathsf{Owner} \times \mathsf{Method} \times \mathsf{Address} \times \mathsf{LocalVar} \times \mathsf{OperandStack} \\
\mathsf{StackValue} &= \mathsf{Value} - \mathsf{ByteValue} \\
S = v_1 :: \ldots :: v_n \in \mathsf{OperandStack} &= \mathsf{StackValue}^* \\
V \in \mathsf{LocalVar} &= [\mathsf{StackValue}_\perp]
\end{aligned}
$$

A frame $F = \langle\!\langle own, m, pc, V, S \rangle\!\rangle$ contains the following components:

- The owning context and applet, $own.context$ and $own.aid$, respectively. The owner of a frame is either an applet or the JCRE.
- The current method structure $m \in \mathsf{Method}$.
- The program counter register $pc \in \mathsf{Address}$, which contains the address of the instruction currently being executed in method m.
- The array of local variables[3] V.
- The operand stack S.

JCVM instructions take their dynamic operands from the operand stack and, in some cases, the array of local variables. The operand stack is defined as a sequence of values $S = v_1 :: \ldots :: v_n$ where values on top of the stack appear on the right-hand side of the sequence i.e. v_n is on top of the stack above. The empty stack is represented with the symbol ϵ. An element is pushed on top of the stack with the **::** operator ($S::v$) and two stacks are concatenated with the **:** operator ($S : S'$). For clarity we will write $v_1::v_2$ instead of $\epsilon::v_1::v_2$.

The operand stack stores values of all types except those of type `byte` (and `boolean`). As it was mentioned in section 3.2, some operations may require

[3] Note that this is not an array object but an array element, as defined in section 1.1

values to be truncated (or sign-extended) after they are popped from (or before they are pushed to) the stack. For such cases, the *fromStack*(v, τ) and *toStack*(v, τ) functions are used to determine if - by inspecting the type τ - type conversion is required and, if so, call the appropriate type conversion function i.e. *toShort* and *toByte*.

Stack operations do not carry type information, they only require that the integrity of the values is maintained i.e. multiword values must not be split apart.

A local variable is not referenced by name but by the position it occupies in the array of local variables. Uninitialised local variables are denoted by $\perp$. The array of local variables can store values of any of the stack types. Local variable slots are not statically typed so values of different types can be stored in the same slot at different times during the execution of a method. Special care must be taken when using multiword values. An `int` value stored in $V[i]$ takes two slots and access to $V[i+1]$ is prohibited. We model this by making $V[i+1] = \perp$ if position i holds an `int` value.

We write $V = v_0{::}\ldots{::}v_n$ when local variable array V is initialised with the sequence of values (usually from the operand stack) $v_0{::}\ldots{::}v_n$. Note that $V[0] = v_0$ but v_i does not necessarily correspond to $V[i]$ due to multi-word values.

3.5 Configurations

A runtime configuration describes the state of execution of the JCVM. We define three kinds of configurations:

$$\begin{aligned}
C &\in \mathsf{Config} = \mathsf{RConfig} \cup \mathsf{EConfig} \cup \mathsf{HConfig} \\
\langle\!\langle (J, K, H), SF \rangle\!\rangle &\in \mathsf{RConfig} = \mathsf{JCRE} \times \mathsf{StaticMem} \times \mathsf{Heap} \times \mathsf{Frame}^* \\
\langle\!\langle (J, K, H), \mathtt{E}(loc) \rangle\!\rangle &\in \mathsf{EConfig} = \mathsf{JCRE} \times \mathsf{StaticMem} \times \mathsf{Heap} \times \mathsf{Location} \\
\langle\!\langle (J, K, H), \mathtt{halt}(v) \rangle\!\rangle &\in \mathsf{HConfig} = \mathsf{JCRE} \times \mathsf{StaticMem} \times \mathsf{Heap} \times \mathsf{Value}_\perp
\end{aligned}$$

A running configuration ($\mathsf{RConfig}$) keeps track of the chain of invoked methods by using a stack of call frames. A running configuration $\langle\!\langle (J, K, H), SF \rangle\!\rangle$ contains the following components:

- The JCRE component $J \in \mathsf{JCRE}$ used to model the interaction between the JCVM and the JCRE via the API.
- The heap H, a mapping from locations to runtime objects (class and array instances).

- Static fields memory K, a mapping $\mathsf{FieldID} \rightarrow \mathsf{Value}$.
- The call stack SF, where $SF \in \mathsf{Frame}^*$.

A new frame is pushed on top of the stack when the JCVM executes an `invoke` instruction (section 6.3.3). A context switch occurs if the owning context of the new frame is different from the preceding frame. The frame of the method currently being executed - the current frame - sits on top of the call stack. Given $SF = SF'::\langle\!\langle own, m, pc, V, S\rangle\!\rangle$, the current instruction I can be extracted with $I = m.instructionAt(pc)$. The next instruction can be found at address $pc + 1$. The currently active context is the owning context of the current frame.

Exception ($\mathsf{EConfig}$) and halt ($\mathsf{HConfig}$) configurations represent terminal states of the JCVM. An (uncaught) exception configuration $\langle\!\langle (J, K, H), \mathtt{E}(loc)\rangle\!\rangle$ describes the state of the JCVM after it was unable to find a suitable exception handler for the exception object referenced by loc thrown during the execution of a running configuration. A halt configuration $\langle\!\langle (J, K, H), \mathtt{halt}(v)\rangle\!\rangle$ is reached after the first invoked method - the one represented by the frame at the bottom of the stack frame - successfully returns with value v. If the initial method has return type `void` then we shall write $\langle\!\langle (J, K, H), \mathtt{halt}()\rangle\!\rangle$.

4 Exceptions

An exception (object) is thrown whenever a program violates the semantic constraints of the JCVM or any other constraint specified by the programmer. All exceptions are instances of extensions of the `Throwable` class and can be thrown explicitly by the programmer using the `throw` instruction.

