

Formal specification and static checking of Gemplus' electronic purse using ESC/Java

Néstor Cataño and Marieke Huisman

INRIA Sophia-Antipolis, France
{Nestor.Catano, Marieke.Huisman}@sophia.inria.fr

Abstract. This paper presents a case study in formal specification of smart card programs, using ESC/Java. It discusses an electronic purse application, provided by Gemplus, that we have annotated with functional specifications (*i.e.* pre- and postconditions, modifies clauses and class invariants), that are as detailed as possible. The specification is based on the informal documentation of the application. Using ESC/Java, the implementation has been checked *w.r.t.* the specification. This revealed several errors or possibilities for improvement in the source code (*e.g.* removing unnecessary tests).

Our paper shows that a relatively lightweight use of formal specification techniques can already have a serious impact on the quality of a program and its documentation. Furthermore, we also present some ideas on how ESC/Java could be further improved, both *w.r.t.* specification and verification.

Keywords: static checking, specification, ESC/Java, Java, smart cards.

1 Introduction

Background When developing a large software application, a significant part of the work is spent on writing clear and concise documentation. This documentation serves several purposes. First of all, it helps the developers of the application to do maintenance, as the documentation helps to understand the implementation decisions taken by a colleague, but also to understand ones own decisions after a certain period of time. Further, software documentation also is useful when somebody else builds a new application, using features provided by the application at hand.

However, such program documentation is only useful if it correctly describes the implementation, thus one would like to have some trust in its appropriateness. A way to achieve this is to write a formal specification, *i.e.* a description of the program behaviour in logic, and then prove the correctness of the implementation *w.r.t.* this specification, but this is difficult (as it requires a good understanding of the semantics underlying the specification and programming language), and labour-intensive (see *e.g.* [9, 2] for examples of full program verification). Thus, although it is feasible to do formal specification and verification, the benefits in general do not outweigh the costs.

Recently, several projects have started developments to overcome these problems. First of all, to encourage application developers to write formal specifications, specifications are written in a language that is close to the language in which the specified programs are written: the specification languages reuse the expression syntax of the programming language. The assertions that can be written in an Eiffel program [15] are a first example of such a specification style, and recently several annotation languages for Java have been proposed, following the same strategy: JML [11], ESC/Java [6], and the Jass annotation language [10]. For JML and ESC/Java, effort has been put into making these specification languages converge [5], so that the respective tools can be used for both languages. Typical for these languages is that expressions are written as Java expressions, extended with some specification-specific constructs.

Secondly, together with the ESC/Java language, a static checker has been developed [6], which can be used to check automatically simple, but useful properties. This static checker tries to check that a program satisfies its annotations, by using a dedicated, automatic theorem prover. This automatic theorem prover has been fine-tuned to find common programming problems like `NullPointerException`, and `ArrayIndexOutOfBoundsException`, but it also can be used to check other annotations. If the theorem prover cannot establish that a certain specification is satisfied, ESC/Java issues a warning. Such a warning does not necessarily mean that the program is wrong, as the ESC/Java approach is neither sound, nor complete. When designing the tool, a compromise has been made between soundness, completeness, and efficiency. The result is an efficient, automatic checker, that can increase the confidence in the correctness of programs, and that finds many common programming errors. However, if one wishes to establish formally the correctness of a complicated algorithm, other (possibly interactive) verification techniques have to be used, as advocated in *e.g.* the LOOP project [13] or the Jive project [16]. But even for such complex algorithms it pays to use ESC/Java first, in order to find quickly and automatically a first approximation of the errors in the algorithm and/or specification, before diving into the complete formal verification.

This paper To demonstrate the usefulness of this approach, this paper describes the ESC/Java annotation of a smart card application. This case study shows that annotating programs with ESC/Java specifications can be helpful to create quickly clear, concise and unambiguous documentation for a software application, which is in correspondence with the implementation. The original source code of the case study – which implements an electronic purse – comes from Gemplus [8]. In this paper we discuss the annotations of the source code, and several possibilities for improvement that we encountered. The result of this work does not give a fully verified specification, but it gives a reasonable description of the electronic purse implementation, which could serve as a basis for further formal verification, *e.g.* by using the LOOP compiler.

The main contribution of this paper is that it shows that by making light-weight use of formal verification techniques, as provided by ESC/Java, it is

feasible (i) to write a formal specification of an application, and (ii) to have the implementation checked *w.r.t.* the specification so as to increase confidence in the correctness of the implementation. When specifying the purse we have found several (simple) properties which are informally documented, but are not satisfied by the implementation. It is straightforward to formally specify these properties and ESC/Java immediately finds the places where these properties are not preserved in the implementation.

Furthermore, to the best of our knowledge this case study is one of the first larger case studies using ESC/Java, and we found several points for improvement in the static checker and its specification language. This leads us to a wish list on improvements in the specification language and to the development of a checker for so-called modifies clauses. This checker will be described in a separate paper [3].

