Correctness of Java Card Method Lookup via Logical Relations

Ewen Denney and Thomas Jensen

Projet Lande, IRISA, Rennes Cedex 35042, France

Abstract. We formalise the Java Card bytecode optimisation from class file to CAP file format as a set of constraints between the two formats, and define and prove its correctness. Java Card bytecode is formalised as an abstract operational semantics, which can then be instantiated into the two formats. The optimisation is given as a logical relation such that the instantiated semantics are observably equal. The proof has been automated using the Coq theorem prover.

Using a high-level language for programming embedded systems may require a transformation phase in order that the compiled code fits on the device. In this paper we describe a method for formally proving the correctness of such a transformation. The method makes extensive use of types to describe the various run-time structures and relies on the notion of logical relation to relate the two representations of the code. We present the method in the setting of mapping Java onto smart cards. The Java Card language [10] is a trimmed down dialect of Java aimed at programming smart cards. As with Java, Java Card is compiled into bytecode, which is then verified and executed on a virtual machine [4], installed on a chip on the card itself. However, the memory and processor limitations of smart cards necessitate a further stage, in which the bytecode is optimised from the standard class file format of Java, to the *CAP file* format [11]. The core of this optimisation is a *tokenisation* in which names are replaced with tokens enabling a faster lookup of various entities.

We describe a semantic framework for proving the correctness of Java Card tokenisation. The basic idea is to give an abstract description of the constraints given in the official specification of the tokenisation and show that any transformation satisfying these constraints is ‘correct’. This is independent of showing that there actually exists a collection of functions satisfying these constraints. This article concentrates on proving the correctness of the specification. The formal development of an algorithm is the subject of another report. The main advantage of decoupling ‘correctness’ into two steps is that we get a more general result: rather than proving the correctness of one particular algorithm, we are able to show that the constraints described in Sun’s official specification [11] (given certain assumptions) are sufficient. We give a formalisation and correctness proof for the part concerned with dynamic method lookup. A comprehensive formalisation appears as a technical report [2].

1 The Conversion

Java source code is compiled on a class by class basis into the *class file* format. By contrast, Java Card *CAP files* correspond to packages. They are produced by the *con-*

version of a collection of class files. In the class file format, methods, fields and so on are referred to using strings. In CAP files, however, tokens are ascribed to the various entities. The idea is that if a method, say, is publically visible¹, then it is ascribed a token. If the method is only visible within its package, then it is referred to directly using an offset into the relevant data structure. Thus references are either internal or external. The conversion groups entities from different class files into the components of a CAP file. For example, all constant pools of the class files forming a package are merged into one constant pool component, and all method implementations are gathered in the same method component. One significant difference between the two formats is the way in which the method tables are arranged. In a class file, the methods item contains all the information relevant to methods defined in that class. In the CAP file, this information is shared between the class and method components. The method component contains the implementation details (*i.e.* the bytecode) for the methods defined in this package. The class component is a collection of class information structures. Each of these contains separate tables for the package and public methods, mapping method tokens to offsets into the method component. The method tables contain the information necessary for resolving any method call in that class.

The conversion is presented in [11] as a collection of constraints on the CAP file, rather than as an explicit mapping between class and CAP formats. For example, if a class inherits a method from a superclass then the conversion can choose to include the method token in the relevant method table or, instead, that the table of the superclass should be searched. There is a choice, therefore, between copying all inherited methods, or having a more compressed table. The specification does not constrain this choice. We adopt a simplified definition of the conversion, only considering classes, constant pools, and methods (with inheritance and overwriting). In particular, we ignore fields, exceptions and interfaces. The conversion also includes a number of mandatory optimisations such as the inlining of final fields, and the type-based specialisation of instructions [10,11], which we do not treat here.

2 Overview of Formalisation

The conversion from class file to CAP format is a transformation between formats of two virtual machines. The first issue to be addressed is determining in what sense, exactly, the conversion to token format should be regarded as an equivalence. We cannot simply say that the JVM and JCVM have the same behaviour for all bytecodes, in class and CAP file format respectively, because, *a priori*, the states of the virtual machines are themselves in different formats. Instead, we adopt a simple form of equivalence based on the notion of *representation independence* [5]. This is expressed in terms of so-called *observable* types. This limits us to comparing the two interpretations in terms of words, but this is sufficient to observe the operand stack and local variables, where the results of execution are stored.

Representation independence may be proven by defining *coupling relations* between the two formats that respect the tokenisation and are the identity at observable types.

¹ We follow the terminology of [11], where a method is *public visible* if it has either a `protected` or a `public` modifier, and *package visible* if it is declared `private` or has no visibility modifier.

This can be seen as formalising a *data refinement* from class to CAP file. We formalise the relations nondeterministically as any family of relations that satisfies certain constraints, rather than as explicit transformations. This is because there are many possible tokenisations and we wish to prove any reasonable optimisation correct.

The virtual machines are formalised in an operational style, as transition relations over abstract machines. We adopt the action semantics formalism of Mosses [6], using a mixture of operational and denotational styles: the virtual machines are formalised operationally, parameterised with respect to a number of auxiliary functions, which are then interpreted denotationally. This modular presentation of the semantics facilitates the comparison between the two formats. We illustrate this for dynamic method lookup, used in the semantics of the method invocation instructions. The lookup function which searches for the implementation of a method is dependent on the layout of the method tables. The operational rule giving the semantics of the method invocation instructions, presented in Section 5, is parameterised with respect to the lookup function. Then in Section 6 two possible interpretations of lookup are given.

