A Refinement Method for Java Programs

Holger Grandy, Kurt Stenzel, Wolfgang Reif

Lehrstuhl für Softwaretechnik und Programmiersprachen
Institut für Informatik, Universität Augsburg
86135 Augsburg Germany

E-Mail: {grandy, stenzel, reif}@informatik.uni-augsburg.de

Abstract. We present a refinement method for Java programs which is motivated by the challenge of verifying security protocol implementations. The method can be used for stepwise refinement of abstract specifications down to the level of code running in the real application. The approach is based on a calculus for the verification of Java programs for the concrete level and Abstract State Machines for the abstract level. In this paper we illustrate our method by the verification of a M-Commerce application for buying movie tickets using a mobile phone written in J2ME. For verification we use KIV, our interactive theorem prover [1].

1   Introduction

Refinement is an established method for proving algorithms correct. A concrete specification is a refinement of a more abstract specification if every state change that can be performed on the concrete level is also possible on the abstract level. State based refinement methods (e.g. [8] [30] [3]) have been used in numerous case studies for the verification of algorithmic correctness. The underlying theory and the methods for applying those approaches, also on the level of tool support, are elaborated and widely used.

Much less work has been done on refinement methods for the verification of Java implementations. Although there are many examples of Java [17] program verification, e.g. [16] [5] [6] [22] [15], the authors are not aware of a larger case study of interactive verification using a refinement framework for proving full functional correctness of a Java program respecting an abstract specification.

In the field of security protocol implementations the past has shown that implementation flaws are very common and can be very subtle. In this paper, we present a general refinement method for Java programs inspired by the challenge of verifying security protocol implementations. The method is illustrated by the verification of a Java M-Commerce application, the Cindy[1] case study. The refinement approach is not limited to the field of security protocols. Using the mechanisms described below we can prove functional correctness for all kinds of programs with input, output and state change.

---

[1] Cinema Handy (Handy is the German word for mobile phone)

The paper is organized as follows: Section 2 presents the case study, Section 3 illustrates the specifications for refinement and proof obligations. Section 4 describes the mapping of abstract data types to Java classes. Section 5 presents some difficulties the refinement method has to solve stemming from this mapping and Section 6 gives some details on the verification of the case study. Finally, Section 7 compares the approach to related work and Section 8 concludes.

2   The Cindy Case Study

With Cindy users can buy cinema tickets using mobile phones. A user can order a ticket using a Java application running on the device. Payment can be done using the usual phone bill. After having ordered a ticket it is sent to the mobile phone as a MMS (Multimedia Messaging Service) message. The ticket contains the movie data and an additional unique
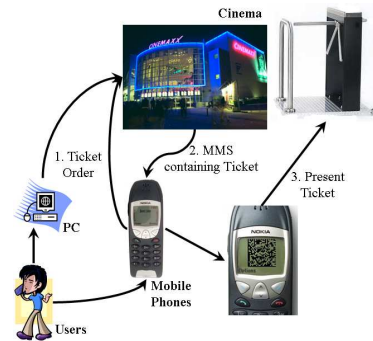

Fig. 1. The Cindy Application

identifier for the ticket. It can be displayed on the phone using a two-dimensional data matrix barcode and is scanned at the entrance to the cinema directly from the display using a barcode scanner. This kind of application exists e.g. in the Netherlands [2]. Additionally, the German railway company, Deutsche Bahn, has recently implemented a similar service for buying train tickets using a mobile phone.

One important question for the cinema is, of course, how to avoid fraud. The idea is simple: Every ticket contains a nonce, a unique random number that is too long to guess. Therefore it is virtually impossible to 'forge' a ticket.

Full details on the abstract model of Cindy as well as the details on the verification of security properties on this abstract level (which follows our approach for the verification of security protocols called PROSECCO) can be found in [10]. The next section describes the approach for verifying an implementation of Cindy running on a mobile phone written in J2ME.

3   Abstract and Concrete Specification Levels

We assume the reader is roughly familiar with data refinement theory, which in this section we will adopt to Java programs using the notation based on [9].

The abstract level is given as a data type $\mathsf{ADT} = (\mathsf{GS}, \mathsf{AS}, \mathsf{AINIT}, \{\mathsf{AOP_i}\}_{i \in I}, \mathsf{AFIN})$ consisting of a set of global states GS and a set of (local) states AS. Total relations $\mathsf{AINIT} \subseteq \mathsf{GS} \times \mathsf{AS}$ and $\mathsf{AFIN} \subseteq \mathsf{AS} \times \mathsf{GS}$ initialize and finalize the data

type. $\mathsf{AOP_i} \subseteq \mathsf{AS} \times \mathsf{AS}$ (using an index $\mathsf{i} \in \mathsf{I}$) are the operations possible on the data type. In the specification of the Cindy example different agents are involved modelling the different protocol participants. Every **agent** has a type **type(agent)** (the type can be **cellphone**, **cinema**, **user** or **attacker**). The index set I of $\mathsf{AOP_i}$ now consists of the different agents, where e.g. $\mathsf{AOP_{cellphone(n)}}$ denotes the protocol steps of the cellphone agent with number n.