The rest of this paper is organised as follows. Section 2 describes the general outline of the case study. Section 3 gives a brief introduction into the static checker ESC/Java. Section 4 describes the annotations of the purse in general, and discusses several aspects of the specification in more detail. Section 5 comments on the use of ESC/Java and gives suggestions for improvement. Finally, Section 6 gives conclusions and presents future work.

2 General outline of the Electronic Purse

The electronic purse is a JavaCard application, published as an advanced smart card programming case study by Gemplus [8]. A JavaCard smart card is capable of running programs developed in JavaCard [18], a dialect of standard Java. JavaCard does not provide concepts such as dynamic class loading, security management, multi-threading and synchronisation, object cloning and large primitive data types (float, double, long and char). JavaCard applications are called *applets*. The electronic purse applet provides the ability to perform banking operations to the card holder. Typical operations are credit, debit, and currency changes.

The *debit* operation. Debit operations which involve an amount greater than `maxDebitWOPIN` are protected by a pin code. During one session, a user can do several of these transactions by presenting his pin code only once. To protect the card against attacks, the number of transactions that can be performed without presenting a pin code is limited to `maxTransactionWOPIN`.

The *credit* operation. If the balance on the card is not sufficient to execute a certain debit operation, the balance can be increased by performing a credit operation. To do this, the point of sale terminal asks the bank of the card holder for credit permission. If the permission is obtained, the account is credited and a confirmation is sent to the bank.

The *currency change* operation. The balance of the purse is expressed in a certain currency. When the card holder travels, the current currency can be

changed. In order to do this, the terminal requests a new exchange rate and a certificate from the bank. The purse verifies that the bank is really the expected bank and validates the exchange rate. After changing the balance value, the purse must modify all variables related to the currency.

The purse applet interacts with so-called *loyalty applets* (implementing *e.g.* a frequent flyer program) that may be present on the card. Within the loyalty applet, the card holder gets loyalty points having made certain purchases, and these points can be used later to make other purchases. Further, a *card issuer applet* should be available on the card, which can initialise the purse. Finally, the purse applet also communicates with the *point of sale terminal*.

The purse application consists of three packages: **utils**, **purse** and **pacap-interfaces**. The **utils** package implements basic classes such as **Annee** (year), **Mois** (month) **Jour** (day) and **Decimal** (floating point numbers). The **pacap-interfaces** package declares shareable interfaces which enable the purse applet to communicate with *e.g.* the loyalty applets on the card. The **purse** package is the core of the purse application. It contains the class **PurseApplet**, which manages the operations related with installation, selection and deselection of the applet, and which communicates with the point of sale terminal. The basic purse functionalities are implemented in the class **Purse**. This class performs the communication with the loyalty applets, using the interfaces described in the **pacapinterfaces** package. Also, this class keeps track of the balance of the purse, the transactions done by the purse (stored as a **TransactionRecord**), the different currency changes that have taken place (in an **ExchangeRecord**) and the different loyalty programs that the card holder is subscribed to (in a **LoyaltiesTable**).

Certain operations can only be performed by a restricted set of users, *e.g.* because a pin code is needed. The class **AccessCondition** defines the different access conditions, and the class **AccessControl** binds the access conditions to the operations. So, when a card holder intends to perform a certain operation, the purse application will check that the card holder has the appropriate permissions. The class **Currencies** stores the different currencies used by the purse application. Finally, the purse application contains several classes implementing cryptographic concepts, namely, **PacapCertificate**, **PacapCipher**, **PacapKey**, **PacapRandom**, **PacapSecureMessaging** and **PacapSignature**. These classes are not studied in full detail in this case study.

3 Static checking of Java programs

ESC/Java is a verification tool developed at Compaq SRC, which permits a user to find common errors in Java programs. The basic idea is that a user specifies the desired behaviour of a class and its methods and the ESC/Java tool checks whether the implementation satisfies the specification. If it cannot establish this, it issues a warning. As explained above, such a warning does not necessarily mean that there is an error, as ESC/Java is neither sound, nor complete.

The specifications are given as *pre-* and *postconditions* of methods and as *class invariants*. The specifications are written as special Java comments, thus they do not change the annotated program. The properties are specified as Java expressions, enriched with several specification-specific constructs. Here, we present some ESC/Java specification constructs, together with an example of their use. Their full description can be found in [12].

