

Enhancing Java Runtime Environment for Smart Cards Against Runtime Attacks

No Author Given

No Institute Given

Abstract. Smart cards are mostly deployed in security-critical environments in order to provide a secure and trusted access to the provisioned services. These services are delivered to a cardholder using the Service Provider's (SPs) applications on his or her smart card(s). These applications are at their most vulnerable state when they are executing. There exist a variety of runtime attacks that can circumvent the security checks implemented either by the respective application or the runtime environment to protect the smart card platform, user and/or application. In this paper, we discuss the Java Runtime Environment and a potential threat model based on runtime attacks. Subsequently, we discussed the counter-measures that can be deployed to provide a secure and reliable execution platform, along with an evaluation of their effectiveness, incurred performance-penalty and latency.

1 Introduction

An application on a smart card relies on the Smart Card Runtime Environment (SCRT) for secure and reliable execution. An SCRT contains a library of Application Programming Interfaces (APIs) that provide a secure and reliable interface between the installed applications and on-card services. An SCRT is used in order to:

1. Provide a secure and reliable program execution.
2. Enforce an execution isolation and access to memory locations.
3. Provide an interface to access cryptographic algorithms.
4. Protect the platform and applications from malicious or ill-formed applications.
5. Handle communication between applications and with external entities.

In early 2000, fault attacks became the modus operandi of adversaries to subvert the implemented cryptographic algorithms in the smart card industry. Since then the technology has evolved to counter these threats to some extent [3–5]. There has been a growing interest in fault injection and combined attacks [6–8] to subvert the protection mechanisms on a smart card. In combined attacks both the software (i.e. attacker's application) and fault injection are used to achieve the objectives. In this paper, we analyse the attacks that target the SCRT and provide counter-measures. The attacks we have considered in this paper are fault and combined attacks targeted at the SCRT. In this paper, we focus on the Java Cards; therefore, we will constantly refer to the Java Card Runtime Environment (JCRE) and it is used synonymously with SCRT. The rationale is that the JCRE has an open specification as compared to Multos, and new attacks mostly target Java Cards.

1.1 Contributions of the Paper

In this paper, we propose and evaluate the following:

1. A JCRE protection framework referred as “Runtime Protection Mechanism (RPM)”.
2. Inclusion of application developer’s security requirements at the compilation of the application. If these requirements do not violate the security requirements of the JCRE, the runtime environment will try to enforce them.
3. A set of countermeasures that include:
 - (a) Operand Stack Integrity: Safeguarding the JCRE’s operand stack from any malicious modifications.
 - (b) Permitted Execution Path Analysis: Evaluate the program flow and verify whether a particular execution path is allowed or not, based on the security requirements; defined by the application developer and/or JCRE.
 - (c) Bytecode Integrity: To verify and validate whether the execution code of an application, in storage (persistent memory) and while on the non-persistent memory during its execution has not been modified.
4. Two variants of “Runtime Security Manager (RSM)” referred serial and parallel mode. The RSM enforces the security requirements defined by the application developers and JCRE as part of the RPM along with deploying the countermeasures. The variants are differentiated based on the architecture of the underlying hardware and the point at which the RSM verify and validate an application during its execution.
5. The proposed framework is implemented, and evaluated for security, incurred performance penalty and latency.

1.2 Structure of the paper:

We begin the discussion with a brief introduction of the JCRE, related work on JCRE security and our motivation for the paper. In section 3, we describe the attacker capabilities and security requirements for a secure JCRE. In this section, we also propose our framework. In section 4, we analyse the proposed framework, which includes security, incurred performance-penalty, latency and comparative analysis. Finally, the paper is concluded in section 5.

2 Smart Card Runtime Environment

In this section, we open the discussion with a brief description of the Java Card Virtual Machine (JCVM) followed by related work, and our motivation for the paper.

2.1 Java Card Virtual Machine

The JCRE consists of APIs, system classes, Java Card Virtual Machine (JCVM), and native methods. The most crucial component of the JCRE is the JCVM that actually interpret the application code to execute on the underlying hardware. The architecture of the JCVM is more or less similar between various Java Card versions.

An application is coded in a subset of Java language that is supported by the JCVM, which is represented as a Java file. The application is then compiled into a

class file, and it is packaged along with any resource files and supporting libraries into an installation package (e.g. CAP, or JAR file [9, 10]) that can be downloaded to a Java Card. On the Java Card, the on-card bytecode verifier would analyse the downloaded application and validate that it conforms to the Java language semantics.