An exception is said to be caught if an appropriate exception handler is found in the current stack of call frames. If an exception is caught, control is transfered to the start address indicated by the exception handler and execution is resumed at that point. An uncaught exception will cause the virtual machine to halt by generating an exception configuration. The *catchException* function, defined below, takes a configuration and a reference to an exception object, and returns the configuration where execution shall resume. Let loc be a reference to an exception object:

$$catchException : \mathsf{RConfig} \times \mathsf{Location} \rightarrow \mathsf{RConfig} \cup \mathsf{EConfig}$$

$$catchException(\langle\!\langle (J,K,H), SF \rangle\!\rangle, loc) = \begin{cases} \langle\!\langle (J,K,H), r \rangle\!\rangle & \text{if } r \in \mathsf{Frame} \\ \langle\!\langle (J,K,H), \mathtt{E}(r) \rangle\!\rangle & \text{if } r \in \mathsf{Location} \end{cases}$$

$$\text{where } r = catchException^*(SF, H, loc)$$

$$catchException^* : \mathsf{Frame}^* \times \mathsf{Heap} \times \mathsf{Location} \rightarrow \mathsf{Frame}^* \cup \mathsf{Location}$$

$$catchException^*(\epsilon, H, loc) = loc$$

$$catchException^*(SF{::}\langle\!\langle own, m, pc, V, S \rangle\!\rangle, H, loc) = \begin{cases} catchException^*(SF, H, loc) & \text{if } pc' = \bot \\ \mathrm{SF}{::}\langle\!\langle own, m, pc', V, loc \rangle\!\rangle & \text{if } pc' \neq \bot \end{cases}$$

$$\text{where } H \vdash loc : \tau \wedge pc' = findHandler(m, pc, H, loc)$$

The $catchException^*$ function searches the stack of call frames until the first appropriate exception handler is found [29, Chapter 7]. The $findHandler(m, pc, \tau)$ function searches for the first exception handler in method m that can handle an exception of class τ at address pc. Information about exception handlers is kept by the $m.Handlers$ structure. The detailed semantics of $findHandler$ is described in [26].

4.1 Runtime Exceptions

Runtime exceptions ($\mathsf{RTException}$) may be raised by the JCVM during the execution of a program. Being memory usage an important concern, the creation of unique runtime exception objects per class is enforced. These are JCRE owned, unique, pre-allocated exceptions objects designated as temporary JCRE entry points. The JCRE keeps a reference to each instance of the Java `RuntimeException` objects. These references can be accessed using the *getException* function. The *getException* function has the following property:

$$getException \in \mathsf{JCRE} \rightarrow \mathsf{RTException} \rightarrow \mathsf{Location}$$

$$\forall cl \in \mathsf{RTException}, \exists loc \in \mathsf{Location} :$$

$$(loc = J.getException(cl)) \wedge (o = H(loc)) \wedge (o.class = cl)$$

$$(o.owner = JCRE) \wedge (o.entryPoint = \mathtt{temp})$$

5 The JCVM Firewall

The Java Card platform partitions the object system into separate protected object spaces called contexts. A `Context` is either a package AID or the special `JCRE` context associated to actions performed by the JCRE. An object is owned[4] by the applet that was active (currently active owner) when the object was created (see `new` instruction, section 6.3.2) or by the JCRE. An applet's owning context is determined by the AID of the package where the applet is declared. An object's owning context is the same as the owning context of its owner, which in general means that the owning context of an object is the package where the object's class is declared. The JCVM firewall provides a security mechanism that prevents an object from being accessed by code running in a different context i.e. by an object with a different owning context, thus providing a coarse type of applet isolation. Object and field accesses among different instances of the same applet or instances of applets declared in the same package (group context) is allowed since they share the same context.

Prior to performing an access operation on an object, the JCVM performs a runtime check that depends on the instruction, the currently active context ($ctxt$), and the type and owner of the referenced object (o) [28, Chapter 6]. Additionally, store operations check the value that is being stored e.g. $canStore(H, v)$. The firewall runtime checks are defined in Figure 5. The JCVM will throw an instance of the class `SecurityException` whenever a firewall violation has been detected. In general, access to an object is disallowed if the current context - the context of the current frame - is different from the owning context of the object that is being accessed (equation 5). In addition the JCRE maintains its own context. This context has special system privileges and can access any object (equation 4). Thus, in the simplest case, access of object o from context $ctxt$ is allowed if any of the following conditions are met:

$$isJCRE(ctxt) \equiv (ctxt = \mathtt{JCRE}) \tag{4}$$

$$sameContext(ctxt, o) \equiv (ctxt = o.owner.context) \tag{5}$$

This is the case of the *checkGetField* predicate: an instance field of an object (o) can only be accessed if the currently active context ($ctxt$) is either the JCRE or the owning applet of the object. These conditions alone would make the firewall system too restrictive. The JCRE provides several mechanisms to allow object access across contexts: global arrays, JCRE entry points and shareable interfaces. JCRE entry point objects - temporary and permanent - and global arrays are objects owned by the JCRE that can be accessed from any context. We say that an object o has global access if the conditions defined

[4] An applet is owned by itself.

$$
\begin{aligned}
&\mathit{checkGetField}(\mathit{ctxt}, o) \equiv \mathit{isJCRE}(\mathit{ctxt}) \vee \mathit{sameContext}(\mathit{ctxt}, o)\\
&\mathit{checkPutStatic}(\mathit{ctxt}, H, v) \equiv \mathit{isJCRE}(\mathit{ctxt}) \vee \mathit{canStore}(H, v)\\
&\mathit{checkPutField}(\mathit{ctxt}, o, H, v) \equiv \mathit{isJCRE}(\mathit{ctxt}) \vee (\mathit{sameContext}(\mathit{ctxt}, o) \wedge\\
&\qquad\qquad \mathit{canStore}(H, v))\\
&\mathit{checkThrow} = \mathit{checkInvokeDefinite} = \mathit{objectAccess}\\
&\mathit{checkArrayStore}(\mathit{ctxt}, \mathit{app}, a, H, v) \equiv \mathit{objectAccess}(\mathit{ctxt}, \mathit{app}, a) \wedge\\
&\qquad\qquad (\neg \mathit{isJCRE}(\mathit{ctxt}) \Rightarrow \mathit{canStore}(H, v))\\
&\mathit{checkArrayLoad} = \mathit{checkInvokeVirtual} = \mathit{objectAccess}\\
&\mathit{checkInvokeIface}(\mathit{ctxt}, \mathit{app}, o, \mathit{iface}) \equiv\\
&\qquad \mathit{objectAccess}(\mathit{ctxt}, \mathit{app}, o) \vee \begin{pmatrix} \mathit{transAccess}(\mathit{ctxt}, \mathit{app}, o) \\ \wedge\ \mathit{isShareable}(o, \mathit{iface}) \end{pmatrix}\\
&\mathit{checkCast}(\mathit{ctxt}, \mathit{app}, o, \tau) \equiv \begin{cases} \mathit{objectAccess}(\mathit{ctxt}, \mathit{app}, o) \vee (\tau \in \mathsf{Interface} \\ \wedge\ \mathit{transAccess}(\mathit{ctxt}, \mathit{app}, o) \wedge \mathit{isShareable}(o, \tau)) \end{cases}
\end{aligned}
$$