In Section 4, we define abstract types for the various entities converted during tokenisation, which are common to the two formats. For example, `Class_ref` and `Environment`. It is this type structure which is used to define the logical relations. In Section 5 we give an operational semantics which is independent of the underlying class/CAP file format. The structure of the class/CAP file need not be visible to the operational semantics. We need only be able to extract certain data corresponding to a particular method, such as the appropriate constant pool. In Section 6, we give the specific details of the class and CAP file formats, defined as interpretations of types and auxiliary functions, $\llbracket \cdot \rrbracket_{name}$ and $\llbracket \cdot \rrbracket_{tok}$. We refer to these as the *name* and the *token* interpretation, respectively.

In Section 7, we define the logical relation, $\{R_\theta\}_{\theta \in Abstract_type}$. It is convenient to group the definition into several levels. First, there are various basic observable types (byte, short, *etc.*), γ , for which we have $R_\gamma = id_\gamma$. Second, there are the references, ι , such as package and class references, for which the relation R_ι represents the tokenisation of named items. Third, the constraints on the organisation into components (which we will call the *componentisation*) are expressed in R_κ , where κ includes method information structures, constant pools, and so on. This represents the relationship between components in CAP files and the corresponding entities in class files. Using the above three families of relations we can define R_θ for each type, θ , where

$$\theta ::= \gamma \mid \iota \mid \kappa \mid \theta \times \theta' \mid \theta \rightarrow \theta' \mid \theta + \theta' \mid \theta^*.$$

The family of relations, $\{R_\theta\}_{\theta \in Abstract_type}$, represents the overall construction of components in the CAP file format from a class file. The relations are ‘logical’ in the sense that the definitions for defined types follow automatically. For example, we define the type of the environment that contains the class hierarchy as

$$Environment = Package_ref \rightarrow Package.$$

and so the definition of $R_{Environment}$ follows from those of $R_{Package_ref}$, $R_{Package}$ and the standard construction of $R_{\rightarrow}$; similarly for R_{Heap} .

3 Related Work

There have been a number of formalisations of the Java Virtual Machine which have some relevance for our work here on Java Card. Bertelsen [1] gives an operational semantics which we have used as a starting point. He also considers the verification conditions, which considerably complicates the rules, however. Pusch has formalised the JVM in HOL [8]. Like us, she considers the class file to be well-formed so that the hypotheses of rules are just assignments. The operational semantics is presented directly as a formalisation in HOL, whereas we have chosen (equivalently) to use inference rules. All these works make various simplifications and abstractions. However, since these are formalisations of Java rather than Java Card they do not consider the CAP file format. In contrast, the work of Lanet and Requet [3] is specifically concerned with Java Card. They also aim to prove the correctness of Java Card tokenisation. Their work can be seen as complementing ours. They concentrate on optimisations, including the type specialisation of instructions, and do not consider the conversion as such. In contrast, we have specified the conversion but ignored the subsequent optimisations. Their formalism is based on the B method, so the specification and proof are presented as a series of refinements. In [7], Pusch proves the correctness of an implementation of Prolog on an abstract machine, the WAM. The proof structure is similar to ours, although there are refinements through several levels. There are operational semantics for each level, and correctness is expressed in terms of equivalence between levels. The differences between the semantics are significant, since they are not factored out into auxiliary functions as here. She uses a big-step operational semantics, which is not appropriate for us because we wish to compare intermediate results. Moreover, she uses an abstraction function on the initial state, the results being required to be identical, whereas we have a *relation* for both initial and final states.

4 Abstract Types

We use types to structure the transformation. These are not the types of the Java Card language, but rather are based on the simply-typed lambda calculus with sums, products and lists. We use record types with the actual types of fields (drawn from the official specification where not too confusing) serving as labels. Occasionally we use terms as singleton types, such as `0xFFFF` and `0`. There are two sorts of types: *abstract* and *concrete*. The idea is that abstract types are those we can think of independently of a particular format. The concrete types are the particular realisations of these, as well as types which only make sense in one particular model. For example, `CP_index` is the abstract type of indices into a constant pool for a given package. In the name interpretation, this is modelled by a class name and an index into the constant pool of the corresponding class file, *i.e.* `Class_name` $\times$ `Index` where `Index` is a concrete type. In the token interpretation, however, since all the constant pools are merged, we have $\llbracket \text{CP_index} \rrbracket_{\text{tok}} = \text{Package_tok} \times \text{Index}$. Another example is the various distinctions that are made between method and field references in CAP files, but not class files, and which are not relevant at the level of the operational semantics, which concerns terms of abstract types. We arrange the types so that as much as possible is common

between the two formats. For example, it is more convenient to uniformly define environments as mappings of type $\text{Package_ref} \rightarrow \text{Package}$, with Package interpreted as $\text{Class_name} \rightarrow \text{Class_file}$ or CAP_file .

There is a ‘type of types’ for the two forms of data type in Java Card — primitive types, *i.e.*, the simple types supported directly on the card, and reference types.

$$\begin{aligned}\text{Type} &= \{\text{Boolean}, \text{Byte}, \text{Short}\} + \text{Reference_type}. \\ \text{Reference_type} &= \text{Array_type} + \text{Class_ref}.\end{aligned}$$

We use a separate type, Object_ref , to refer to objects on the heap. The objects themselves contain a reference to the appropriate class or array of which they form an instance.