On the conrete level, one agent in the protocol model is replaced by his Java implementation. So the concrete level is given similarly as $\mathsf{CDT} = (\mathsf{GS}, \mathsf{CS}, \mathsf{CINIT}, \{\mathsf{COP_i}\}_{i \in I}, \mathsf{CFIN})$, where one $\mathsf{COP_i}$ is a Java implementation. Details will be given later in Sect. 3.2.

Our operations are total so we use the approach of [13] and a forward simulation $\mathsf{R} \subseteq \mathsf{AS} \times \mathsf{CS}$ leading to the following proof obligations for refinement correctness:

- $\mathsf{CINIT} \subseteq \mathsf{AINIT} \, \mathring{,} \, \mathsf{R}$ ("initialization")
- $\forall \, \mathsf{i} \in \mathsf{I}. \, \mathsf{R} \, \mathring{,} \, \mathsf{COP_i} \subseteq \mathsf{AOP_i} \, \mathring{,} \, \mathsf{R}$ ("correctness")
- $\mathsf{R} \, \mathring{,} \, \mathsf{CFIN} \subseteq \mathsf{AFIN}$ ("finalization")

3.1 The Abstract Level

The state $\mathsf{as} : \mathsf{AS}$ consists of a function $\mathsf{astate} : \mathsf{agent} \to \mathsf{A_{type(agent)}}$ that maps each **agent** to its internal state in $\mathsf{A_{type(agent)}}$ For an agent of type **cellphone** this is e.g. the list of current tickets stored on a phone and its phone number. Additionally, **as** contains the current context $\mathsf{actxt} : \mathsf{context}$ of the communication infrastructure (connections and inputs for every agent that represent the messages that are currently in transit). Together $\mathsf{as} = \mathsf{astate} \times \mathsf{actxt}$. The global state GS contains only the list of tickets of the phones, since we want to show that this list is the same on both levels. GS is ignored in AINIT, AFIN extracts the list of tickets sold so far from GS.

The abstract specification of the functionality of the protocol in Cindy is given as an Abstract State Machine (ASM) [4] consisting of models for all the different agents in the scenario. Although not being used directly by the refinement theory, we use the different rules of this ASM to define the operations $\mathsf{AOP_{agent}}$. The ASM for Cindy is described in [10], so we only give a slight introduction here.

The interesting part of the abstract ASM specification for this paper is the step of an agent of type **cellphone** because this is the agent that will be refined to Java. An excerpt of the according ASM rule for the **cellphone** agent which actually loads a ticket on the mobile phone is:

```
APROG_cellphone(agent, tickets, inputs){
 let indoc = first(inputs(agent)) in
        inputs(agent) := rest(inputs(agent))
        if is_load_message(indoc) ∧ #tickets(agent) < MAXTICKETS
        then tickets(agent) := tickets(agent) + getPart(2, indoc)
        else ... // other protocol steps }
```

In this example, astate for the cellphone agent is given by the state function tickets, which stores the list of tickets of every agent. The context actxt is given by the inputs state function, which maps every agent to his current input messages. First an input message indoc is taken from the input (APROG$_{cellphone}$ is only called when the input is non-empty) and the list of input messages is shortened. If the input message has the correct structure of a message to load a ticket (is_load_message(indoc)) and there is space in the list of tickets of the actual agent (#tickets(agent) < MAXTICKETS) then the ticket contained in the input document (getPart(2, indoc)) is added to the list of tickets. For the refinement theory presented in this paper it is sufficient to know that the specification of Cindy consists of ASM rules APROG$_{agent}$ for every agent, which define the input/output behavior and the state changes of agent for every protocol step.

We use the Theorem Prover KIV [1] for our approach. In KIV, Abstract State Machines are modeled using Dynamic Logic (DL). In DL, the formula $\langle \alpha \rangle\, \phi$ states, that ϕ holds after the execution of program α. APROG$_{agent}$ is in fact a DL procedure. To integrate this into the data refinement theory presented above we define the operation AOP$_{agent}$ of ADT using APROG$_{agent}$:

AOP$_{agent}$(astate, actxt, astate′, actxt′) $\leftrightarrow$
$\langle$APROG$_{agent}$(astate, actxt)$\rangle$ (astate = astate′ $\wedge$ actxt = actxt′)

3.2 The Concrete Level

We now refine our abstract agent specification to Java. This works by step-wise replacement of an agent type and its abstract protocol step specification AOP$_{agent}$ by a Java implementation for agent, preserving every other part of the specification. In this paper, this is illustrated by the refinement of the cellphone agent type. Accordingly, the concrete level is a mixture of steps of agents, that are already replaced by a Java program (cellphone agent here) and other agents (the cinema server or the attacker), that are still preserved as on the concrete level. So the concrete state cs and the concrete operations COP$_{agent}$ are a mixture of Java implementation and abstract specification.

A concrete state cs : CS is defined as cs = cstate × cctxt with cctxt : context and cstate : agent $\rightarrow$ B$_{type(agent)}$. The context needs to be preserved like in the abstract level because the communication infrastructure is not implementable (it is a model of messages currently in transit). The state of a Java program is stored in an algebraic data type called store in KIV. A store can be seen as the equivalent of the heap of a Java virtual machine (in our case the JVM running on a mobile phone). All the runtime information about pointer structures is contained inside the store. Full details on the store and on the Java Calculus implemented in KIV can be found in [27] [26]. On the concrete level the state of a refined agent is now replaced by a store st : store. The state of non-refined agents remains the same as on the abstract level. This means that B$_{cellphone}$ = store and B$_{agenttype}$ = A$_{agenttype}$ for agenttype $\neq$ cellphone. Because we now integrate a Java implementation of an agent in our model, we have to do a data transformation step from the abstract data types specifying input and output of the agent into

the Java store and vice versa. The inputs of the cellphone agent (given by actxt on the abstract level) need to be mapped to Java data types representing the same input on the programming language level. This is done by a ASM rule called TOSTORE. The reverse transformation has to be done for the output, called FROMSTORE. More details on this transformation will be discussed later in section 4.