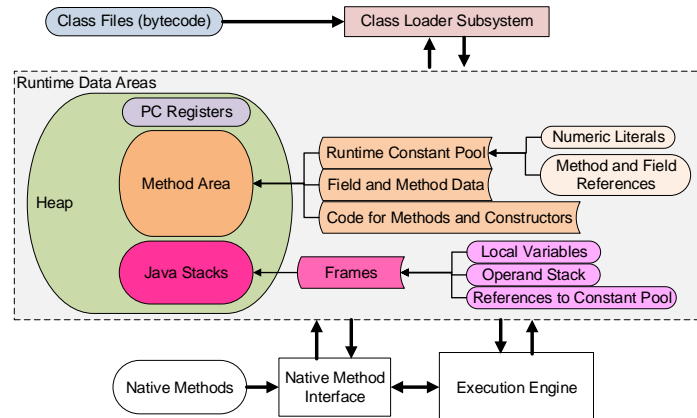


Fig. 1. Architecture of the Java Card Virtual Machine

Figure 1 illustrates the architecture of a typical JCVM. Various components and their functions are described subsequently with emphasis on how they interact during the execution of an application.

The JCVM mainly deals with an abstract storage unit called word that is the smallest storage unit that it can process. The actual size of a word is left to the JCVM implementers. However, the JCVM specification [9] states that a word should be large enough to hold a value of **byte**, **short**, **reference**, or **returnAddress**.

When an application is initiated, the bytecode representation of an application is loaded to the JCVM memory by a “class loader subsystem”. The class loader is responsible for locating and loading the class onto the memory areas used by the JCVM. This memory is divided into sub-areas, where each of them contains specific information regarding the application. The JCVM memory area is heap and all data/code related to an application is loaded onto it. The three main storage structures defined on the heap that are of relevance here are Program Counter (PC) registers, method area, and Java stacks. These storage structures are briefly discussed here as they are referred to in the remaining paper (i.e., when we discuss our proposed counter-measures).

The PC registers store the memory address of the bytecode instruction currently executing. If the JCVM supports multiple threading then each thread will have its own PC register.

The method area is a memory space that consists of structures that include runtime constant pool, field and method data, and code related to methods and constructors. The runtime constant pool stores the constant field values (e.g. numeric literals) and references to the memory address related to methods and fields. The other two structures store the data and code related to fields and methods, etc.

A frame is created by the JCVm each time a method is invoked during the execution of an application. A frame is a construct that stores data, partial results, return values, and dynamically resolved links, associated with a single method (not the related class). These frames are stored on a last-in first-out (LIFO) stack called Java Stack. For each thread, there will be a different Java Stack. For security reasons, only the JCVm can issue the push and pop instructions to Java Stacks. The data structures that reside on a frame include an array of local variables, operand stack, and references to constant pool. The operand stack is a LIFO stack and it is empty when a frame is created. During the execution of a method, the JCVm will load data values (of either constant or non-constant variables/fields) onto the operand stack. The JCVm will operate on the values at the top of the operand stack and push the results back on it.

The JCVm provide well-defined interfaces to access native methods; however, contrary to traditional Java virtual machines they do not allow user-defined native methods. Each JCVm has an execution engine that is responsible for the execution of the individual instructions (opcodes) in an application code. The design of the execution engine is dependent on the underlying hardware platform and in a simple way, it can be considered as a software interface to the platform's processor.

This section do not exhaustively explain JCRE and the rationale for covering the aforementioned topics is to make it easy to follow the subsequent discussion in the paper.

2.2 Related Work on JCRE Security

Earlier work on Java Cards was mainly related to the semantic and formal modelling of the JCVm [11–14], Java Card firewall mechanism [15, 16], and applets [17–19]. The JCRE countermeasure against ill-formed applications was based on on-card bytecode verification [20–23], which became compulsory in the Java Card version 3 [9].

In the early 2000s, side channel analysis and fault attacks on smart card platforms were mainly focussed on the cryptographic algorithms [2, 24–28]. However, in recent years, logical and fault attacks are combined to target the JCRE [29–31].

In 2008, Mostowski and Poll [32] loaded an ill-typed bytecode on various smart cards to test their security and reliability mechanisms. They also noted that smart cards that had an effective on-card bytecode verifier were less susceptible than others. In 2009, Hogenboom and Mostowski [33] managed to read arbitrary contents of the memory. They performed this attack even in the presence of the Java Card firewall mechanism. Similar results were also shown by Lanet and Iguchi-Cartigny [34]. Sere et al. [35] use the similar attack of modifying the bytecodes to gain unauthorised access or skip the security mechanism on a platform. However, Sere et al. relied on fault attacks to modify the bytecodes rather than modifying them off-card as done by [32–34]. This way, Sere et al. managed to bypass the on-card bytecode verification. A countermeasure to this attack provided by Sere et al. relied on tagging the bytecode instructions with integrity values (i.e. integrity bits) and during the execution, the JCVm checks these bits and if it fails, the execution terminates.