The type Word is an abstract unit of storage and is platform specific. All we need know is that object references and the basic types, Byte , Short and Boolean , can be stored in a Word . Rather than use an explicit coercion, we assume

$$\text{Word} = \text{Object_ref} + \text{Null} + \text{Boolean} + \text{Byte} + \text{Short}.$$

Thus a word is (*i.e.* represents) either a reference (possibly null) or an element of a primitive type. Furthermore, we define $\text{Value} = \text{Word}$. Although this is not strictly necessary, there is a conceptual distinction. If we were to introduce values of type int , then a value could be either a word or a double word.

There are several forms of references used during tokenisation, *viz.*, Package_ref , Class_ref and Method_ref . We distinguish Package from Package_ref , and similarly for the other items. Note that a *reference* is a composite entity which can be context dependent (*e.g.* in the CAP format a class reference can be in internal or external forms). We assume, however, that sufficient information is given so that references make sense globally. For example, class names are fully qualified, and class tokens are paired with a package token. We take field and method references to be to particular members of some class, and so contain a class reference. In contrast, an *identifier* is a name or a token (these are not used at the abstract level though). Using these basic types, we can then construct complex types using the usual type constructors: (non-dependent) sum, product, function and list types (denoted θ^*) as we did when defining the environment at the end of Sect. 2.

5 Operational Semantics

We define an operational semantics framework that allows us to model the execution of both class and CAP files. This is obtained by parameterising the semantics on a number of *auxiliary functions* that embody the differences between the two formats. This factorisation of the semantics reduces the equivalence proof considerably.

The official specification of the JCVM (and JVM) is given in terms of *frames*. A frame represents the state of the current method invocation, together with any other useful data. We introduce the notion of *configuration*, consisting of (the abstract syntax of) the code of the current method still to be executed, the operand stack, the local variables, and the current class reference. We write these as $\text{Config}(b, o, l, c)$ or just $\langle b, o, l, c \rangle$. To account for method invocations, we allow a configuration itself to be

considered as an instruction. When a method is invoked, the current instruction becomes a new configuration. Instead of a stack of frames, then, we have a single piece of ‘code’ (in this general sense). This form of closure is more general than the traditional idea of a call stack but helps simplify the proof. Method invocation is modelled by replacing the invoking instruction with a configuration that contains the code of the invoked method (see the detailed description of `invokevirtual` below). Execution of a method body is modelled by allowing transitions inside a configuration.

$$\frac{f \Rightarrow f'}{\langle \text{Config } f, ops, l, c \rangle \Rightarrow \langle \text{Config } f', ops, l, c \rangle}$$

The method invocation instructions (and others) take an argument which is an index into either the constant pool of a class file, or into the constant pool component of a CAP file. This means that the ‘concrete’ bytecode is itself dependent on the implementation and is therefore modelled by an abstract type. Formally, we define a transition relation

$$\Rightarrow \subseteq \text{Config} \times \text{Arrow} \times \text{Config}$$

using the types

$$\begin{aligned} \text{Config} &= \text{Bytecode} \times \text{Word}^* \times \text{Locals} \times \text{Class_ref} \\ \text{Arrow} &= \text{Global_state} \rightarrow \text{Global_state} \\ \text{Global_state} &= \text{Environment} \times \text{Heap} \\ \text{Bytecode} &= \text{Instruction} + (\text{Bytecode} \times \text{Bytecode}) + \text{Config} \end{aligned}$$

As mentioned above, the structure of the class/CAP file need not be visible to the operational semantics. We use a number of auxiliary functions, some of which have preconditions that we take as concomitant with the well-formedness of the class file. The definition of method invocation uses the lookup function

$$\text{lookup} : \text{Class_ref} \times \text{Method_ref} \rightarrow \text{Class_ref} \times \text{Bytecode}$$

that takes the actual class reference, together with the declared method reference (which contains the class where the method is declared), and returns the class reference where the method is defined together with the code. Function `method_nargs : Method_ref → Nat` returns the number of arguments for a given method reference. The instruction for virtual method invocation is evaluated as follows:

1. The two byte index, i , into the constant pool is resolved to get the declared method reference containing the declared class reference and a method identifier (either a signature or token).
2. The number of arguments to the method is calculated.
3. The object reference, r , is popped off the operand stack.
4. Using the heap, we get $\text{heap}(r) = \langle \text{act_cref}, _ \rangle$, the actual class reference (fully qualified name or a package/class token pair).
5. We then do $\text{lookup}(\text{act_cref}, \text{dec_mref})$, getting the class where the method is implemented, and its bytecode. The lookup function is used with respect to the class hierarchy (environment).
6. A configuration is created for this method and evaluation proceeds from there.

$dec_mref := \text{constant_pool}(c)(i)$	get declared method reference
$n := \text{method_nargs}(stat_mref)$	get number of arguments
$\langle act_cref, - \rangle := \text{heap}(r)$	get actual class reference from heap
$\langle m_cl, m_cd \rangle := \text{lookup}(act_cref, dec_mref)$	look up method
<hr/>	
$\langle \text{invokevirtual } i, a_1 \dots a_n :: r :: s, l, c \rangle \Rightarrow \langle \langle m_cd, \langle \rangle, a_1 \dots a_n :: r, m_cl \rangle, s, l, c \rangle$	

In the following sections we show how to instantiate the semantic framework (in particular the lookup function) to obtain a class file and a CAP file semantics.