The Java method step() is the protocol implementation of the cellphone agent. For the sake of understandability the implementation itself will be presented later in Sect. 6. Java method calls are written in the Java calculus in KIV as $\langle st;\ step() \rangle\ \phi$, which states that formula ϕ holds after the execution of method $step()$ in the context of store st. Together with TOSTORE and FROMSTORE, we now define $\mathsf{COP_{agent}}$ as:

$\mathsf{COP_{agent}}(\mathsf{cstate}, \mathsf{cctxt}, \mathsf{cstate'}, \mathsf{cctxt'}) \leftrightarrow$
    if $\neg$ is_refined(agent) **then**
        $\mathsf{AOP_{agent}}(\mathsf{cstate}, \mathsf{cctxt}, \mathsf{cstate'}, \mathsf{cctxt'})$
    else $(\exists\, \mathsf{st}, \mathsf{st'}.\ \mathsf{st} = \mathsf{TOSTORE}(\mathsf{cctxt}, \mathsf{cstate}(\mathsf{agent})) \wedge$
               $\langle \mathsf{st};\ \mathsf{step}() \rangle\ (\mathsf{st} = \mathsf{st'}) \wedge$
               $\mathsf{cstate'} = \mathsf{cstate}[\mathsf{agent} \mapsto \mathsf{st'}] \wedge$
               $\mathsf{cctxt'} = \mathsf{FROMSTORE}(\mathsf{st'}, \mathsf{cctxt}))$

$\mathsf{COP_{agent}}$ is defined to be the same operation as on the abstract level ($\mathsf{AOP_{agent}}$) for all agents, that are not refined ($\neg$ is_refined(agent), for example the cinema). When agent is one of the agents, that are refined (is_refined(agent), here the cellphone), the $\mathsf{COP_{agent}}$ is defined using a Java implementation and TOSTORE and FROMSTORE operations: the inputs are transformed into Java objects in the store (TOSTORE(cctxt, cstate(agent))). Then a Java method call step() implementing the protocol and starting in this store st must result in a store st', which is given by cstate' (cstate' = cstate[agent $\mapsto$ st']). The output of the Java program is extracted from the store using FROMSTORE and this output forms the new concrete context cctxt'.

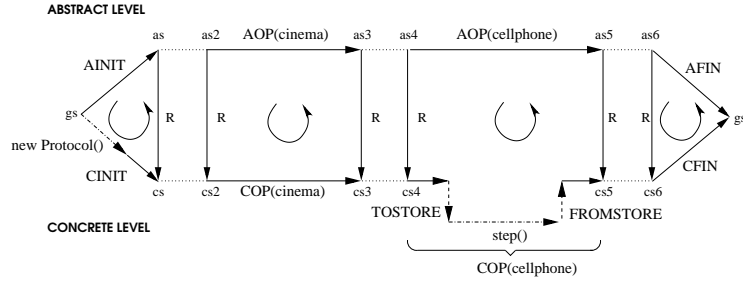### 3.3 Proof Obligations for the Example

Fig. 2. Refinement diagram

Fig. 2 gives an overview of the refinement proof obligations in Cindy for initialization, finalization and for the steps of the cinema agent (that is not refined in the example) and of the cellphone agent (which is refined to Java). The circle-like arrows illustrate the refinement proof obligations of commuting sub-diagrams. Fig. 2 also shows the operations TOSTORE and FROMSTORE before and after the Java method step() of the cellphone implementation is executed.

All together the main proof obligation for the refinement of the cellphone agent now is:

$$
\begin{aligned}
& \mathsf{R}(\mathsf{astate}, \mathsf{actxt}, \mathsf{cstate}, \mathsf{cctxt}) \\
& \wedge\ \mathsf{st} = \mathsf{TOSTORE}(\mathsf{cctxt}, \mathsf{cstate}(\mathsf{agent})) \\
& \wedge\ \langle\mathsf{st};\ \mathsf{step}()\rangle\ (\mathsf{st} = \mathsf{st}') \\
& \wedge\ \mathsf{cstate}' = \mathsf{cstate}[\mathsf{agent} \mapsto \mathsf{st}'] \\
& \wedge\ \mathsf{cctxt}' = \mathsf{FROMSTORE}(\mathsf{st}', \mathsf{cctxt}) \rightarrow \\
& \exists\,\mathsf{astate}', \mathsf{actxt}'.\ \mathsf{AOP}_{\mathsf{agent}}(\mathsf{astate}, \mathsf{actxt}, \mathsf{astate}', \mathsf{actxt}') \\
& \qquad\qquad\quad \wedge\ \mathsf{R}(\mathsf{astate}', \mathsf{actxt}', \mathsf{cstate}', \mathsf{cctxt}')
\end{aligned}
$$

If the retrieve relation holds for two states and the concrete level performs a sequence of TOSTORE, the actual protocol step step() and FROMSTORE, resulting in state $\mathsf{cstate}' \times \mathsf{cctxt}'$, then there must be the possibility to perform a similar step on the abstract level (AOP) which leads to a state $\mathsf{astate}' \times \mathsf{actxt}'$ in which the retrieve relation holds again. More details on the proof of this property will be given in Sect. 6.