Fig. 5. JCVM Firewall Checks

in equation 6 are true. Objects with global access have a special restriction. References to objects designated as temporary JCRE entry points and global arrays cannot be stored in class fields, instance fields or array components, as defined by the *canStore* predicate in (7).

$$
\begin{aligned}
\mathit{globalAccess}(o) \equiv\ & (o \in \mathsf{ClassInst} \wedge o.\mathit{entryPoint} \neq \mathtt{no}) \vee \\
& (o \in \mathsf{ArrayObject} \wedge o.\mathit{global})
\end{aligned} \tag{6}
$$

$$
\mathit{canStore}(H, v) \equiv (v \in \mathsf{ReferenceValue}) \Rightarrow \mathit{canStoreObject}(H(v)) \tag{7}
$$

$$
\begin{aligned}
\mathit{canStoreObject}(o) \equiv\ & (o \in \mathsf{ArrayObject} \wedge \neg o.\mathit{global}) \vee \\
& (o \in \mathsf{ClassInst} \wedge o.\mathit{entryPoint} \neq \mathtt{temp})
\end{aligned}
$$

The *objectAccess* predicate (equation 8) spells out the general conditions under which access to object o is allowed. The predicate includes an additional check for `CLEAR_ON_DESELECT` transient objects (*transAccess*, equation 9). Transient objects of `CLEAR_ON_DESELECT` type can only be accessed when the currently active context (*ctxt*) is the context of the currently selected applet (*app*) [5] . We have:

[5] The currently selected applet can be obtained from the JCRE element by calling the *appletContext* function.

$$\begin{aligned} &\mathit{objectAccess}(\mathit{ctxt}, \mathit{app}, o) \equiv \mathit{isJCRE}(\mathit{ctxt}) \vee \\ &\quad \left(\mathit{transAccess}(\mathit{ctxt}, \mathit{app}, o) \wedge \left(\begin{array}{c} \mathit{sameContext}(\mathit{ctxt}, o) \\ \vee\ \mathit{globalAccess}(o) \end{array} \right) \right) \end{aligned} \tag{8}$$

$$\begin{aligned} &\mathit{transAccess}(\mathit{ctxt}, \mathit{app}, o) \equiv (o \in \mathsf{ArrayObject}) \ \Rightarrow \\ &\quad ((o.\mathit{transient} = \texttt{CLEAR\_ON\_DESELECT}) \ \Rightarrow \ (\mathit{ctxt} = \mathit{app})) \end{aligned} \tag{9}$$

Currently, only arrays with primitive type components and arrays with references to `Object` can be transient objects. It may seem an abuse of notation to include *transAccess* as part of the firewall checks of instructions that do not deal with arrays but, by doing so, we provide a general criterion for most firewall checks. Later versions of the Java Card platform may allow transient class instances. In the worst case, the *transAccess* condition is trivially true. Still, we would like to point out that it is possible to invoke methods of the `Object` class on an array.

Another mechanism that allows object access across contexts is the use of shareable interfaces. A shareable interface is an interface that extends the `javacard.framework.Shareable` interface. Methods of shareable interfaces can be invoked from one context even if the object implementing them is owned by an applet in another context. An interface method declared in *iface* can be invoked on object o iff:

$$\mathit{isShareable}(o, \mathit{iface}) \equiv \begin{cases} \texttt{Shareable} \in \mathit{implements}^*(o.\mathit{class}) \wedge \\ \texttt{Shareable} \in \mathit{superInterfaces}^*(\mathit{iface}) \end{cases}$$

that is, if the object's class implements a `Shareable` interface i.e. o is a Shareable Interface Object, and if the interface in question extends the `Shareable` interface. The same conditions are used when the `checkcast` or `instanceof` instructions are applied to an interface type.

6 The Operational Semantics of Carmel

6.1 The Evaluation Relation

The operational semantics of Carmel is defined as a relation $\Rightarrow$ between program configurations. We say that the judgement $C \ \Rightarrow \ C'$ is true if configuration C' is the result of executing the current instruction (section 3.5) in configuration C.

We give the evaluation rules of Carmel in section 6.3. The evaluation rules detail the set of conditions under which the assertion $C \Rightarrow C'$ is true. The first condition in all the rules determines the instruction that is being executed: the current instruction. For example, in the rule for the `nop` instruction:

$$\frac{I = m.instructionAt(pc) \ \wedge \ I = \texttt{nop} \ \wedge \ G = (J, K, F)}{\langle G, SF{::}\langle own, m, pc, V, S\rangle\rangle \ \Rightarrow \ \langle G, SF{::}\langle own, m, pc+1, V, S\rangle\rangle}$$

the current instruction is obtained with $m.instructionAt(pc)$, where m and pc are the method and program counter kept by the current frame. The `nop` does not modify the state of the program except for the program counter. The program counter is incremented to point to the next instruction ($pc+1$) where execution should resume. Since the process of reading the instruction is the same for all rules, from now on we will consider $I = m.instructionAt(pc)$ as the default and write $I = \texttt{instruction}$. We will also assume $G = (J, K, F)$ unless stated otherwise.

The semantics presented in this section correspond to the semantics of the offensive virtual machine defined in [29]. Therefore, we only show the conditions and changes to the program state that result when the current execution succeeds. Successful execution of the instructions is ensured by the static conditions checked by the bytecode verified. Runtime violations of the semantics of the language - that can not be detected by the verifier - are reported by the JCVM by throwing an exception. The process of throwing an exception is explained in the next section.