6 Interpretations

The name interpretation gives semantics using Java class files (see Figure 1). Since this is fairly standard we give a brief description. Classes are described by fully qualified names, whereas methods and fields are given signatures, consisting of an unqualified name and a type, together with the class of definition. We assume a function *pack_name* which gives the package name of a class name. The data is arranged into class files, each of which contains all the information corresponding to a particular class. We only give the interpretation of those parts used here. We group the class files by package into a global environment so *env_name(p)(c)* denotes the class file in package *p* with name *c*.

$$\begin{aligned}
\llbracket \text{Package} \rrbracket_{name} &= \text{Class_name} \rightarrow \text{Class_file} \\
\llbracket \text{Class_ref} \rrbracket_{name} &= \text{Class_name} \\
\llbracket \text{Method_ref} \rrbracket_{name} &= \text{Class_name} \times \text{Sig} \\
\text{Sig} &= \text{Method_name} \times \llbracket \text{Type} \rrbracket_{name}^* \\
\llbracket \text{Class} \rrbracket_{name} &= \text{Class_file} \\
\text{Class_file} &= \text{Class_flags} \times \text{Super} \times \text{Methods_item} \times \text{Constant_pool_item} \times \text{Class_name} \\
\text{Super} &= \text{Class_name} + \text{Void} \\
\llbracket \text{Pack_methods} \rrbracket_{name} &= \text{Class_name} \rightarrow \text{Methods_item} \\
\text{Methods_item} &= \text{Sig} \rightarrow \text{Method_info} \\
\text{Method_info} &= \\
\text{Method_flags} \times \text{Sig} \times (\llbracket \text{Type} \rrbracket_{name} + \text{Void}) \times \text{Maxstack} \times \text{Maxlocals} \times \text{Bytecode}
\end{aligned}$$

Fig. 1. Name Interpretation

Method signatures are not considered to include the return type. We assume that the signature in the result of a methods item is the same as the argument.

There are a number of possibilities for how method lookup should be defined, depending on the definition of inheritance. For example, [1,8] use a ‘naive’ lookup which does not take account of visibility modifiers. A fuller discussion of this appears in [9].

In the JCVM, data is arranged by packages into CAP files. Each CAP file consists of a number of components, but not all are used for method lookup (or, indeed, the rest

```

lookup_name (act_class, (sig, dec_class)) =
let dec_pk  = pack_name(dec_class)
    act_pk  = pack_name(act_class)
    (_,_,meth_dec,_,_) = env_name (dec_pk) (dec_class)
    (_,super,_,meth_act,_,_) = env_name (act_pk) (act_class)
    (dec_flags,_,_,_,_) = meth_dec(sig) in
if meth_act(sig) = undefined
then lookup_name(super, (sig, dec_class))
else if
    dec_flags(protected) or dec_flags(public) or act_pk = dec_pk
then let (_,_,_,_,_,code) = meth_act(sig) in (act_class,code)
else lookup_name(super, (sig, dec_class))

```

Fig. 2. The lookup function for the class file format.

of the operational semantics). We just include those components we need here, namely, the constant pool, class and method components.

References to items external to a package are via tokens — for packages, classes, and virtual methods — each with a particular range and scope. These are then used to find internal offsets into the components. For example, a class reference is either an internal offset into the class component of the CAP file of the class’ package, or an external reference composed of a package token and a class token. However, since we need to relate the reference to class names, we will assume that all references come with package information, even though this is superfluous in the case of internal references.

$$\begin{aligned}
 \llbracket \text{Package} \rrbracket_{tok} &= \text{CAP_file} \\
 \text{CAP_file} &= \text{Constant_pool_comp} \times \text{Class_comp} \times \text{Method_comp} \\
 \llbracket \text{Package_ref} \rrbracket_{tok} &= \text{Package_tok} \\
 \llbracket \text{Class_ref} \rrbracket_{tok} &= \text{Package_tok} \times (\text{Class_tok} + \text{Offset}) \\
 \llbracket \text{Method_ref} \rrbracket_{tok} &= \llbracket \text{Class_ref} \rrbracket_{tok} \times \text{Virtual_method_tok} \\
 \llbracket \text{Class} \rrbracket_{tok} &= \text{Class_info} \\
 \text{Class_comp} &= \text{Offset} \rightarrow \text{Class_info} \\
 \text{Class_info} &= \text{Class_flags} \times \text{Super} \times \text{Public_table} \times \text{Package_table} \times \text{Class_ref} \\
 \text{Public_table} &= \text{Public_base} \times \text{Public_size} \times (\text{Index} \rightarrow \text{Offset} + \{0\text{x}\text{FFFF}\}) \\
 \text{Package_table} &= \text{Package_base} \times \text{Package_size} \times (\text{Index} \rightarrow \text{Offset}) \\
 \llbracket \text{Pack_methods} \rrbracket_{tok} &= \text{Method_comp} \\
 \text{Method_comp} &= \text{Offset} \rightarrow \text{Method_info} \\
 \text{Method_info} &= \text{Method_flags} \times \text{Maxstack} \times \text{Nargs} \times \text{Max_locals} \times \text{Bytecode}
 \end{aligned}$$

Fig. 3. Token Interpretation

The class component consists of a list of class information structures, each of which has method tables, giving offsets into the method component, where the method implementations are found. The lookup algorithm uses tokens to calculate the corresponding method table index. There are separate tables for public and package methods. Method access information is given implicitly by the tokens rather than by flags. The two method tables each contain a base, size and ‘list’ of entries. The entries are defined from the *base* to *base + size - 1* inclusive. The entry for a public method will be 0xFFFF if the method is defined in another package.