3.1 ESC/Java pragmas to specify method and class behaviour

- **requires** P. This pragma specifies a method precondition P. When ESC/Java checks the body of the method, it assumes that P holds initially, but when ESC/Java checks a method call, it will issue a warning if it can not establish that P holds at the call site.
- **ensures** Q. This pragma specifies a method postcondition Q. The postcondition is supposed to hold if the method terminates normally, *i.e.* without throwing an exception.
- **exsures** (E) R. This pragma specifies a exceptional condition. This condition is supposed to hold if the method finishes abruptly and if the exception e that is thrown is a subclass of E.
- **modifies** L. This pragma specifies that a method *may* modify the state components listed in L, where these state components are variable names, field or array accesses and array range expressions (denoting the elements within an array). Within a method body, the method parameters and the local variables always may be modified. When checking a method call, ESC/Java assumes that only the state components denoted by the modifies clauses may have been changed, but it does not check the correctness of the modifies clause.
- **assert** P. This pragma states that the property P should be true whenever control reaches this program point.
- **invariant** I. This pragma specifies a class invariant, *i.e.* the property I has to be established by the constructor of the class and it has to be preserved by all the methods in the class.

3.2 Specification expressions

- **==>** is the logical implication. So, $P \implies Q$ is true if and only if P is false or Q is true, where P and Q are specification expressions of `boolean` type. Further, **<==>** denotes logical equivalence and **<!=>** specifies non-equivalence.
- **(\forall T V; E)** and **(\exists T V; E)** are quantifier expressions of type `boolean`. The first one denotes that E is true for all substitutions of values of type T for the bound variable V. The second one denotes that E is true for at least one substitution of a value of type T for the bound variable V.
- **\old(E)** is used within a postcondition, where it denotes the value of E in the pre-state of the method invocation.
- **\result** represents the value returned by a non-void method. It can only be used within an **ensures** clause.

```

/*@
  modifies nbData, data[nbData];
  ensures (\old(nbData) < MAX_DATA) ?
         (nbData == \old(nbData) + 1 && data[\old(nbData)] == cur) :
         (nbData == \old(nbData));
*/
void addCurrency(byte cur){
  if(nbData < MAX_DATA) {
    data[nbData] = cur ;
    nbData++ ;
  }
}

```

Fig. 1. Example ESC/Java specification

Fig. 1 shows a typical annotation example using ESC/Java. This example comes from the specification of the electronic purse [4]. The `addCurrency` method belongs to the class `Currencies`. This class stores all currencies supported by the purse application. The method `addCurrency` adds a new currency to the list of valid currencies. This list is represented by the array `data`. The `modifies` clause declared in the method's header specifies that this method may modify `nbData` and `data` in the position `nbData`¹. The postcondition of the method `addCurrency` – written as `ensures` clause – expresses that if `nbData` has not yet reached the threshold value `MAX_DATA`, `nbData` will increase its value by one and the value of the formal parameter `cur` will be assigned to `data[\old(nbData)]`, otherwise `nbData` remains unchanged. Inside the postcondition, the expression `\old(nbData)` refers to the value of `nbData` before the method invocation.

4 Specification of the Electronic Purse

4.1 The general specification approach

ESC/Java forces one to start writing specifications for the classes that are ‘used’ by many other classes, either because they are used as components or because they are inherited from. In the electronic purse case study most classes inherit directly from classes as *e.g.* `Object`, `Exception` or – in the case of interfaces – `Shareable`, so the inheritance structure is not very complex. Therefore, we started by specifying classes that provide basic (and general) features, *e.g.* those in the `utils` package, that are used by the classes in the `purse` package. The specifications for these basic classes form the basis for the specification of the

¹ More precisely, it specifies that the method body only may modify these instance variables and the local variables and formal parameters of the method.

more application-specific classes, so it is important that they are sufficiently detailed.

For every method, we specify the precondition (as a `requires` clause in ESC/Java), the postcondition (`ensures`), the modifies clause (`modifies`), and the exceptional postcondition (`exsures`). ESC/Java does not have a keyword to specify that a method may not modify any variables, but this is implied by the absence of a modifies clause. To make our specifications explicit about this, in such a case we added a comment `modifies \nothing; –` as in JML [11]. Further, ESC/Java requires that every exception that is mentioned in the exceptional postcondition is also mentioned in the `throws` clause of the method. To avoid having to add `throws` clauses to every method, in many cases we chose to have the assertion `exsures (Exception) false; –` meaning that no exception will be thrown – as a comment, without having it checked by ESC/Java. However, everywhere where there can be any doubt about the correctness of the `exsures` clause, we add the `throws` clauses and have it checked by ESC/Java.

When writing method specifications, two different styles can be used: either a precondition is given which ensures that no exceptions will be thrown, or one specifies a light precondition (*e.g.* `true`), and an exceptional postcondition which describes under which conditions an exception will be thrown. For example, given the left specification, one has to show that `P` is satisfied before the method is called, and then it is guaranteed that the method cannot produce an exception, while the right hand specification makes no requirements on the method call, but specifies that if an exception occurs, this is because `P` did not hold in the pre-state.

```

/*@ modifies M;
   requires P;
   ensures Q;
   exsures (E) false;
*/
void m() {
...
}

/*@ modifies M;
   requires true;
   ensures Q;
   exsures (E) !\old(P);
*/
void m () {
...
}

```

In our specifications, we usually follow the first approach, unless the informal documentation clearly suggests that the second approach is intended.