Fig. 2 also shows the constructor call of the Java class implementing the protocol (new Protocol()), which is called during CINIT. We have to prove that the constructor call of the Java implementation performs the same initialization steps as AINIT for the refined agent type. This proof obligation is omitted here because it is very similar to the main proof obligation above (excepting TOSTORE and FROMSTORE because there is no input or output for the constructor).

One important point for the proof of our obligations is the definition of the retrieve relation R. It has to express how the state of the Java program and the abstract state of the protocol ASM relate to each other. Since we focus on security protocols, we can give a generic template for this relation. It is:

$$
\begin{aligned}
& \mathsf{R}(\mathsf{astate}, \mathsf{actxt}, \mathsf{cstate}, \mathsf{cctxt}) \leftrightarrow \\
& \mathsf{actxt} = \mathsf{cctxt} \wedge \mathsf{AINV}(\mathsf{astate}, \mathsf{actxt}) \wedge \mathsf{CINV}(\mathsf{cstate}, \mathsf{cctxt}) \wedge \\
& (\forall\ \mathsf{agent}.\mathbf{if}\ \mathsf{is\_refined}(\mathsf{agent})\ \mathbf{then}\ \mathsf{extract}(\mathsf{cstate}(\mathsf{agent})) = \mathsf{astate}(\mathsf{agent}) \\
& \qquad\quad \mathbf{else}\ \mathsf{cstate}(\mathsf{agent}) = \mathsf{astate}(\mathsf{agent}))
\end{aligned}
$$

The relation states the following: The extract function gets the state of the agent from the store (more precisely it looks at the fields of the classes implementing the protocol and converts those fields back into an abstract state). The state on the abstract level (astate(agent)) must be equal to the corresponding value in the store (extract(cstate(agent))), if agent is one of the agents that have a Java implementation. For the other agents the state on the concrete level must be exactly equal to the abstract level. The context (like the inputs of the agents) must be equal in every case. Additionally we need an invariant on the abstract state (AINV) and an invariant on the concrete state (CINV) that is preserved

by every step. The invariants basically state that everything is well-formed and reasonable for our application, e.g. the list of tickets contains only tickets, not other entries.

By proving the refinement, security properties of the abstract ASM specification level now can be transfered to the implementation level via the retrieve relation R. If a property is e.g. invariant for astate(agent) it is also invariant for extract(cstate(agent)) because of the refinement. In general, it is known that not all security properties are preserved under refinement (see e.g. [19]), but those problems arise only when the granularity changes during refinement. This is not the case in our refinement approach, because in our model both the abstract ASM rules and the concrete implementation steps are atomic operations of the same granularity, which last from the receiving of input to the sending of output for every agent. We do not consider attacks on the implementation which take place *during* the execution of a protocol step. This would mean changing of memory contents of the devices during execution and would of course allow a lot more attacks. Also we do not consider problems like power failures of the mobile phone in the middle of a protocol step execution.

4 Data Type Mapping to the Concrete Level

Java programs and Abstract State Machines use different internal types. On the one hand we have the Java class hierarchy (consisting of interfaces and classes) and primitive types, on the other hand we have algebraically specified abstract data types and state functions for the abstract specification level.

For our M-Commerce example same external behavior means sending of the same output messages in reply to the same input messages. On the abstract level input and output are specified using an abstract data type called document. This data type is quite similar to the messages used in [23] or [7]. It is specified algebraically as follows:

document = intdoc(.int : int)
          | keydoc(.key : key)
          | noncedoc(.nonce : nonce)
          | secretdoc(.secret : secret)
          | hashdoc(.doc : document)
          | encdoc(.key : key; .doc : document)
          | sigdoc(.key : key; .doc : document)
          | doclist(.list : documentlist)

A document can contain an arbitrary large integer (intdoc). The intdoc type is also used to model arbitrary data since every data can be represented as an integer. Documents can also contain a key (keydoc), a nonce (noncedoc) or a secret (secretdoc). Furthermore a document can be the result of a cryptographic hashing operation (hashdoc) or can be an encrypted document with a certain key (encdoc) or a signature of a document with a certain key (sigdoc). To model composition of messages our document type also contains a type doclist containing a list of other documents. In our ASM model the inputs of all agents are

represented as an ASM state function inputs : agent → documentlist (which is a part of the context described in section 3).

On the concrete level a natural representation of the abstract document data type is a class hierarchy which is directly implementing our abstract data type. The Cindy application relies on the security of GSM communication which already supports encryption of all sent messages. Therefore the protocol of Cindy only uses the type intdoc for modelling the ticket data or concepts like phone numbers, and noncedoc for modelling the unique identifier of the ticket. Additionally, the doclist type is used for composing those basic documents to MMS messages.