For a given class reference, the function `class_info` finds the corresponding class information structure in the global environment. The variant, `class_info'` returns the class information structure in a particular CAP file. The function `method_array` simply finds the method component for a given class reference. We assume the existence of functions `class_offset` and `method_offset` for resolving external tokens to internal offsets. It follows from the definition of the abstract type `Environment`, that the environment in the token format consists of a mapping from package tokens to their corresponding CAP file *i.e.*, $env_{tok} : \text{Package_tok} \rightarrow \text{CAP_file}$. The lookup function takes a class reference (the declared class), a method reference (in the actual class), and returns the reference to the class where the code is defined, together with the bytecode itself. The main steps of the algorithm (see Fig. 4) are:

1. Get method array for the package of the actual class.
2. Get class information for the actual class.
3. If public: if defined then *get info* else *lookup super*.
If package: if defined and visible then *get info* else *lookup super*.

7 Formalisation of Equivalence

We formalise the equivalence between the class and CAP formats as a family of relations, $\{R_\theta : [\![\theta]\!]_{name} \leftrightarrow [\![\theta]\!]_{tok}\}_{\theta \in \text{Abstract_type}}$ indexed by abstract type, θ . The idea is that $x R_\theta y$ when y is a *possible* transformation of x . The relations are not necessarily total, *i.e.* for some $x : [\![\theta]\!]_{name}$, there may not be a y such that $x R_\theta y$. Formally, the relations are defined as a mutually inductive collection of constraints, R_θ , for each type θ , where the types, θ , are given by the grammar:

$$\begin{aligned} \gamma &::= \text{Bool} \mid \text{Nat} \mid \text{Object_ref} \mid \text{Boolean} \mid \text{Byte} \mid \text{Short} \mid \text{Value} \mid \text{Word} \\ \iota &::= \text{Package_ref} \mid \text{Ext_class_ref} \mid \text{Class_ref} \mid \text{Method_ref} \\ \kappa &::= \text{CP_index} \mid \text{CP_info} \mid \text{Method_info} \mid \text{Package} \mid \text{Class} \mid \\ &\quad \text{Constant_pool} \mid \text{Pack_methods} \\ \theta &::= \gamma \mid \iota \mid \kappa \mid \theta \times \theta' \mid \theta \rightarrow \theta' \mid \theta + \theta' \mid \theta^* \end{aligned}$$

where the observable types are built up inductively from the γ , *i.e.* do not contain the ι and κ . There are two sources of underspecification. First, the relations really can be non-functional. Second, there is a choice for what some of the relations are. For example, $R_{\text{Class_ref}}$ is *some* bijection satisfying certain constraints. The relations between the ‘large’ structures, however, are completely defined in terms of those between smaller ones. There are two parts to the transformation itself: the tokenisation, defined as the relations R_ι , and the ‘componentisation’, defined as the R_κ .

```

lookup_tok (act_class_ref, (dec_class_ref, method_tok)) =

let methods = method_array (act_class_ref)
  (_, super, (public_base, _, public_table),
    (package_base, _, package_table), _) : Class_info =
  class_info(act_class_ref) in
if method_tok div 128 = 0 then /* public */
  if method_tok >= public_base then
    let method_offset = public_table[method_tok-public_base] in
    if method_offset <> 0xFFFF
    then (act_class_ref, methods[method_offset].Bytecode)
    else /* look in superclass */
      lookup_tok(super, (dec_class_ref, method_tok))
  else /* look in superclass */
    lookup_tok(super, (dec_class_ref, method_tok))
else /* package */
  if method_tok >= package_base /\
    same_package(dec_class_ref, act_class_ref)
  then let method_offset =
    package_table[method_tok mod 128 - package_base]
    in (act_class_ref, methods[method_offset].Bytecode)
  else /* look in superclass */
    lookup_tok(super, (dec_class_ref, method_tok))

```

Fig. 4. The lookup function for the CAP file format.

7.1 Tokenisation

The relations, R_L , represent the tokenisation of items. The general idea is to set up relations between the names and tokens assigned to the various entities, subject to certain constraints described in the specification.

In order to account for token scope, we relate names to tokens paired with the appropriate context information. For example, method tokens are scoped within a class, so the relation $R_{\text{Method_ref}}$ is between pairs of class names and signatures, and pairs of class references and method tokens. We must add a condition, therefore, to ensure that the package token corresponds to the package name of this class name.

We assume that each of these relations is a bijection, modulo the equivalence between internal and external references (with one exception to account for the copying of virtual methods, explained below). Formally,

$$a R b \wedge a' R b \Rightarrow a = a'$$

$$a R b \Rightarrow (a R b' \iff \text{Equiv}(b, b'))$$

where equivalence, Equiv , of class references is defined as the reflexive symmetric closure of:

$$\text{Equiv}(\langle p_tok, offset \rangle, \langle p_tok, c_tok \rangle) \iff \text{class_offset}(p_tok, c_tok) = offset$$

The second condition contains two parts: that the relation is functional modulo *Equiv*, and that it is closed under *Equiv*. We say that R is an *external bijection* when these conditions hold. We extend the definition of *Equiv* and external bijection to the other references.