In 2010, Barbu et al. [7] along with Vétillard and Ferrari [6] used a similar attack methodology to Sere et al. [35] that later came to be known as combined attacks. Later, the combined attack technique was extended to target various components of

JCVM in [36–39]. These attacks are significant; nevertheless, they require the loading of an application designed specifically to accomplish the attack goals.

Dubrile et al. [54] discussed the fault enabled mutants in the Java Cards and proposed a countermeasure based on the typed stack. In [55] Julien Lancia illustrated a combined attack on the memory references (object and variable references) and proposed a countermeasure based on a defensive virtual machine.

The discussion in this section is by no means exhaustive but it accurately introduces the challenges faced by the JCRE.

2.3 Motivation

During an application’s lifetime, the application’s security is dependent on the security of the runtime environment. As discussed in section 2.2, a smart card runtime environment is increasingly facing the convergence of various attack techniques (e.g. fault and logical attacks). Although, physical protection mechanisms regarding fault attacks are proposed [40]; however, we consider that the necessary software protection for the runtime environment cannot be understated. The software protection can augment the hardware mechanism to protect against the combined attacks, as a similar approach has yielded successful results in the secure design of cryptographic algorithms for smart cards [41–43]. Therefore, in this paper, we will focus on the software protection mechanism.

In literature, several methods are described for software protection mechanism, including application slicing in which an application is partitioned for performance [44, 45] or to protect intellectual property [46]. Such partitioning can be used to tag individual segments of an application with adequate security requirements. The runtime environment can then take into account the security requirements, tagged with individual segments during the execution; thus providing configurable runtime security architecture. A similar approach is proposed by Java Card 3 [9] and as part of the counter-measures to combined attacks proposed by Sere et al. [39] and Bouffard et al. [47]. These proposals are based on using Java annotations to tag segments of an application with required security or reliability levels.

Developers can use Java annotations to provide information regarding an application (or its segment), which is used by either the compiler, or runtime environment (i.e. JCVM). Based on Java annotations, Bouffard et al. [47] and Sere et al. [39] proposed mechanisms to prevent control flow attacks. In addition, Loining et al. [48] used the Java annotations to ensure a secure and reliable development of applications for embedded devices (e.g. smart cards). Furthermore, Java Card 3 Connected Edition also makes provision for Java annotations [9]. The defined annotations by Java Card 3 are integrity, confidentiality, and full (which corresponds to both integrity and confidentiality). In addition, the specification also allows proprietary annotations that can be used to invoke specific protection mechanisms implemented by the respective card manufacturer. The Java Card 3 specification does not detail what operations a JCVM should perform when encountering a particular annotation, which are left to the discretion of the card manufacturers.

These proposals are useful but a malicious user can use the annotations to his advantage in order to accomplish his malicious goals. To avoid this we in our proposal have an on-card analyser that checks the security and reliability requirements of an application, validate the associated Java annotations (tags) with each segment

of the application, and modify the security annotations where adequate. In such a scenario, we may assume that tagging segments of an application with security annotations might be useful. Nevertheless, such an on-card analyser is not available on smart cards. In this paper, we solely focus on adequately hardening the runtime environment.

In our proposed framework, we tackle the problem from three aspects: application compilation, runtime protection, and trusted component. The Java annotations are used to tag properties of individual segments of an application. Runtime commands (opcodes) that might be subverted to gain unauthorised access are hardened with additional protection (security checks), and finally a trusted component is included to complement the runtime environment for security verification and validation of an application's execution.

3 Runtime Protection Mechanism

In this section, we describe the attacker's capability along with the security requirements for an reliable and safe JCRE. Subsequently, we discuss the proposed runtime protection mechanism and how it provides a secure and reliable framework for JCRE.