6.2 Throwing Exceptions

The JCVM performs a series of runtime checks that preserve the semantic constraints of the language. These runtime checks appear in the evaluation rules as predicates prefixed with the word *check*. We have already defined the predicates that enforce the firewall conditions in section 5. The violation of a firewall predicate will result in a `SecurityException`. The rest of the predicates that determine the conditions for throwing a runtime exception are:

$$
\begin{aligned}
&\mathit{checkNotNull}(loc) &&\equiv (loc \neq \texttt{null})\\
&\mathit{checkArithmetic}(op, v_1, v_2) &&\equiv op \notin \{\texttt{div}, \texttt{rem}\} \vee v_2 \neq 0\\
&\mathit{checkNotNegative}(n) &&\equiv (n \geq 0)\\
&\mathit{checkCanCast}(H, loc, \tau) &&\equiv H \vdash loc :\preceq \tau\\
&\mathit{checkArrayBounds}(a, i) &&\equiv (0 \leq i < a.\mathit{length})
\end{aligned}
\tag{10}
$$

We associate the runtime checks defined above to the the null pointer, arithmetic, negative size array, class cast and array index out of bounds exceptions [6] , respectively. The violation of any of these runtime checks will trigger the JCVM to throw the runtime exception associated with them. For example, if the *checkNotNull*(loc) condition is not true in a particular rule then the null pointer exception is thrown by the JCVM. If more than one runtime check is violated, then the exception associated to the runtime check that appears first in the rule will be thrown.

Every instruction that can potentially throw an exception has a separate rule that describes the detailed conditions of such event. Due to lack of space we have not included all the exception rules but show an example instead. Let's have a look at the rule for `getfield` of Figure 7. The corresponding rule that describes the process of throwing a runtime exception is shown below. Once the runtime violation is found, the JCVM throws the JCRE-owned exception object by first calling the *getException* function. If an exception handler is found, execution is resumed from the frame returned by *catchException*. Otherwise, an exception configuration is created.

$$
\frac{
\begin{array}{l}
I = \texttt{getfield}\ f \ \wedge\ S = S'{::}loc \ \wedge \\
e = \begin{cases} \texttt{NullPointerException} & \text{if } \neg \mathit{checkNotNull}(loc) \\ \texttt{SecurityException} & \text{if } \neg \mathit{checkGetField}(\mathit{own.context}, H(loc)) \end{cases} \wedge \\
loc' = \mathit{getException}(J, e) \ \wedge \\
r = \mathit{catchException}(\langle\!\langle (J, K, H), SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\rangle\!\rangle, loc')
\end{array}
}{
\langle\!\langle (J, K, H), SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\rangle\!\rangle \ \Rightarrow\ C'
}
$$

[6] `NullPointer`, `ArithmeticException`, `NegativeArraySizeException`, `ClassCastException` and `ArrayIndexOutOfBoundsException`.

6.3 Evaluation Rules

6.3.1 The Core Language

The core language supports instructions for stack manipulation, local variable access, numerical operation and branching. Stack instructions deal explicitly with the operand stack. They do not carry type information - with the sole exception of `push` that uses the operand type to put the right kind of value on top of the stack - and operate ignoring the type of data on the stack. Instead, stack instructions operate at the word level and their operands indicate the number of words to be manipulated. For example, the execution of the `pop` n instruction results on the the top n words being popped from the operand stack. A more elaborate rule is shown below:

$$\frac{I = \texttt{dup}\ n\ d\ \wedge\ S = S_3 : S_2 : S_1\ \wedge \quad n = nbWords(S_1) \wedge d = nbWords(S_2 : S_1)\ \wedge\ S' = S_3 : S_1 : S_2 : S_1}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S'\rangle\!\rangle\rangle\!\rangle}$$

Execution of the `dup` instruction results on a new configuration where the top n words on the operand stack have been duplicated and inserted d words down from the top. When d equals 0, the top n words are copied and placed on top of the stack. The rules respect the restriction that preserve the integrity of two-word values by making sure that only whole values are moved. This is also true for the instructions that manipulate the array of local variables V where, as we pointed out in section 3.4, the special symbol $\perp$ is used to "block" the index that corresponds to the second word of an `int` value. The `load` instruction places the value of local variable i on top of the stack and `store` updates local variable i with the value found on top of the stack. The success of the rules is ensured by the bytecode verification process which we assume has taken place.

The `numop` instruction uses the top two elements of the stack - after removing them - as arguments to the numeric operation op and places the result back on the stack. The `goto` instruction jumps to the instruction found in address $addr$. The `if` instruction - shown below - jumps to $addr$ only if the result of performing the comparison operation cmp between the top two elements of

$$\frac{I = \texttt{push}\ \ t\ \ c}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S{::}c^t\rangle\!\rangle}$$

$$\frac{I = \texttt{pop}\ \ n\ \wedge\ n = nbWords(S')}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S : S'\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S\rangle\!\rangle}$$

$$\frac{\begin{array}{c} I = \texttt{swap}\ \ n_1\ \ n_2\ \wedge\ S = S_3 : S_2 : S_1\ \wedge \\ n_1 = nbWords(S_1) \wedge n_2 = nbWords(S_2)\ \wedge\ S' = S_3 : S_1 : S_2 \end{array}}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S'\rangle\!\rangle}$$

$$\frac{\begin{array}{c} I = \texttt{numop}\ \ t\ \ op \wedge \\ checkArithmetic(op, v_1, v_2)\ \wedge\ r = applyBinary(op, v_1, v_2, t) \end{array}}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}v_1{::}v_2\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S{::}r\rangle\!\rangle}$$

$$\frac{I = \texttt{load}\ \ t\ \ i}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S{::}V[i]\rangle\!\rangle}$$

$$\frac{I = \texttt{store}\ \ t\ \ i\ \wedge\ V' = \begin{cases} V[i \mapsto v][i+1\ \mapsto\ \bot] & \text{if } (t = \texttt{i}) \\ V'[i\ \mapsto\ v] & \text{otherwise} \end{cases}}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}v\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V', S\rangle\!\rangle}$$

$$\frac{I = \texttt{inc}\ \ t\ \ i\ \ c\ \wedge\ r = V[i] + c}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V[i\ \mapsto\ r], S\rangle\!\rangle}$$

$$\frac{I = \texttt{goto}\ \ addr}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, addr, V, S\rangle\!\rangle}$$

Fig. 6. The Core Language Semantics

the stack is true, otherwise it continues to the next instruction.

$$\frac{I = \texttt{if}\ \ t\ \ cmp\ \ \texttt{goto}\ \ addr \wedge \quad pc' = \begin{cases} addr & \text{if } applyComparison(cmp, v_1, v_2) \\ pc + 1 & \text{otherwise} \end{cases}}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}v_1{::}v_2\rangle\!\rangle\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc', V, S\rangle\!\rangle\rangle\!\rangle}$$

The rest of the evaluation rules for the instructions of the core language are shown in Figure 6. We have not included the evaluation rules of the `if` and `op` instructions that handle null comparisons and unary operations. We have also excluded the `keyswitch` and `indexswitch` instructions. The complete formalisation of these instructions can be found in [26].