The class hierarchy we use in the implementation of Cindy is shown in Figure 3. We implement every constructor of the abstract data type document by a separate Java class type for exactly that type of document. For our general refinement approach to security protocols the other document types are implemented as well but omitted here. In addition


Fig. 3. Document Classes

to input/output behavior we furthermore have to prove that the same state changes are performed on both levels. In the Cindy example the state of the mobile phone consists of a list of documents representing tickets which are currently stored on the phone. This list is specified using the doclist abstract type on the abstract level, respectively implemented by the *Doclist* class for the concrete state. The state function tickets : agent → documentlist specifies this for the abstract level (part of astate(cellphone) as explained in Section 3). In addition the state function inputs : agent → documentlist is relevant for the refinement because it contains the input messages of each agent. Those two functions have to be taken into account for the refinement and have to be transformed to Java data types. Using the abstract data types and the store we define mapping functions for the transformation of the abstract data type into the concrete pointer structure inside the store and vice versa. The store defines a mapping of keys to values. Store keys are a combination of a reference (a memory address) and a class field or a array index. Getting the value for the field f of the instance at reference r is written as st[r.f]. The lookup for static fields can be written as st[.f]. The value can be a primitive value or a reference to another class instance or an array. The operations for the transformation of documents are called addDoc : document × store → reference × store and getDoc : reference ×store → document (all operations below are specified algebraically). addDoc for e.g. the *IntDoc* class type works as follows:

addDoc-intdoc:
$[r_1, r_2] = \mathsf{newrefs}(2, \mathsf{st}) \rightarrow$
    $\mathsf{addDoc}(\mathsf{intdoc}(i), \mathsf{st}) =$
        $r_1 \times \mathsf{addobj}(r_1, \mathsf{IntDoc}, .\mathsf{value} \times r_2,$
                $\mathsf{addarray}(r_2, \mathsf{byte\_type}, \mathsf{int2bytes}(i), \mathsf{st}))$

Adding an *Intdoc* with value i to the store works by adding an object of class *IntDoc* via the operation $\mathsf{addobj} : \mathsf{reference} \times \mathsf{type} \times \mathsf{fieldvalues} \times \mathsf{store} \rightarrow \mathsf{store}$. The reference r_1 of this new object must not be already contained in the store ($[r_1, r_2] = \mathsf{newrefs}(2, \mathsf{st})$). The actual value i of the *Intdoc* is encoded as an array of bytes. This array must also be added to the store via the operation $\mathsf{addarray} : \mathsf{reference} \times \mathsf{type} \times \mathsf{arrayvalues} \times \mathsf{store} \rightarrow \mathsf{store}$. The reference r_2 of this array must also be a new reference in the store ($\ldots = \mathsf{newrefs}(2, \mathsf{st})$). The array values are obtained by transforming the integer i to a sequence of bytes ($\mathsf{int2bytes}(i)$). The function addDoc additionally returns the reference r_1 of the *IntDoc* instance as well as the store because we have to know where the new instance is placed inside the store.

The getDoc function for the IntDoc type works the other way:

getDoc-intdoc:
$r \neq \mathsf{null} \wedge \mathsf{st}[r.\mathsf{type}] = \mathsf{IntDoc} \rightarrow$
    $\mathsf{getDoc}(r, \mathsf{st}) = \mathsf{intdoc}(\mathsf{bytes2int}(\mathsf{getbytearray}(\mathsf{st}[r.\mathsf{value}], \mathsf{st})))$

Getting the document of type *IntDoc* ($\mathsf{st}[r.\mathsf{type}] = \mathsf{IntDoc}$, where .type is a special field containing the type information of a reference) back from the store is done by first getting the byte array representing the value from the store ($\mathsf{getbytearray}(\mathsf{st}[r.\mathsf{value}])$). The resulting byte sequence is transformed to an integer using the operation $\mathsf{bytes2int}$ and the resulting integer value is used to construct the *Intdoc*.

The operations $\mathsf{TOSTORE}$ and $\mathsf{FROMSTORE}$ basically use addDoc and getDoc to transform the input messages of the agents into the Java store. Additionally getDoc implements the $\mathsf{extract}$ function described in Section 3 in the retrieve relation of the refinement for the list of tickets of an agent. This works because in Cindy both input/output messages and the state are specified using documents.

5 Additional Attacks on the Concrete Level

An interesting observation is the fact that when implementing the data types by pointer structures there are more possible values on the concrete level than on the abstract level. The reason is that on the concrete level there can be pointer structures that do not have any abstract counterpart. One example for this fact are instances of class *IntDoc* which contain a null pointer in their value field. Since the value field is the counterpart of the abstract value of the integer contained in the *IntDoc* and since null does not represent a number this document has no counterpart. In the following we will call those additional inputs invalid. A refinement respecting only valid inputs would not be correct because in the

real world other inputs than the abstract ones may be sent by an attacker and may cause implementation errors or security leaks.