Further, we specify appropriate class invariants for each class, typically restricting the set of legal values for the instance variables. In some cases, the class invariant immediately follows from the informal documentation (*e.g.* the documentation in class `Decimal` states: **the decimal part must be done in the interval [000,999]** [8]) and in other cases the appropriate class invariant follows from closer inspection of the code, *e.g.* a variable is never `null`. Section 4.2 discusses the specification of class invariants in more detail.

Sometimes discrepancies between the informal documentation and the implementation occur. In general we try to follow the informal documentation, and we correct the implementation where necessary (and document these changes).

In several cases we consulted the case study developers at Gemplus, to get a better understanding which behaviour was actually intended.

In the case study, several functions from the JavaCard API [17] are used. When we specify methods using API functionalities, we use the API specification as constructed by Erik Poll and Hans Meijer (see [7, 14]). In the classes `Purse` and `PurseApplet`, several classes are used that we do not have access to. To overcome this problem, we construct specification files, declaring the methods and fields that we need, but without making any assumptions about their behaviour.

Our aim is to give a functional specification of the behaviour of the purse. However, we did not study the algorithms to manage secret keys, and therefore we only give a lightweight specification (*i.e.* specifying the precondition and modifies clauses, but no postcondition) of the classes dealing with key generation and certification. This enables us to write and check the specifications of the classes `Purse` and `PurseApplet`. How to specify and verify cryptographic algorithms is a topic of future research.

We aim at giving specifications which describe the behaviour of the application as complete as possible. As ESC/Java is not complete, it will sometimes produce a warning for a correct specification. Typically, if a complex control structure occurs in a method (*e.g.* loops in which method calls are made) ESC/Java is unable to establish complicated postconditions. However, in the case study at hand such complex control structures are not very frequent and ESC/Java is able to check most of the specifications without any problems. If one wishes to certify these methods, other verification techniques, as advocated *e.g.* in the LOOP project [13], should be used. As an example of such a verification, the addition and multiplication methods of class `Decimal` have been verified within the LOOP project [2]. In the final version of the specifications, the only remaining warnings are caused by ESC/Java's incompleteness. When we encountered other warnings during the specification and checking process, we adapted the implementation or specification appropriately.

At [4] the full annotated version of the purse case study can be found. In the code it is documented which postconditions cannot be established by ESC/Java. It is also documented which changes we have had to make to the code.

4.2 Interesting aspects of the specification

Below, several interesting aspects of the specification are discussed in more detail. First we elaborate on some implementation errors that we found in the purse application. Then we discuss the specification of class invariants, and how this can help to simplify the code. Finally, we discuss miscellaneous aspects of the case study, and present some possible improvements. The problems that we have found probably also would have been found by doing thorough testing, but using theorem proving techniques one is sure not to forget some cases, without having to put much effort in developing test scenarios. Also, writing the formal specifications forces one to think very precisely about the intended behaviour of programs, which helps in finding errors.

```

/*@
  requires d != null;
  ensures \result == (intPart>d.getIntPart() ||
                    (intPart == d.getIntPart() &&
                     (decPart == d.getDecPart() || decPart > d.getDecPart())));
*/
public boolean isGreaterEqualThan(Decimal d){
  boolean resu = false ;
  if(intPart>d.getIntPart()) resu = true ;
  else if(intPart<d.getIntPart()) resu = false ;
  else if(intPart==d.getIntPart()){
    if((decPart>d.getDecPart()) || (decPart>d.getDecPart())) resu=true ;
    else if(decPart<d.getDecPart()) resu = false ;
  }
  return resu ;
}

```

Fig. 2. Method `isGreaterEqualThan`

Implementation mistakes This section presents some examples of common programming errors, and how we found them using ESC/Java.

The method `isGreaterEqualThan` The class `Decimal` represents a floating point number composed of a decimal part and an integer part, denoted by instance variables `decPart` and `intPart`, respectively. The method `isGreaterEqualThan` (see Fig. 2) belongs to this class and, as suggested by its name and the informal documentation, it is supposed to decide whether the decimal represented by `this` is greater or equal than the decimal represented by parameter `d`. This behaviour is specified in the method specification.

However, after running ESC/Java on this asserted method, a warning is issued, suggesting that the postcondition might not hold. Inspection of the code reveals a “copy paste” error in the fourth `if` statement, where the condition `decPart > d.getDecPart()` is tested twice, on both side of an `||` (or) operator. Replacing the whole expression by `decPart >= d.getDecPart()` would solve the problem, although it would probably be better to rewrite the method in such a way that it simply tests the condition as expressed in the postcondition.