6.3.2 The Object Language

The instructions described in this section deal with the creation and manipulation of class instances. The evaluation rules for the instructions of the object language are shown in figure 7. The `new` *cl* instruction creates a new instance object of class *cl* by calling the *newObject* relation defined in (1). The owner of the current frame is designated as the owner of the new object. The `getstatic` and `putstatic` instructions use the field ID $f.id$ to access the static field f in K. The rest of the instructions take the location of the object to be read, modified or inspected from the stack. For example, in the case of the `putfield` instruction shown below,

$$\frac{\begin{array}{l} I = \texttt{putfield}\ \ f \ \wedge\ v' = fromStack(v, f.type) \ \wedge\ o = H(loc) \wedge \\ checkNotNull(loc) \wedge checkPutField(own.context, o, H, v) \wedge \\ H' = updateEle(H, (loc, f.id), v') \end{array}}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S'{::}loc{::}v\rangle\!\rangle\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G', SF{::}\langle\!\langle own, m, pc+1, V, S'\rangle\!\rangle\rangle\!\rangle}$$

the location of the object to be modified and the new value of field f are taken from the operand stack. A type conversion is performed, if needed, by the *fromStack* function (section 3.4). The new value is stored in field f of the resolved object o using the *updateEle* function (equation 3). The `putfield this` and `getfield this` instructions (not included here) take the location of the object from $V[0]$, the location of the current object.

The semantics of `checkcast` and `instanceof` are given by the definition of $H \vdash loc :\preceq \tau_{\Gamma}$ (Figure 4). The difference between `checkcast` and `instanceof` is that the former throws a `ClassCastException` when the operation is not

Let $app = J.appletContext$

$$\frac{I = \texttt{new}\ \ cl \wedge cl \in \mathsf{Class}\ \wedge \quad \langle H', loc\rangle \in newObject(own, cl, \texttt{no}, H)}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G', SF{::}\langle\!\langle own, m, pc+1, V, S{::}loc\rangle\!\rangle}$$

$$\frac{I = \texttt{getstatic}\ \ f \wedge v = toStack(K(f.id), f.type)}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S{::}v\rangle\!\rangle}$$

$$\frac{I = \texttt{putstatic}\ \ f\ \wedge\ checkPutStatic(own.context, H, v) \wedge \quad v' = fromStack(v, f.type)\ \wedge\ K' = K[f.id\ \mapsto\ v']}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}v\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G', SF{::}\langle\!\langle own, m, pc+1, V, S\rangle\!\rangle}$$

$$\frac{I = \texttt{getfield}\ \ f \wedge checkNotNull(loc) \wedge o = H(loc) \wedge \quad checkGetField(own.context, o) \wedge v = toStack(o.fieldValue(f.id), f.type)}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S'{::}loc\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S'{::}v\rangle\!\rangle}$$

$$\frac{I = \texttt{checkcast}\ \ \tau_{\mathrm{r}}\ \wedge\ checkCanCast(H, loc, \tau_{\mathrm{r}}) \wedge \quad loc \neq \texttt{null}\ \Rightarrow\ checkCast(own.context, app, \tau_{\mathrm{r}}, H(loc))}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}loc\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S{::}loc\rangle\!\rangle}$$

$$\frac{I = \texttt{instanceof}\ \ \tau_{\mathrm{r}} \wedge v = \begin{cases} 1 & \text{if } loc \neq \texttt{null} \wedge H \vdash loc :\preceq \tau_{\mathrm{r}} \\ 0 & \text{otherwise} \end{cases} \wedge \quad loc \neq \texttt{null}\ \Rightarrow\ checkCast(own.context, app, \tau_{\mathrm{r}}, H(loc))}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}loc\rangle\!\rangle \ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S{::}v\rangle\!\rangle}$$

Fig. 7. The Object Language Semantics

valid (*checkCanCast* predicate, equation 10). Firewall security checks take as arguments the current context, the current applet context, the object being accessed and, in the case of `putstatic` and `putfield`, the value being stored.

6.3.3 Method Language

The evaluation rules for the instructions of the method language are shown in figure 8. The `invokedefinite` instruction is used for static methods and for non-static methods that have been resolved statically. The `invokedefinite` instruction contains a resolved method as argument. No method search through the class hierarchy is needed, even if the original call was made to an inherited method. The method has been resolved statically i.e. before execution when loaded and linked.

The `invokevirtual` and `invokeinterface` instructions perform a method lookup based on the class of the object where the method is being invoked (target object). This search is performed by the *lookupID* function. All methods sharing the same signature are assigned the same method ID. Thus, if a method overrides another one or if an instance method implements an interface method, their method IDs are equal. Given a method ID and a class cl, the function $lookupID(id, cl)$ returns the first method structure with the same method ID found in the superclass hierarchy of class cl.