The solution for this problem is to consider the invalid inputs on the concrete level by implementing a check on the input which checks whether the concrete input has an abstract counterpart. We add an additional document type $\perp$ (representing all the invalid inputs) and specify that the abstract level performs an error treatment (e.g. a reset operation on the internal state) when receiving $\perp$. Then the concrete step which receives an invalid input (and discovers this using the input check) has to be a refinement of the abstract error treatment step. With such a refinement nothing bad can happen on the concrete level when receiving invalid inputs. The TOSTORE operation now relates $\perp$ to all invalid documents. An attacker sending $\perp$ on the abstract level is now able to send any invalid document on the concrete level. Formally, the predicate validDoc : reference $\times$ store specifies whether a pointer structure represents an abstract document. The result $r \times st$ of addDoc always satisfies validDoc(r, st). The check for valid inputs is done in the receive() method in the Java implementation. Therefore the implementation of receive() must satisfy:

Receive-correct:
$\dots //$ *reference r is a valid communication interface in st*
$\wedge$ st $=$ st$_0$ $\rightarrow$
$\langle$st; $r_0 =$ r.receive(); $\rangle$
$\qquad$ st $=$ st$_0$[.input, null] $\wedge$
$\qquad$ ((validDoc(st$_0$[.input], st$_0$) $\rightarrow$ $r_0 =$ st$_0$[.input]) $\wedge$
$\qquad$ ($\neg$ validDoc(st$_0$[.input], st$_0$) $\rightarrow$ $r_0 =$ null))

If the input is a valid representation (validDoc(...)) of an abstract document, the return value r_0 of receive is the reference which was added in the TOSTORE operation (st$_0$[.input]). Otherwise null is returned. Additionally receive sets the input buffer to null (st[.input, null]).

It is not desirable to verify the correctness of a concrete input/output checker again for every single application. E.g. all our security protocol implementations use the document class type as the input type. We have used this type for the implementation of Cindy and also e.g. for the implementation of the Mondex [28] application. Also, a real implementation would not directly send pointer structures but do some kind of encoding (e.g. to byte arrays or XML, which is then sent by MMS). The data checker can be integrated in such a transformation function. We provide an implementation for such a transformation and data check layer which can be verified separately. This enables us to split the refinement proof into two layers. In the first layer the refinement of an abstract specification of the protocol into an implementation working on the document class type is shown using *receive-correct* as an assumption. The second refinement adds the transformation and data check layer. Then TOSTORE has to add an encoding of the input document instead of a pointer structure to the store. The receive method has to check this input and transform it into a pointer structure. Then the property of receive above can be proven using correctness properties of the check and transformation layer.

6   Details on the Cindy Refinement and Implementation

Sect. 3 showed an excerpt of the ASM specification for the cellphone agent, which covers storing of new tickets on the cell phone. The J2ME implementation[2] of this protocol step is:

```
public class Protocol {
 private Doclist tickets; // bought tickets
 ...

 public void step(){
  if(comm.available()){
   Document inmsg = comm.receive();
   phoneStep(inmsg);}}

 private void phoneStep(Document inmsg) {
  Document originator = inmsg.getPart(1);
  inmsg = inmsg.getPart(2);
  Doclist ticket = getTicket(inmsg);
  if(ticket != null && tickets.len() < MAXTICKETLEN){
   tickets = tickets.attach(ticket);}
  ... //other protocol steps}

 private Doclist getTicket(Document indoc) {
  if(indoc != null && indoc.is_comdoc()){
   byte[] ins = indoc.getPart(1).getValue();
   if(ins.length == 1 && ins[0] == LOADTICKET){
    Document indoc2 = indoc.getPart(2);
    if(indoc2 != null && indoc2.len() == 2){
     Document indoc21 = indoc2.getPart(1);
     Document indoc22 = indoc2.getPart(2);
     if(indoc21 != null && indoc21.is_intdoc() &&
        indoc22 != null && indoc22.is_noncedoc()){
      return indoc2;}}}}
  return null;}}
```

The method step() is the top-level method for executing a protocol step. First it tests whether input is available. If there is an input available the receive method is executed and phonestep() is called with the input. This method now tests the structure of the input using getTicket() method. getTicket() returns the data part of the input document if it was a valid representation of a ticket and null otherwise. phonestep() then adds the returned data to the list of actual tickets if the input was valid.

The proof structure now is the following: Starting with the proof obligation given by the refinement theory in Sect. 3 we first symbolically execute the abstract and the concrete level. The cases for the non-refined agents (such as the

---

[2] This source code is running on any J2ME mobile phone. We have tested it on Nokia 3250 and Sony Ericsson W550i. The receive operation uses the J2ME API to access the MMS messages of the mobile phone.

attacker) are trivial because they are the same in both specifications. For the refined agent we come to the proof obligation shown in Sect. 3.3. We then formulate theorems for each Java method which relate the behavior of the method to the abstract counterpart of its input. The corresponding theorem for the load-ticket protocol step is for example:

is_load_message(first(inputs(agent))) $\land$ st_1 = store(agent) $\land$
st = TOSTORE(inputs, st_1) $\land$ INV(st_1) $\land$ $\ldots$
$\rightarrow \langle st; \ \text{Protocol.step}(); \ \rangle$
    (getDoc(st[Protocol.tickets], st) =
        tickets(agent) + first(inputs(agent))
    $\land$ st[.input] = null $\land$ INV(st))

If the actual input document (first(inputs(agent))) is a correct load message (is_load_message) on the abstract level and if this document is added to the store via TOSTORE then the step method performs the correct state change: It computes the correct ticket list (the new ticket attached to the old tickets). Also the input was deleted (st[.input] = null). Additionally an invariant that holds before the execution of the method (INV(st)) holds again afterwards.