These relations are defined with respect to the environment (in name format). We use a number of abbreviations for extracting information from the environment. We write $c < c'$ for the subclass relation (*i.e.* the transitive closure of the direct subclass relation) and $\leq$ for its reflexive closure. In the token interpretation this is modulo *Equiv*. We write $m_tok \in c_ref$ when a method with token m_tok is declared in the class with reference c_ref , and $pack_name(c)$ for the package name of the class named c .

We define function *Class_flag* for checking the presence of attributes such as *public*, *final*, *etc.* The tokenisation uses the notion of *external visibility*.

$$\text{Externally_visible}(c_name) = \text{Class_flag}(c_name, \text{Public})$$

We will also write $public(sig)$ and $package(sig)$ according to the visibility of a method.

Package_ref : As mentioned above, we take package tokens to be externally visible. The relation $R_{\text{Package_ref}}$ is simply defined as any bijection between package names and tokens.

Ext_class_ref : In order to define the relation for class references we first define the relation for external class references. We define $R_{\text{Ext_class_ref}}$ as a bijection between class names and external class references such that:

$$c_name R_{\text{Ext_class_ref}} (p_tok, c_tok) \Rightarrow$$

$$\text{Externally_visible}(c_name) \wedge pack_name(c_name) R_{\text{Package_ref}} p_tok$$

Method_ref : This is not a bijection because of the possibility of copying. Although ‘from names to tokens’ we do have:

$$\begin{aligned} \langle c_name, sig \rangle R_{\text{Method_ref}} \langle c_ref, m_tok \rangle \wedge \\ \langle c'_name, sig' \rangle R_{\text{Method_ref}} \langle c_ref, m_tok \rangle \Rightarrow \begin{cases} c_name = c'_name \wedge \\ sig = sig' \end{cases} \end{aligned}$$

for a converse we have:

$$\begin{aligned} \langle c_name, sig \rangle R_{\text{Method_ref}} \langle c_ref, m_tok \rangle \wedge \\ \langle c_name, sig \rangle R_{\text{Method_ref}} \langle c'_ref, m'_tok \rangle \Rightarrow \begin{cases} (c_ref \leq c'_ref \vee c'_ref \leq c_ref) \\ \wedge m_tok = m'_tok \end{cases} \end{aligned}$$

The first condition says that if a method overrides a method implemented in a superclass, then it gets the same token. Restrictions on the language mean that overriding cannot change the method modifier from *public* to *package* or vice versa.

$$\left. \begin{aligned} \langle c_name, sig \rangle R_{\text{Method_ref}} \langle c_ref, m_tok \rangle \wedge \\ \langle c'_name, sig \rangle R_{\text{Method_ref}} \langle c'_ref, m'_tok \rangle \wedge \\ c'_name < c_name \wedge \\ (package(sig) \Rightarrow same_package(c_name, c'_name)) \end{aligned} \right\} \Rightarrow m_tok = m'_tok$$

The second condition says that the tokens for *public* introduced methods must have higher token numbers than those in the superclass. We assume a predicate, *new_method*, which holds of a method signature and class name when the method is defined in the class, but not in any superclass.

$$\begin{aligned} public(sig) \wedge new_method(sig, c_name) \wedge \\ (c_name, sig) R_{\text{Method_ref}} (c_ref, m_tok) \Rightarrow \begin{cases} \forall m'_tok \in super(c_ref). \\ m_tok > m'_tok \end{cases} \end{aligned}$$

Package-visible tokens for introduced methods are similarly numbered, if the superclass is in the same package.

$$\begin{aligned} & package(sig) \wedge new_method(sig, c_name) \wedge \\ & (c_name, sig) R_{Method_ref} (c_ref, m_tok) \wedge \quad \Rightarrow \quad \begin{cases} \forall m'_tok \in super(c_ref). \\ m_tok > m'_tok \end{cases} \\ & same_package(c_name, super(c_name)) \end{aligned}$$

The third condition says that public tokens are in the range 0 to 127, and package tokens in the range 128 to 255.

$$\begin{aligned} & \langle c_name, sig \rangle R_{Method_ref} \langle c_ref, m_tok \rangle \Rightarrow \\ & (public(sig) \Rightarrow 0 \leq m_tok \leq 127) \wedge (package(sig) \Rightarrow 128 \leq m_tok \leq 255) \end{aligned}$$

The specification [11] also says that tokens must be contiguously numbered starting at 0 but we will not enforce this.

7.2 Componentisation

The relations in the previous section formalise the correspondence between named and tokenised entities. When creating the CAP file components, all the entities are converted, including the package visible ones. Thus at this point we define R_{Class_ref} as a relation between named items and either external tokens or internal references, subject to coherence constraints.

We must ensure that if a name corresponds to both an external token and to an internal offset, then the token and the offset correspond to the same entity. We ensure this by using the offset function $class_offset : Package_tok \times Class_tok \rightarrow Offset$ which returns the internal offset corresponding to an external token, and then *define* R_{Class_ref} from this and $R_{Ext_class_ref}$. Clearly, therefore, R_{Class_ref} is not a bijection.