An array may be the target of an `invokevirtual` instruction. If that's the case then the search must be performed on the `Object` class structure. This check is performed by the *methodLookup* function which, given a method m and an object o, returns the correct method. The definitions of *lookupID* and *methodLookup* are listed below:

$$lookupID : \mathsf{MethodID} \times \mathsf{Class}_\perp \to \mathsf{Method}_\perp$$

$$lookupID(id, \perp) = \perp$$

$$lookupID(id, cl) = \begin{cases} m & \text{if } \exists m \in cl.methods.\ m.id = ID \\ lookupID(id, cl.superclass) & \text{otherwise} \end{cases}$$

$$methodLookup : \mathsf{Method} \times \mathsf{Object} \to \mathsf{Method}_\perp$$

$$methodLookup(m, o) = \begin{cases} lookupID(m'.id, \texttt{Object})\ o \in \mathsf{ArrayObject} \\ lookupID(m'.id, o.class) \ \ \text{otherwise} \end{cases}$$

We describe the execution of the `invokevirtual` instruction (see rule below). The top $n+1$ values - corresponding to the n method arguments and object location (`this`) - are popped from the operand stack of the current frame and placed in the local variables array of the new frame F'. The owner of the

Let $app = J.appletContext$:

$$\frac{I = \texttt{invokedefinite}\ \ m' \ \wedge\ m'.isStatic \ \wedge\ n = |m'| \wedge \quad S_0 = S{::}v_1{::}\ldots{::}v_n \wedge V' = v_1{::}\ldots{::}v_n \wedge F' = \langle own, m', 0, V', \epsilon\rangle}{\langle G, SF{::}\langle own, m, pc, V, S_0\rangle\rangle \ \Rightarrow\ \langle G, SF{::}\langle own, m, pc, V, S\rangle{::}F'\rangle}$$

$$\frac{\begin{array}{c} I = \texttt{invokedefinite}\ \ m' \ \wedge\ \neg m'.isStatic \wedge \\ n = |m'| \ \wedge\ S_0 = S{::}loc{::}v_1{::}\ldots{::}v_n \ \wedge\ checkNotNull(loc) \wedge \\ o = H(loc) \ \wedge\ checkInvokeDefinite(own.context, app, o) \wedge \\ V' = loc{::}v_1{::}\ldots{::}v_n \wedge F' = \langle o.owner, m', 0, V', \epsilon\rangle \end{array}}{\langle G, SF{::}\langle own, m, pc, V, S_0\rangle\rangle \ \Rightarrow\ \langle G, SF{::}\langle own, m, pc, V, S\rangle{::}F'\rangle}$$

$$\frac{\begin{array}{c} I = \texttt{invokeinterface}\ \ m' \ \wedge\ n = |m'| \ \wedge\ S_0 = S{::}loc{::}v_1{::}\ldots{::}v_n \wedge \\ checkNotNull(loc) \wedge o = H(loc) \ \wedge\ m_v = methodLookup(m', o) \wedge \\ checkInvokeIface(own.context, app, m'.classI, o) \wedge \\ V' = loc{::}v_1{::}\ldots{::}v_n \wedge F' = \langle o.owner, m_v, 0, V', \epsilon\rangle \end{array}}{\langle G, SF{::}\langle own, m, pc, V, S_0\rangle\rangle \ \Rightarrow\ \langle G, SF{::}\langle own, m, pc, V, S\rangle{::}F'\rangle}$$

$$\frac{I = \texttt{return}\ \ t}{\langle G, \langle own, m, pc, V, S{::}v\rangle\rangle \ \Rightarrow\ \langle G, \texttt{halt}(v)\rangle}$$

$$\frac{I = \texttt{return} \ \wedge\ F' = \langle own', m', pc', V', S'\rangle}{\langle G, SF{::}F'{::}\langle own, m, pc, V, S\rangle\rangle \ \Rightarrow\ \langle G, SF{::}\langle own', m', pc' + 1, V', S'\rangle\rangle}$$

$$\frac{I = \texttt{return}}{\langle G, \langle own, m, pc, V, S\rangle\rangle \ \Rightarrow\ \langle G, \texttt{halt}()\rangle}$$

Fig. 8. The Method Language Semantics

target object o is set as owner of the new frame. The new operand stack is set to empty and, if both runtime checks succeed, the new frame is pushed onto the call frame stack. The current method is now m_v and execution will continue at the first instruction of the method ($pc = 0$). A context switch occurs when a method of an object owned by a different context is invoked i.e. two consecutive frames have different contexts.

$$\frac{\begin{array}{c} I = \texttt{invokevirtual}\ \ m' \ \wedge\ n = |m'| \ \wedge \\ F = \langle own, m, pc, V, S{::}loc{::}v_1{::}\ldots{::}v_n\rangle \wedge \\ \mathit{checkNotNull}(loc) \ \wedge\ o = H(loc) \ \wedge\ m_v = \mathit{methodLookup}(m', o) \wedge \\ \mathit{checkInvokeVirtual}(own.context, J.appletContext, o) \wedge \\ V' = loc{::}v_1{::}\ldots{::}v_n \ \wedge\ F' = \langle o.owner, m_v, 0, V', \epsilon\rangle \end{array}}{\langle G, SF{::}F\rangle \ \Rightarrow\ \langle G, SF{::}\langle own, m, pc, V, S\rangle{::}F'\rangle}$$

Method return is handled by the `return` instruction. We have identified four cases. The first case,

$$\frac{\begin{array}{c} I = \texttt{return}\ \ t \ \ \wedge \\ F = \langle own, m, pc, V, S{::}v\rangle \ \wedge\ F' = \langle own', m', pc', V', S'\rangle \end{array}}{\langle G, SF{::}F'{::}F\rangle \ \Rightarrow\ \langle G, SF{::}\langle own', m', pc' + 1, V', S'{::}v\rangle\rangle}$$

corresponds to the return from a method with return type different from `void`. The value v is popped from the operand stack of the current frame and pushed onto the operand stack of the frame of the invoker. The virtual machine then removes the current frame and returns control to the invoker. If the current frame is the last remaining frame in the configuration then the JCVM copies the return value v to the special halt configuration and stops execution, as shown by the `return` t rule in Figure 8. The last two cases - shown in Figure 8 - correspond to `void` returns.