With such theorems the refinement proof obligation is divisible in different proof obligations for every protocol step. After applying those theorems we symbolically execute the corresponding abstract ASM step. This results in an updated abstract state which has to be proven to relate to the Java store which is given by the theorem above via retrieve relation R. Using this technique the whole proof becomes feasible. The whole case study consists of around 1000 lines of code. The implementation of Cindy itself consists of around 350 lines of code. The rest is the implementation of the document classes and some utility classes (e.g. for handling byte arrays). The verification of the refinement starting with the creation of the concrete and abstract specification of the protocol and ending with the refinement proof took around one and a half man months with KIV. The case study consists of 329 theorems which took 11408 proof steps. 4655 of those steps were done by the user. The degree of automation thereby is nearly 60 %. We expect a much higher degree of automation for upcoming case studies because of the high re-usability of the Document implementation and the corresponding library.

7 Related Work

Related work concerning the verification of Java programs was already mentioned in Section 1. Here we focus on related work concerning refinement approaches for security protocols:

[20] describes a similar approach for Java Smart Cards. The authors specify protocols using a high level specification language for proving security properties and a more concrete one which works on the level of byte arrays. They specify lengths and contents of messages using byte arrays and then use static program analysis on the JavaCard implementation to decide whether the implementation is correct. This approach is limited to the very specific class of protocols the

specification language allows while our approach allows any abstract specification using all the possibilities of algebraic specifications on KIV [18]. Additionally, because of the automated analysis and the fact that implementation correctness is undecidable this approach cannot give reliable answers in every case.

[29] uses the Spi Calculus for specifying security protocols and a code generation engine to transform this specification to an implementation, also mapping abstract messages to Java objects. Code generation yields large implementations that are less readable than our code and cannot be optimized without losing correctness guarantess. Their mapping to concrete data types is not formally verified and does not address the problem of invalid inputs on the concrete level.

[14] presents an approach to verify that a JavaCard implementation respects a protocol specification given by a finite state machine. This approach cannot directly transfer security proofs from the abstract specification to the implementation level, because they basically show that the Java program sends certain message types in the right order but do not show that those messages and the internal state of the implementation have the right contents.

The Mondex [21] case study has recently received a lot of attention because its tool supported verification has been set up as a challenge for today's verification tools [31]. The original refinement proofs using Z have been done on a very detailed level by hand [28]. In [25] and [24] we show that the same verification can be done with good tool support and in a short period of time using KIV. An extension of Mondex using our PROSECCO approach can be found in [12]. The Mondex refinement basically splits a world view of an application into components implementing a protocol. But even the lowest level of the Mondex case study is a only an abstract specification of the communication protocol of the involved parties that does not contain cryptographic operations. The approach presented here can be used to do an additional refinement for Mondex adding a real implementation. Details on our implementations of Mondex can be found in [11].

8   Conclusion

We presented a refinement method for Java programs instantiating data refinement. The method is based on a calculus for Java verification and Abstract State Machines using the interactive theorem prover KIV. While the approach is not bounded to KIV only and the method itself could be transfered to other Java verification systems, KIV's strong support for ASM verification, Java verification and algebraic specifications as well as its large library for security protocol verification makes it an efficient tool for this approach.

As discussed in Sect. 3 our approach transfers security properties for the abstract specification down to running Java code. Furthermore, we have shown how to handle invalid inputs that only exist on the concrete level of Java pointer structures. We have demonstrated that the method is suitable for handling case studies of relevant size. Further work includes the incorporation of the method into further verification case studies like Mondex.

References

1. Michael Balser, Wolfgang Reif, Gerhard Schellhorn, Kurt Stenzel, and Andreas Thums. Formal system development with KIV. In T. Maibaum, editor, *Fundamental Approaches to Software Engineering*. Springer LNCS 1783, 2000.
2. *Tickets on your Mobile*. URL: `http://www.beep.nl` [last seen 2007-03-16].
3. C. Bolton, J. Davies, and J.C.P. Woodcock. On the refinement and simulation of data types and processes. In K. Araki, A. Galloway, and K. Taguchi, editors, *Proceedings of the International conference of Integrated Formal Methods (IFM)*, pages 273–292. Springer, 1999.
4. Egon Börger and Robert F. Stärk. *Abstract State Machines—A Method for High-Level System Design and Analysis*. Springer-Verlag, 2003.
5. C. Breunesse, B. Jacobs, and J. van den Berg. Specifying and verifying a decimal representation in Java for smart cards. In *Proceedings AMAST 2002*, Reunion Island, France, 2002. Springer LNCS 2422.
6. L. Burdy, Yoonsik Cheon, David Cok, Michael Ernst, Joe Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of jml tools and applications. In Thomas Arts and Wan Fokkink, editors, *Eighth International Workshop on Formal Methods for Industrial Critical Systems (FMICS '03)*. Volume 80 of Electronic Notes in Theoretical Computer Science, Elsevier, 2003.
7. Michael Burrows, Martín Abadi, and Roger M. Needham. A Logic of Authentication. Technical report, SRC Research Report 39, 1989.
8. W. de Roever and K. Engelhardt. *Data Refinement: Model-Oriented Proof Methods and their Comparison*, volume 47 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1998.
9. J. Derrick and E. Boiten. *Refinement in Z and in Object-Z : Foundations and Advanced Applications*. FACIT. Springer, 2001.
10. Holger Grandy, Dominik Haneberg, Wolfgang Reif, and Kurt Stenzel. Developing Provably Secure M-Commerce Applications. In Günter Müller, editor, *Emerging Trends in Information and Communication Security*, volume 3995 of *LNCS*, pages 115–129. Springer, 2006.
11. Holger Grandy, Nina Moebius, Markus Bischof, Dominik Haneberg, Gerhard Schellhorn, Kurt Stenzel, and Wolfgang Reif. The Mondex Case Study: From Specifications to Code. Technical Report 2006-31, University of Augsburg, December 2006. URL: `http://www.informatik.uni-augsburg.de/lehrstuehle/swt/se/publications/`.
12. Dominik Haneberg, Gerhard Schellhorn, Holger Grandy, and Wolfgang Reif. Verification of Mondex Electronic Purses with KIV: From Transactions to a Security Protocol. Technical Report 2006-32, University of Augsburg, December 2006. URL: `http://www.informatik.uni-augsburg.de/lehrstuehle/swt/se/publications/`.
13. He Jifeng, C. A. R. Hoare, and J. W. Sanders. Data refinement refined. In B. Robinet and R. Wilhelm, editors, *Proc. ESOP 86*, volume 213 of *Lecture Notes in Computer Science*, pages 187–196. Springer-Verlag, 1986.
14. Engelbert Hubbers, Martijn Oostdijk, and Erik Poll. Implementing a Formally Verifiable Security Protocol in Java Card. In *Proceedings of SPC'2003, First International Conference on Security in Pervasive Computing*, number 2802 in LNCS, Boppard, Germany, 2004. Springer.
15. M. Huisman. Verification of java's abstractcollection class: a case study. In *MPC'02: Mathematics of Program Construction*. Springer LNCS 2386, 2002.