Final modifiers The class `Annee` represents a *year*. It declares two static variables called `MIN` and `MAX`, which represent the minimum and maximum year allowed by the application. Its declarations are as follows:

```

public static byte MIN = (byte)99 ;
public static byte MAX = (byte)127 ;

```

The class `Annee` also defines a method `check`, which is used to determine whether a value is between `MIN` and `MAX`. The class `Date` has three instance

```

/*@
  modifies jour, mois, annee;
  requires j >= Jour.MIN && j <= Jour.MAX;
  requires m >= Mois.MIN && m <= Mois.MAX;
  requires a >= Annee.MIN && a <= Annee.MAX;
  ensures jour == j && annee == a && mois == m;
*/
public void setDate(byte j, byte m, byte a) throws DateException{...}

```

Fig. 3. Fragment of class `Date`

variables, representing the components of a date: `jour` (day), `mois` (month), and `annee` (year). The method `setDate` (see Fig. 3) in this class assigns its arguments to these instance variables, provided they are in a valid interval (see the `requires` pragma). Surprisingly, ESC/Java complains when it finds a statement such as `date.setDate((byte)1, (byte)1, (byte)110);` (where `date` is an instance of class `Date`). The warning message states that the the third precondition of this call might not hold, even though 110 is between 99 and 127.

The problem is caused by the erroneous declaration of the variables `MIN` and `MAX` in class `Annee`. Since these variables are not declared `final`², their values can be changed at runtime by a direct assignment (as they are declared public), and thus ESC/Java warns correctly that the precondition of `setDate` might not be satisfied.

Class invariants Typically, invariants are used to restrict the state space of a class, *i.e.* the set of allowed values for its instance variables. The most common example is an invariant which states that a reference may never be a null pointer, *e.g.* the variable `purse`, as declared in the class `PurseApplet` should never be null.

```
/*@ invariant purse != null;
```

Another common example of an invariant is to restrict the possible values of a numeric variable to a certain range. As remarked above, the class `Decimal` says that the value of the decimal fraction part must be between 0 and 999. Inspection of the code reveals that the integer part of the decimal number is supposed to be a positive short, and combining this gives the following class invariant:³

```
/*@
```

² According to the Java semantics, final variables may only be assigned to when they are initialised, and afterwards they remain constant.

³ `MAX_DECIMAL_NUMBER` is equivalent to the maximal value of a short and the clause `intPart <= MAX_DECIMAL_NUMBER` of the invariant will thus be ensured by the type of the variable. We chose to state this explicitly for clarity of specification.

```

    invariant decPart >= 0 && decPart < PRECISION ;
    invariant intPart >= 0 && intPart <= MAX_DECIMAL_NUMBER;
*/

```

Another way to use class invariants, is to improve the simulation of enumeration types, which are not available in Java(Card). To simulate them, typically several constants with suggestive names are defined and a variable is silently assumed to contain always one of these values. This implicit assumption can be made explicit by specifying invariants. For example, the class `Transaction` contains the following declarations.

```

public static final byte INDETERMINE           = (byte)0;
public static final byte TYPE_CREDIT          = (byte)50;
public static final byte TYPE_DEBIT           = (byte)51;

/* the transaction type: debit or credit*/
/*@ spec_public*/ private byte type;

```

The documentation above suggests that the variable `type` always should have a value `TYPE_CREDIT` or `TYPE_DEBIT`. However, in the code (in the method `reset()`), an assignment `type = INDETERMINE;` occurs, suggesting that this is also a correct value for `type`. Having a specification which states the allowed values for this variable avoids all confusion⁴.

```

/*@ invariant type == INDETERMINE ||
           type == TYPE_CREDIT ||
           type == TYPE_DEBIT;
*/

```

Invariants of this kind often occur in the specification of the electronic purse. It is easy to specify them, and useful as well, as there are examples in the electronic purse where such implicitly assumed invariants are violated. For example, the class `AccessCondition` declares constants to state the different access conditions for the actions in the purse. Following [1], variables that denote access conditions should be restricted as follows.

```

/*@ invariant condition == FREE ||
           condition == LOCKED ||
           condition == SECRET_CODE ||
           condition == SECURE_MESSAGING ||
           condition == (SECRET_CODE | SECURE_MESSAGING);
*/

```

However, in the constructor of this class, the variable `condition` is set to 0, which breaks this invariant⁵. Correcting this and maintaining the invariant also

⁴ However, notice that this does not give type safety, in contrast to real enumeration types.