Class_ref : We define R_{Class_ref} as an external bijection which respects $R_{Ext_class_ref}$, that is, such that

$$c_name R_{Class_ref} (p_tok, c_tok) \iff c_name R_{Ext_class_ref} (p_tok, c_tok).$$

Thus R_{Class_ref} extends $R_{Ext_class_ref}$ to internal references.

Method_info: We only treat certain parts of the method information here:

$$\begin{aligned} & \langle flags, sig, -, -, maxstack, maxlocals, code, - \rangle \quad \begin{matrix} flags R_{Method_flags} flags' \wedge \\ maxstack = maxstack' \wedge \end{matrix} \\ & \quad \quad \quad R_{Method_info} \quad \iff \quad \begin{matrix} size(sig) = nargs' \wedge \\ maxlocals = maxlocals' \wedge \\ code R_{Bytecode} code' \end{matrix} \\ & \langle flags', maxstack', nargs', maxlocals', code' \rangle \end{aligned}$$

In the name interpretation, information is grouped by the package and so, for example, $\llbracket Pack_methods \rrbracket_{name} : Class_name \rightarrow Methods_item$ is the ‘set’ of method data for all classes. In the token format the method information is spread between the two components. The coupling relations reflect this: the relation R_{Class} ensures that a named method corresponds to a particular offset, and $R_{Pack_methods}$ ensures that the entry at this offset is related by R_{Method_info} .

Pack_methods: The method item and method component contain the implementations of both static and virtual methods.

$$\begin{aligned} \text{methods_name} & \quad \forall \langle c_name, sig \rangle R_{\text{Method_ref}} \langle p_tok, c_tok, m_tok \rangle. \\ R_{\text{Pack_methods}} & \quad \Leftrightarrow \text{methods_name}(c_name, sig) R_{\text{Method_info}} \\ \text{method_comp} & \quad \text{methods_comp.methods}(\text{method_offset}(p_tok, c_tok, m_tok)) \end{aligned}$$

Class: We define R_{Class} . There are a number of equivalences expressing correctness of the construction of the class component. For the lookup, the significant ones are those between the method tables. These say that if a method is defined in the name format, then it must be defined (and equivalent) in the token format. Since the converse is not required, this means we can copy method tokens from a superclass. Instead, there is a condition saying that if there is a method token, then there must be a corresponding signature in some superclass.

If a method is visible in a class, then there must be an entry in the method table, indicating how to find the method information structure in the appropriate method component. For package visible methods this implies that the method must be in the same package. For public methods, if the two classes are in the same package, then this entry is an offset into the method component of this package. Otherwise, the entry is $0xFFFF$, indicating that we must use the method token to look in another package.

The class component only contains part of the information contained in the class files. The full definition is given in Figure 5. (writing c_name for $cf.Class_name$ and c_ref for $ci.Class_ref$): The offset functions link the various relations. We make a global assumption (in fact, local to an environment) of the existence of class_offset and method_offset .

Equivalence proof: The full proof establishes that the auxiliary functions preserve the appropriate relations [2]. Here, we state the main lemma for the function lookup whose type is $\text{Class_ref} \times \text{Method_ref} \rightarrow \text{Class_ref} \times \text{Bytecode}$.

Lemma 1. *If the heap and environment are related in the two formats, then:*

$$\llbracket \text{lookup} \rrbracket_{\text{name}} R_{\text{Class_ref} \times \text{Method_ref} \rightarrow \text{Class_ref} \times \text{Bytecode}} \llbracket \text{lookup} \rrbracket_{\text{tok}}$$

In order to use the operational semantics with the logical relations approach it is convenient to view the operational semantics as giving an interpretation. We define $\llbracket \text{code} \rrbracket(\langle env, heap, op_stack, loc_vars, m_ref \rangle)$ as the resulting state from the (unique) transition from $\langle \text{code}, op_stack, loc_vars \rangle$ with environment env and heap $heap$. Thus interpreted bytecode has type $\text{State} \rightarrow \text{Bytecode} \times \text{State}$ where State is

$$\text{State} = \text{Global_state} \times \text{Operand_stack} \times \text{Local_variables} \times \text{Class_ref}$$

Now, the following fact is trivial to show: if $R_B = id_B$ for all basic observable types, then $R_\theta = id_\theta$ for all observable θ . In combination with the following theorem, then, this says that if a transformation satisfies certain constraints (formally expressed by saying that it is contained in R) then it is correct, in the sense that no difference can be observed in the two semantics. In particular, we can observe the operand stack (of observable type Word^*) and the local variables (of observable type $\text{Nat} \rightarrow \text{Word}$) so these are identical under the two formats.