6.3.4 The Array, Exception and Subroutine Language

The evaluation rules of the instructions that deal with the creation and manipulation of arrays, the explicit throwing of exceptions by the programmer and the call of (and return from) subroutines, are listed in Figure 9. A new array is created by the `new` τ instruction when the parameter τ denotes an array type. New arrays are initialised as non-global, non-transient objects by the *newArray* relation (equation 2). The owner of the current frame is set as the owner of the new array. The location of the new array is placed on top of

Let $app = J.appletContext$ and $G' = (J, H', K)$:

$$\frac{I = \texttt{new}\ \ \tau\texttt{[ ]} \wedge checkNotNegative(n) \wedge \quad \langle H', loc\rangle \in newArray(own, \tau, n, \texttt{false}, \texttt{NOT\_TRANSIENT}, H)}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}n\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G', SF{::}\langle\!\langle own, m, pc+1, V, S{::}loc\rangle\!\rangle}$$

$$\frac{I = \texttt{arraylength}\ \wedge\ checkNotNull(loc) \wedge \quad checkArrayLoad(own.context, app, H(loc))\ \wedge\ v = H(loc).length}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}loc\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc+1, V, S{::}v\rangle\!\rangle}$$

$$\frac{I = \texttt{arraystore}\ \ t\ \wedge\ checkNotNull(loc)\ \wedge\ a = H(loc) \wedge \quad checkArrayBounds(a, i)\ \wedge\ checkArrayStore(own.context, app, a, H, v) \wedge \quad v' = fromStack(v, t)\ \wedge\ H' = updateEle(H, (loc, i), v')}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}loc{::}i{::}v\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G', SF{::}\langle\!\langle own, m, pc+1, V, S\rangle\!\rangle}$$

$$\frac{I = \texttt{throw}\ \wedge \quad checkNotNull(loc) \wedge checkThrow(own.context, app, H(loc)) \wedge \quad C' = catchException(\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle, loc)}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S{::}loc\rangle\!\rangle\ \Rightarrow\ C'}$$

$$\frac{I = \texttt{jsr}\ \ addr}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, addr, V, S{::}pc+1\rangle\!\rangle}$$

$$\frac{I = \texttt{ret}\ \ i\ \wedge\ pc' = V[i]}{\langle\!\langle G, SF{::}\langle\!\langle own, m, pc, V, S\rangle\!\rangle\ \Rightarrow\ \langle\!\langle G, SF{::}\langle\!\langle own, m, pc', V, S\rangle\!\rangle}$$

Fig. 9. The Array, Exception and Subroutine Language Semantics

the operand stack of the current frame. The other array instructions take their arguments - array location and index - from the operand stack. Transfer of data from the operand stack to an array object, and vice versa, is performed using the *fromStack* and *toStack* functions, respectively, as shown by the rule for `arrayload` below.

$$\frac{\begin{array}{l} I = \texttt{arrayload}\ t \ \wedge\ \mathit{checkNotNull}(loc)\ \wedge\ a = H(loc) \wedge \\ \mathit{checkArrayBounds}(a,i)\ \wedge\ v = \mathit{toStack}(a.\mathit{values}[i],t) \wedge \\ \mathit{checkArrayLoad}(own.context, J.appletContext, a) \end{array}}{\langle G, SF{::}\langle own, m, pc, V, S{::}loc{::}i\rangle\rangle\ \Rightarrow\ \langle G, SF{::}\langle own, m, pc+1, V, S{::}v\rangle\rangle}$$

The rest of the rules deal with the Exception and Subroutine language. The `throw` instruction is used by the programmer to throw the exception object referenced by the location placed on top of the stack. The location *loc* on top of the stack must be a reference to an object that is an instance of class `Throwable` or of a subclass of `Throwable`. The process of catching the exception referenced by *loc* is performed by *catchException* (section 4), where *loc* has been popped from the operand stack of the top frame. The `jsr` instruction jumps to the subroutine address indicated by its operand and places the address of the next instruction on top of the stack. The `ret` instruction jumps to the address stored in local variable i.

7 Related Work

Since its official release, the Java language has attracted considerable interest from the research community. Hartel and Moreau [15] provide a comprehensive review of the substantial amount of work dedicated to the study of each of the main components of Java: the Java Language, the Java Virtual Machine (JVM) and Java to JVM compiler; and most of its interesting features: object-orientation, dynamic class loading, garbage collection, multi-threading, its type system and bytecode verification. However, as pointed out in [15], no single attempt has been made at specifying full Java, the full JVM, or the full compiler.

Most research of the semantics of Java Card has been produced by - or is related to - the formal methods and program verification community. Formal specifications of the semantics of Java Card can be traced back to [14] (virtual machine) and [2] (source code). Hartel et al.provide a complete specification of the Java Secure Processor (JSP) [14], a precursor of the JCVM. JSP is a virtual machine designed to fit on a smart card and, as such, does not support multi-threading, garbage collection and exception handling. The specifications are defined and validated using the LETOS tool [13]. Attali et al. adapt their work with Java and the Centuar tool [3] to Java Card and provide a programming environment for Java Card applications [2]. The research group at Gemplus has produced some work related to the application of the B-method [1] to the formalisation of the semantics of Java Card [17,24,7] .

CertiCartes [5,4] constitutes one of the most in-depth machine-checked accounts of the Java Card platform up to date. It contains formal executable specifications written in Coq [9] of a defensive JCVM, an offensive JCVM and an abstract JCVM together with the specification of a Java Card Bytecode Verifier (BCV) presented as a data-flow analyser based on the abstract virtual machine.

The LOOP project [19] project is involved in the application of formal methods to object oriented languages. The aim is to specify and verify properties of classes in object-oriented languages, using proof tools like PVS [23] and Isabelle/HOL [22].Its main contribution to the formalisation of the semantics of Java Card has been the formal specification of the Java Card API in JML and ESC/Java [21,16].

8 Conclusions and Future Work

The work presented in this paper has been developed in the context of the Secsafe project. The Secsafe project is concerned with the application of static analysis technology to the validation of security and safety aspects of realistic languages and applications. The project has focused a substantial part of its efforts on the development of methods that can be applied to the domain of smart cards and, in particular, Java Card and its applications (applets). In this context, a formalisation of the semantics of Java Card provides the framework where we can specify security properties, derive control and data flow analyses [12] that safely approximate such properties, and prove the analyses correct. Further refinements to the semantics will be directed - among other things - towards this goal.

Our first step have been to formalise the JCVM by defining the semantics of Carmel, a complete representation of the JCVML. The results of this work are presented in this paper. A more in-depth account of this specification can be found at our technical report [26].

We have seen that the specification presented in this paper relies on elements of the JCRE. Other important features of the JC platform are implemented by the API. In order for this specification to be useful, we need to placed it in context with the runtime environment and formalise its interaction with the off-card applications. The next step has been to provide a formal specification of the JCRE and the API. This is work in progress and can be found in [25].

References

[1] J.-R. Abrial, The B-Book: Assigning Programs to Meanings, Cambridge University Press, 1996.