⁵ As none of these constants is equal to 0.

allows to improve other parts of the implementation in this class. For example, in the method `verify()`, the following statement occurs:

```
switch(condition) {
  case FREE: ...
  case SECRET_CODE: ...
  case SECURE_MESSAGING: ...
  case SECRET_CODE | SECURE_MESSAGING: ...
  case LOCKED: ...
  default: /*@ assert false;
           t = AccessConditionException.CONDITION_COURANTE_INVALIDE;
           AccessConditionException.throwIt(t);
}
```

Because of the invariant we know that the default case will never be reached (as signalled by the `/*@ assert false;` annotation, which states that false should hold, every time this program point is reached), and thus that the exception never will be thrown. Therefore the default case can be removed from the code.

Similar cases occur frequently with `try-catch` statements. An operation is executed within a `try`, but as the class invariants assure that the operation never will throw an exception, the `catch` clause will never be executed. In the specification, we have annotated these cases with `/*@ assert false;`. We think that the removal of this “dead code” can improve the readability of the class and, importantly for smart cards, it reduces the size of the byte code.

Miscellaneous aspects of the specification There are many other aspects of the specification that are worth mentioning. Here we mention some.

- As explained above, in the class `Decimal`, two shorts are maintained denoting the integer and the decimal part (`intPart` and `decPart`, respectively) of a decimal number. The integer part ranges between 0 and `MAX_DECIMAL_NUMBER` (which is 32767, the maximal value for shorts). It is left unspecified whether numbers such as `MAX_DECIMAL_NUMBER.999` are allowed. However, a method `round()` is defined, which according to the documentation returns a decimal number with `decPart` set to 0 and `intPart` set to the closest integer value. An obvious specification of this method reads as follows:

```
/*@
  modifies intPart, decPart;
  ensures decPart == 0;
  ensures intPart == (\old(decPart) >= (PRECISION/2) ?
                    (short)(\old(intPart) + 1) :
                    (short)(\old(intPart)))
*/
public Decimal round(){ ... }
```

But, as pointed out by ESC/Java, an implementation of this specification breaks the class invariant `intPart >= 0`. The counterexample that is produced has `intPart` set to `MAX_DECIMAL_NUMBER` and `decPart` *e.g.* to 999. Possible solutions are to specify explicitly the outcome of `round()` in the case that `intPart == MAX_DECIMAL_NUMBER`, or to restrict the set of valid decimal numbers by further strengthening of the class invariant through addition of the following clause:

```
/*@ invariant intPart == MAX_DECIMAL_NUMBER ==> decPart == 0;
```

We chose this last solution.

- Among the developers of the electronic purse application there apparently have been different ideas about the implementation of the class `Decimal` (which could have been avoided if the class invariants immediately would have been specified explicitly in the class). The implementation of several `setValue(...)` methods reveal that `intPart` is assumed to be greater or equal than 0, but on the other hand there are methods `isNegatif()` and `isPositif()`, which test whether a decimal value is negative or positive, respectively. As we specify⁶ that `intPart` should be greater or equal than 0 these methods become obsolete. We can show this by specifying that their results can be predicted, *e.g.* `isNegatif()` we specify as follows:

```
/*@
  ensures \result == false;
*/
public boolean isNegatif(){ ... }
```

- Two classes, `TransactionRecord` and `ExchangeRecord` implement a cyclic table (of `Transactions` and `ExchangeSessions` (currency changes), respectively). These implementations are clearly copied from each other, but this is nowhere documented. Also the fact that a cyclic data structure is implemented is not clearly documented. Class `TransactionRecord` contains a single remark that it is implemented as a cyclic table, and in class `ExchangeRecord` this is only stated in the documentation of a private method. Also, no specification of the operations on the cyclic data structure are given. As a result, in class `ExchangeRecord`, part of the code that is crucial for its behaviour has been commented out by other developers of the electronic purse. Having a formal specification would probably have been helpful to explain the complexity of the implementation to the other developers, and the “wrong correction” would have been signaled earlier⁷. Finally, when writing the formal specifications of the cyclic tables, we found an error in the implementation. When a delete operation is called for an element that is not in the range of the table, the operation nevertheless will be executed and as a side-effect it will corrupt the table by erroneously moving its first element outside the range of the table.

⁶ *cf.* our email exchange with H. Martin, Gemplus.

⁷ Of course, having a general implementation of a cyclic table and instantiating this for the different kinds of data would have been even more elegant.

```

/*@
  modifies \fields_of(this), \fields_of(date), \fields_of(heure),
          id[*], terminalTC[*], terminalSN[*];
  requires es != null ;
  requires es.id != terminalTC & es.id != terminalSN &
          es.terminalTC != terminalSN;
  ensures this.equal(es);
  exsures (TransactionException e)
          e._reason == TransactionException.BUFFER_FULL
          && JCSystem._transactionDepth == 1;
  exsures (NullPointerException) false;
  exsures (ArrayIndexOutOfBoundsException) false;
*/
void clone(ExchangeSession es) { ...
}

```

Fig. 4. Specification of `clone` in `ExchangeSession` in ideal ESC/Java

5 On the use of ESC/Java

We find ESC/Java a useful tool, which is pleasant to work with, but nevertheless we have some suggestions for improvements, both for the specification language and for the checker.