$$\begin{aligned}
& cf : \text{Class_file} \ R_{\text{Class}} \ ci : \text{Class_info} \iff \\
& \left\{ \begin{array}{l}
cf.\text{Class_flags} \ R_{\text{Class_flags}} \ ci.\text{Class_flags} \wedge \\
cf.\text{Super} \ R_{\text{Class_ref}} \ ci.\text{Super} \wedge \\
\\
\forall sig \in cf.\text{Methods_item}. \\
public(sig) \Rightarrow \\
\exists m_tok . ci.\text{Public_base} \leq m_tok < ci.\text{Public_base} + ci.\text{Public_size} \wedge \\
\langle c_name, sig \rangle \ R_{\text{Method_ref}} \langle c_ref, m_tok \rangle \wedge \\
ci.\text{Public_table}[m_tok - ci.\text{Public_base}] = \text{method_offset}(c_ref, m_tok) \\
\wedge \\
package(sig) \Rightarrow \\
\exists m_tok . ci.\text{Package_base} \leq m_tok \ \& \ 127 < ci.\text{Package_base} + ci.\text{Package_size} \wedge \\
\langle c_name, sig \rangle \ R_{\text{Method_ref}} \langle c_ref, m_tok \rangle \wedge \\
ci.\text{Package_table}[m_tok \ \& \ 127 - ci.\text{Package_base}] = \text{method_offset}(c_ref, m_tok) \\
\wedge \\
\forall m_tok \in ci.\text{Public_table} \cup ci.\text{Package_table}. \exists sig. \exists c_name. \\
\langle c_name, sig \rangle \ R_{\text{Method_ref}} \langle c_ref, m_tok \rangle \wedge c_name \leq c_name \wedge \\
public(sig) \Rightarrow [(same_package(c_name, c_name) \iff \\
ci.\text{Public_table}[m_tok - ci.\text{Public_base}] \neq 0xFFFF)]
\end{array} \right.
\end{aligned}$$

Fig. 5. Definition of R_{Class}

Theorem 1. Assume that $env_{name} \ R_{\text{Environment}} \ env_{tok}$, $heap_{name} \ R_{\text{Heap}} \ heap_{tok}$, $ls \ R_{\text{Local_state}} \ ls'$, and code $R_{\text{Bytecode}} \ code'$. Then

$$[[code]]_{name}(env_{name}, heap_{name}, ls) \ R_{\text{Bytecode} \times \text{State}} \ [[code']]_{tok}(env_{tok}, heap_{tok}, ls')$$

8 Conclusion

We have formalised the virtual machines and file formats for Java and Java Card, and the optimisation as a relation between the two. Correctness of this optimisation was expressed in terms of observable equivalence of the operational semantics, and this was deduced from the constraints that define the optimisation. Although the framework we have presented is quite general, the proof is specific to the instantiations of auxiliary functions we chose. It could be argued that we might have proven the equivalence of two incorrect implementations of lookup. The remedy for this would be to specify the functions themselves, and independently prove their correctness. Furthermore, we have made a number of simplifications which could be relaxed. We have used a simple definition of R_{Bytecode} here, which just accounts for the changing indexes into constant pools (as well as method references in configurations). We have not considered inlining or the specialisation of instructions, however. We expressed equivalence in terms of an identity at observable types but we should also account for the difference in word size, as in [3]. Although the specialisation of instructions could be handled by our technique, the extension is less clear for the more non-local optimisations.

We emphasise that the particular form of operational semantics used here is orthogonal to the rest of the proof. This version suffices for the instructions considered here, but

could easily be changed (along with the definition of R_{Bytecode}). The auxiliary functions could be given different definitions; for example, an abstract interpretation or, going in the opposite direction, including error information.

The definitions have been formalised in Coq, and the lemmas verified [9]. The discipline this imposed on the work presented here was very helpful in revealing errors. Even just getting the definitions to type-check uncovered many errors. We take the complexity of the proofs (in Coq) as evidence for the merit in separating the correctness of a particular algorithm from the correctness of the specification. In fact, the operational semantics, correctness of the specification, and development of the algorithm are all largely independent of each other.

As mentioned in the introduction, there are two main steps to showing correctness:

1. Give an abstract characterisation of all possible transformations and show that the abstract properties guarantee correctness.
2. Show that an algorithm implementing such a transformation exists.

We are currently working on a formal development of a tokenisation algorithm using Coq's program extraction mechanism together with constraint-solving tactics.

Acknowledgments

This work was partially supported by the INRIA *Action de recherche coopérative* Java Card

References

1. P. Bertelsen. Semantics of Java byte code. Technical report, Department of Information Technology, Technical University of Denmark, March 1997.
2. Ewen Denney. Correctness of Java Card Tokenisation. Technical Report 1286, Projet Lande, IRISA, 1999. Also appears as INRIA research report 3831.
3. Jean-Louis Lanet and Antoine Requet. Formal proof of smart card applets correctness. In *Third Smart Card Research and Advanced Application Conference (CARDIS'98)*, 1998.
4. T. Lindholm and F. Yelling. *The Java Virtual Machine Specification*. Addison-Wesley, 1997.
5. J. Mitchell. *Foundations for Programming Languages*. Foundations of Computing Series. MIT Press, 1996.
6. Peter D. Mosses. Modularity in structural operational semantics. Extended abstract, November 1998.
7. Cornelia Pusch. Verification of Compiler Correctness for the WAM. In J. von Wright, J. Grundy, and J. Harrison, editors, *Theorem Proving in Higher Order Logics (TPHOLs'96)*, pages 347–362. Springer-Verlag, 1996.
8. Cornelia Pusch. Formalizing the Java Virtual Machine in Isabelle/HOL. Technical Report TUM-19816, Institut für Informatik, Technische Universität München, 1998.
9. Gaëlle Segouat. Preuve en Coq d'une mise en oeuvre de Java Card. Master's thesis, Projet Lande, IRISA, 1999.
10. Sun Microsystems. *Java Card 2.0 Language Subset and Virtual Machine Specification*, October 1997. Final Revision.
11. Sun Microsystems. *Java Card 2.1 Virtual Machine Specification*, March 1999. Final Revision 1.0.