16. Bart Jacobs, Claude Marche, and Nicole Rauch. Formal verification of a commercial smart card applet with multiple tools. In C. Rattray, S. Maharaj, and C. Shankland, editors, *Algebraic Methodology and Software Technology (AMAST) 2004, Proceedings*, Stirling Scotland, July 2004. Springer LNCS 3116.

17. Bill Joy, Guy Steele, James Gosling, and Gilad Bracha. *The Java (tm) Language Specification, Second Edition*. Addison-Wesley, 2000.

18. KIV homepage. `http://www.informatik.uni-augsburg.de/swt/kiv`.

19. H. Mantel. Preserving Information Flow Properties under Refinement. In *Proceedings of the IEEE Symposium on Security and Privacy,*, Oakland, CA, USA, 2001.

20. R. Marlet and D. Le Metayer. Verification of Cryptographic Protocols Implemented in JavaCard. In *Proceedings of the e-Smart conference (e-Smart 2003)*, Sophia Antipolis, 2003.

21. MasterCard International Inc. *Mondex*. URL: `http://www.mondex.com`.

22. W. Mostowski. Rigorous development of java card applications. In T. Clarke, A. Evans, and K. Lano, editors, *Proceedings, Fourth Workshop on Rigorous Object-Oriented Methods*, London, U.K., 2002.

23. L. C. Paulson. The Inductive Approach to Verifying Cryptographic Protocols. *J. Computer Security*, 6, 1998.

24. Gerhard Schellhorn, Holger Grandy, Dominik Haneberg, Nina Moebius, and Wolfgang Reif. A systematic verification Approach for Mondex Electronic Purses using ASMs. Technical Report 2006-27, Universität Augsburg, 2006. URL: `http://www.informatik.uni-augsburg.de/lehrstuehle/swt/se/publications/`.

25. Gerhard Schellhorn, Holger Grandy, Dominik Haneberg, and Wolfgang Reif. The Mondex Challenge: Machine Checked Proofs for an Electronic Purse. In J. Misra, T. Nipkow, and E. Sekerinski, editors, *Formal Methods 2006, Proceedings*, volume 4085 of *LNCS*, pages 16–31. Springer, 2006.

26. Kurt Stenzel. A formally verified calculus for full Java Card. In C. Rattray, S. Maharaj, and C. Shankland, editors, *Algebraic Methodology and Software Technology (AMAST) 2004, Proceedings*, Stirling Scotland, July 2004. Springer LNCS 3116.

27. Kurt Stenzel. *Verification of Java Card Programs*. PhD thesis, Universität Augsburg, Fakultät für Angewandte Informatik, URL: `http://www.opus-bayern.de/uni-augsburg/volltexte/2005/122/`, or `http://www.informatik.uni-augsburg.de/forschung/dissertations/`, 2005.

28. S. Stepney, D. Cooper, and J. Woodcock. AN ELECTRONIC PURSE Specification, Refinement, and Proof. Technical monograph PRG-126, Oxford University Computing Laboratory, July 2000. URL: `http://www-users.cs.york.ac.uk/~susan/bib/ss/z/monog.htm`.

29. B. Tobler and A. Hutchison. Generating Network Security Protocol Implementations from Formal Specifications. In *CSES 2004 2nd International Workshop on Certification and Security in Inter-Organizational E-Services at IFIPWorldComputerCongress*, Toulouse, France, 2004.

30. J. C. P. Woodcock and J. Davies. *Using Z: Specification, Proof and Refinement*. Prentice Hall International Series in Computer Science, 1996.

31. Jim Woodcock. First steps in the verified software grand challenge. *IEEE Computer*, 39(10):57–64, 2006.