Concerning the specification language, we feel that certain specification constructs that are available in JML [11] should be provided in ESC/Java as well, in order to be able to write clear and concise specifications.

First of all, it would be convenient to have some extra specification constructs for `modifies` clauses, *e.g.* `\fields_of(E)`, to denote all the fields of an object, and `\nothing`. This could easily be implemented as syntactic sugar, in particular `modifies \nothing`.

Another improvement that would be easy to implement, would be to enable the specification of runtime exceptions in `exsures` clauses, without mentioning them explicitly in the `throws` clause of the method.

Also we feel that having some extra quantifiers, such as `\min`, `\max`, and `\choose` could be useful to increase expressiveness of the specification language. However, to implement this would require an extension of the theorem prover underlying ESC/Java, so that it also can deal with these language constructs.

Finally, we would like to be able to use method names in specifications, as is allowed in JML for so-called *pure* methods, *i.e.* methods without side-effects, but this would require a major change to ESC/Java.

Fig. 4 shows as an example the specification of the method `clone()` in class `ExchangeSession` the way we would prefer it (see [4] for the ESC/Java specification as it is). We only specify that all the fields of the current class may be modified, without explicitly mentioning them. As the fields of the component

classes `date` and `heure` may be modified as well, we mention this explicitly. Similarly, we mention explicitly that the elements in the arrays `id`, `terminalTC` and `terminalSN` may be modified⁸. Further, instead of having postconditions stating that all the fields are ensured to be equal to the corresponding fields of `es`, this is denoted by writing `this.equal(es)`, where `equal` is overwritten appropriately in `ExchangeSession`.

With respect to the verification that is done by ESC/Java, we found that it is unfortunate that ESC/Java does not try to check the modifies clauses, because an incorrect modifies clause can influence the acceptance of other specifications. For example, suppose one has the following (annotated) methods:

```
/*@ modifies x;
    ensures x == 3;
*/
void m() { x = 3;
          n (); }

void n() { x = 4; }
```

Remember that a method without any modifies clause is assumed to modify only freshly allocated memory, if any [12]. The specification for method `m()` is thus accepted by ESC/Java, although it is incorrect. When annotating existing programs, as we did, it is easy to forget to mention that a variable may be modified, and we felt the need to overcome this problem. Therefore, we have implemented a static checker for modifies clauses. In the tradition of ESC/Java, this checker is designed to be efficient, but it is neither sound, nor complete. It does a syntactic analysis of the annotated program to recognise the various assignment statements and then checks whether the variables that are the “destination” of an assignment are appropriately specified in the corresponding modifies clauses. This checker will be described in more detail in a separate paper [3].

Finally, we feel that it would be an important improvement if the ESC/Java theorem prover would deal more precisely with arithmetic operations (also on bytes and shorts). For example, the current version of ESC/Java issues a warning for the following specification.

```
/*@ requires b == (byte)4 & d == (byte)8;
    ensures \result >= 0;
*/
byte m(byte d) {return (byte)(b | d);}
```

We found that almost all spurious warnings that are produced by ESC/Java are caused by arithmetic operations⁹ and it would be a significant improvement if less of these warnings would be generated.

⁸ In JML this whole modifies clause also can be written as `\fields_of(\reach(this))`, which would probably also be useful in many cases, but has a more complex semantics.

⁹ Warnings about loops are not considered to be spurious, as they are inherent to how ESC/Java works.

6 Conclusions

We have presented a case study in formal specification of smart card programs, using ESC/Java. We have taken an electronic purse application and annotated it with a functional specification, describing its behaviour, basing ourselves on the informal documentation of the application. We have checked the implementation *w.r.t.* the specification, using ESC/Java, thereby revealing several errors in the implementation. Using ESC/Java we were also able to find that some parts of the program will never be reached, thus allowing reduction of the code size – which is important for smart card applications.

The whole case study consists of 42 classes and 432 kB in total (736 kB with annotations). It has taken approximately three months to write the complete specifications. Most of the specifications are written by the first author, who beforehand did not have any experience with ESC/Java, or with writing formal specifications in general. The second author – who had experience in writing formal specifications, but not with ESC/Java – supervised the work and made suggestions to extend the specifications. Of the time spend on writing the specifications, approximately one third was used on getting to know ESC/Java, and understanding the electronic purse application, the remaining time was used on writing and checking the actual specifications.

The errors that we found in general are not very intricate, they could have been found by careful code-inspection or testing. But, writing the formal specification (for existing code) forces one to do code-inspection, and having the theorem prover ensures that all cases are considered, when checking the specification, without having to put effort in writing appropriate test scenarios.