[2] I. Attali, D. Caromel, C. Courbis, L. Henrio, H. Nilsson, Smarttools for Java Card, in: J. Domingo-Ferrer, D. Chan and A. Watson (Eds): 4th Smart Card Research and Advanced Application Conf. (CARDIS), Kluwer Academic Publishers, 2000, pp. 155–174.

[3] I. Attali, D. Caromel, M. Russo, A formal executable semantics for java, in: Proceedings of Formal Underpinnings of Java - An OOPSLA'98 Workshop, Vancouver, Canada, 1998.

[4] G. Barthe, G. Dufay, L. Jakubiec, S. M. de Sousa, A formal correspondence between offensive and defensive JavaCard virtual machines, in: A. Cortesi (Ed.), Verification, model checking and abstract interpretation: third international workshop, VMCAI 2002, Vol. 2294 of Lecture Notes in Computer Science, Springer-Verlag, 2002.

[5] G. Barthe, G. Dufay, L. Jakubiec, B. Serpette, S. M. de Sousa, A formal executable semantics of the JavaCard platform, in: D. Sands (Ed.), Programming languages and systems: 10th European Symposium on Programming, ESOP 2001, held as part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2001, Vol. 2028 of Lecture Notes in Computer Science, 2001.

[6] P. Bertelsen, Semantics of Java byte code, Technical report, Department of Information Technology, Technical University of Denmark, http://www.dina.kvl.dk/~pmb/ (Mar. 1997).

[7] L. Casset, Development of an embedded verifier for Java card byte code using formal methods, in: L.-H. Eriksson, P. A. Lindsay (Eds.), FME 2002: formal methods-getting it right: International Symposium of Formal Methods Europe, Vol. 2391 of Lecture Notes in Computer Science, 2002.

[8] Z. Chen, Java Card technology for Smart Cards: architecture and programmer's guide, Java series, Addison-Wesley, Reading, MA, USA, 2000.

[9] The Coq proof assistant, http://coq.inria.fr.

[10] M. Éluard, T. Jensen, E. Denne, An operational semantics of the Java card firewall, in: I. Attali, T. Jensen (Eds.), Smart card programming and security: International Conference on Research in Smart Cards, E-smart 2001, Vol. 2140 of Lecture Notes in Computer Science, 2001.

[11] S. N. Freund, J. C. Mitchell, A formal framework for the Java Bytecode Language and Verifier, in: Proceedings of the 1999 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages & Applications, OOPSLA '99, Vol. 34(10) of SIGPLAN Notices, ACM Press, Denver, Colorado, USA, 1999, pp. 147–166.

[12] R. R. Hansen, Flow logic for Carmel, Tech. Rep. SECSAFE-IMM-001, DTU, http://www.doc.ic.ac.uk/˜siveroni/secsafe (Jun. 2002).

[13] P. H. Hartel, LETOS - a lightweight execution tool for operational semantics, Software - Practice and Experience 29 (15) (1999) 1379–1416.

[14] P. H. Hartel, M. J. Butler, M. Levy, The operational semantics of a Java secure processor, in: J. Alves-Foss (Ed.), Formal Syntax and Semantics of Java, Vol. 1523 of Lecture Notes in Computer Science, Springer-Verlag, Berlin Germany, 1999, pp. 313–352.

[15] P. H. Hartel, L. A. V. Moreau, Formalizing the safety of Java, the Java Virtual Machine and Java Card, ACM Computing Surveys 33 (4) (2001) 517–558.

[16] E. Hubbers, E. Poll, ESC/Java and JML specifications of the Java Card API 2.1.1, http://www.cs.kun.nl/˜erikpoll/publications/jc211_specs.html.

[17] J.-L. Lanet, A. Requet, Formal proof of smart card applets correctness, in: J.-J. Quisquater, B. Schneier (Eds.), Smart card research and applications: third international conference, CARDIS'98, Louvain-la-Neuve, Belgium, September 1998: proceedings, Louvain-la-Neuve, (be), 1998.

[18] T. Lindholm, F. Yellin, The Java Virtual Machine Specification, 2nd Edition, Addison-Wesley Publishing Co., Reading, MA, 1999.

[19] The LOOP project, http://www.cs.kun.nl/ bart/LOOP/.

[20] R. Marlet, Syntax of the JCVM language to be studied in the SecSafe project, Tech. Rep. SECSAFE-TL-005, v1.7, Trusted Logic, secSafe public (May 2001).

[21] H. Meijer, E. Poll, Towards a full formal specification of the JavaCard API, in: I. Attali, T. Jensen (Eds.), Smart card programming and security: International Conference on Research in Smart Cards, E-smart 2001, Vol. 2140 of Lecture Notes in Computer Science, 2001.

[22] T. Nipkow, L. C. Paulson, M. Wenzel, Isabelle/HOL — A Proof Assistant for Higher-Order Logic, Vol. 2283 of LNCS, Springer-Verlag, 2002.

[23] The PVS specification and verification system, http://pvs.csl.sri.com/.

[24] A. Requet, A B model for ensuring soundness of the Java card virtual machine, in: S. Gnesi, I. Schieferdecker, A. Rennoch (Eds.), 5th International ERCIM Workshop on Formal Methods for Industrial Critical Systems, FMICS 2000, GMD Report No. 91, Berlin, 2000.

[25] I. Siveroni, Formal specification of Java Card Runtime Environment and API, Technical Report SECSAFE-ICSTM-006, Imperial College London, http://www.doc.ic.ac.uk/˜siveroni/secsafe (2003).

[26] I. Siveroni, C. Hankin, A proposal of the JCVMLe operational semantics, Technical Report SECSAFE-ICSTM-001, Imperial College London, http://www.doc.ic.ac.uk/~siveroni/secsafe (2002).

[27] Sun Microsystems, Inc., Palo Alto, CA, USA, Java Card 2.1.1 Platform API Specification (May 2000).

[28] Sun Microsystems, Inc., Palo Alto, CA 94303 USA, The Java Card^{TM} 2.1.1 Runtime Environemnt (JCRE) Specification, revision 1.0 Edition (May 18 2000).

[29] Sun Microsystems, Inc., Palo Alto, CA 94303 USA, The Java Card^{TM} 2.1.1 Virtual Machine Specification, revision 1.0 Edition (May 18 2000).