The specifications that we have constructed for the electronic purse application are not very complex, but describe the functional behaviour of methods as precisely as possible. Nevertheless, we found errors in the code, and we would like to emphasise that even simple formal verification can help significantly to increase confidence in a program. In particular, explicitly specifying class invariants – which are often implicitly assumed by the program – turns out to be very useful.

Future work We plan to work in the field of specification languages for Java: how to improve them, and how to develop appropriate verification techniques for them. In particular, we will focus on the following points.

- We plan to develop a full smart card application from scratch, with annotations. We are interested whether this will affect the quality of the specification and/or the program. Also we would like to know how easy it is to construct the specifications at the same time as developing the code. We would like to evaluate the tradeoff between usability – because one gets immediate feedback on an implementation – and the extra time that is spend on keeping the specification up-to-date.

- Future versions of JavaCard will probably allow multi-threading. Therefore we plan to study how ESC/Java (and JML) can be used to specify (and check) concurrent programs.
- Related with this is an extension of JML with temporal logic. Currently we are studying how to integrate temporal logic in the specification language [19], future work will be to study appropriate verification techniques.
- Most loop structures that are used in typical smart card programs are very restricted and it is relatively easy to show their termination. We plan to develop an automatic verification technique for termination of loops in the tradition of ESC/Java, covering the most common cases.
- We skipped the cryptographic aspects of the application at hand. It is future work to see whether ESC/Java (or JML) is useful to specify such algorithms more precisely, and to develop appropriate (automatic) verification techniques.

Acknowledgements

We thank Erik Poll, Dilian Gurov and Arnd Poetzsch-Heffter for useful feedback on the specifications and on earlier versions of this paper. Also we would like to thank Rustan Leino and his team for their help with ESC/Java.

References

1. E. Bretagne, A. El Marouani, P. Girard, and J.-L. Lanet. Pacap purse and loyalty specification. Technical Report V 0.4, Gemplus, 2000.
2. C. Breunese, B. Jacobs, and J. van den Berg. Specifying and Verifying a Decimal Representation in Java for Smart Cards. In H. Kirchner and C. Ringeissen, editors, *Algebraic Methodology And Software Technology (AMAST'02)*, number 2422 in LNCS, pages 304–318. Springer, 2002.
3. N. Cataño and M. Huisman. A static checker for JML's *assignable* clause, 2002. Manuscript.
4. N. Cataño and M. Huisman. Annotated files of the Electronic Purse case study, 2002. http://www-sop.inria.fr/lemme/verificard/electronic_purse.
5. Differences between Esc/Java and JML, 2000. Comes with JML distribution, in file `esc-jml-diffs.txt`.
6. Extended static checking for Java. <http://research.compaq.com/SRC/esc/>.
7. ESC/Java specifications for the JavaCard API. http://www.cs.kun.nl/~erikpoll/publications/jc211_specs.html.
8. Gemplus. Applet benchmark kit. http://www.gemplus.com/smart/r_d/publications/case-study/.
9. M. Huisman, B. Jacobs, and J. van den Berg. A Case Study in Class Library Verification: Java's Vector Class. *Software Tools for Technology Transfer*, 3/3:332–352, 2001.
10. The JASS project. <http://semantik.informatik.uni-oldenburg.de/~jass/>.
11. G.T. Leavens, A.L. Baker, and C. Ruby. Preliminary design of JML: a behavioral interface specification language for Java. Technical Report 98-06i, Iowa State University, Department of Computer Science, 2000.

12. K.R.M. Leino, G. Nelson, and J.B. Saxe. ESC/Java user's manual. Technical Report SRC 2000-002, Compaq System Research Center, 2000.
13. The LOOP project. <http://www.cs.kun.nl/sos/research/loop>.
14. H. Meijer and E. Poll. Towards a Full Formal Specification of the Java Card API. In I. Attali and T. Jensen, editors, *Smart Card Programming and Security (E-smart 2001)*, number 2140 in LNCS, pages 165–178. Springer, 2001.
15. B. Meyer. *Object-Oriented Software Construction*. Prentice Hall, 2nd rev. edition, 1997.
16. J. Meyer and A. Poetsch-Heffter. An architecture of interactive program provers. In S. Graf and M. Schwartzbach, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2000)*, number 1785 in LNCS, pages 63–77. Springer, 2000.
17. Sun Microsystems, Inc. Java Card 2.1. Platform Application Programming Interface (API) Specification. <http://java.sun.com/products/javacard/htmldoc/>.
18. Sun Microsystems, Inc. JavaCard Technology. <http://java.sun.com/products/javacard/>.
19. K. Trentelman and M. Huisman. Extending JML Specifications with Temporal Logic. In H. Kirchner and C. Ringeissen, editors, *Algebraic Methodology And Software Technology (AMAST'02)*, number 2422 in LNCS, pages 334–348. Springer, 2002.