3.1 Attacker's Capability

Due to the advancement in the chip technology and hardware protection mechanisms [49], we have taken a realistic approach in defining the attacker's capability, taking into consideration the current state-of-the-art in attack methodologies for smart cards. The attacker's capabilities taken into consideration in the proposed runtime protection mechanism are listed as below:

1. Has the knowledge of the underlying (hardware and software) architecture.
2. Has the ability to load a customised application onto a given smart card.
3. Has the capability to induce a fault attack at a precise clock cycle.
4. Has the limited capability of changing a byte value to either 0x00 or 0xFF, or a random value in between.
5. Has the potential to change values stored in a non-volatile memory permanently within the limits of the capability four.
6. Has the ability to inject multiple faults; however, only in serial fashion (i.e. after injecting a fault, the attacker waits for the results before injecting the next fault). The adversary cannot inject multiple faults in parallel — injecting two faults simultaneously.
7. Can overwrite a whole or part of memory like Electrically Erasable Programmable Read-Only Memory (EEPROM) or off-card storage.

Capability four restricts an adversary to induce a precise byte error rather than the precise bit error. This restriction is based on the underlying smart card hardware architecture. This is not to say that precise bit errors are not possible in smart cards. On the contrary, they are technically possible but increasing the density of packaging (i.e. chip fabrication) makes it challenging to change a value of a bit in comparison to changing the value of a byte.

The rationale behind the choice of multiple fault attacks in serial fashion than parallel is to give precise control and reproducibility of the attack. In fault attacks where a malicious user injects multiple faults simultaneously (parallel), it is difficult to assess whether the first fault injection was successful; therefore, injecting the second fault may be less productive.

In our proposed framework, we intend to protect the underlying runtime environment and applications hosted on it. However, an application can be designed in a way that it reveals its sensitive data (e.g. its user’s private key that is specific to the application and not related to the platform or other applications); there is a limit to what a protection mechanism can do to prevent such leakages.

3.2 Security Requirements for a Runtime Protection Mechanism

In this section, we discuss the set of requirements appropriate for a runtime protection mechanism to defend against the attacker discussed in the previous section.

1. Customisable: Enables the application developers to define the security requirements (if preferred) for their applications, which will be enforced as long as they do not violate the platform’s and/or other application’s security requirements.
2. Developer Independent: Does not require the application developers to evaluate the security risks of their application and adequately tag it.
3. Code Integrity: Detect any unauthorised modification to the application code before it is executed.
4. Stack Integrity: Detect any modification to the values stored on the Java stacks (e.g. operand stack).
5. Execution Flow Evaluation: Detect any illegal jumps to either restricted areas (e.g. data or code locations for a different application) or violating the secure execution flow of the application.

These requirements are revisited in section 4.5, when our proposal is compared with the existing proposals discussed in section 2.2.

3.3 Overview of the Proposed Runtime Protection Mechanism

The proposed architecture of the runtime protection mechanism is involved at various stages of the application lifecycle - including the application compilation, on-card bytecode verification, and execution as shown in figure 2.




Fig. 2. Generic overview of the proposed runtime protection mechanism

During the compilation/packaging process additional information regarding individual methods, classes, and objects of an application is generated as part of the

property file, discussed in section 3.4. The property file assists the runtime environment to provide a security and reliability service during the execution of the application. The off-card bytecode verification checks whether the downloaded application conforms to the (given) language's semantics. The on-card bytecode verifier can also request the trusted component to validate the property file. The trusted component is the proposed Runtime Security Manager (RSM) discussed in section 3.6 that actively enforces the security and reliability policy of the platform - taking into account the information included in the property file.

The proposed framework does not require that application developers perform security assessment of their application(s) to adequately tag application segments. The framework only requires at minimum that developers compile their applications in a way that it has a property file that stores information related to the respective application. The second requirement of the proposed framework is to adequately harden the runtime environment discussed in section 3.5 along with introducing the RSM that will enforce the platform security policy (section 3.6).

In subsequent sections, we will extend the generic architecture discussed in this section and explain how these different components come together.

3.4 Application Compilation

A Java compiler will take a Java file and convert it to a (bytecode) class file. The class file not only has opcodes, but it also includes information about various segments (e.g. methods, and classes) of an application that is necessary for the JCVM to execute the application. However, for our proposal we introduce a property file that includes additional information about an application. If a JCVM knows how to process property files then it will proceed with them; otherwise, it will silently ignore them. In our proposal a property file is stored and used by the RSM during the execution of the associated application. In order to integrate the RSM into the runtime environment, the JCVM is required to be modified so it can communicate with the RSM in order to safeguard the execution environment.

```
1 ApplicationInfo{
2     Application_Identifier ApplicationIdentifier;
3     ClassInformation ClassInfo[class_count];}
4 ClassInfo{
5     Class_Identifier ClassIdentifier;
6     MethodInformation MethodInfo[method_count];}
7 MethodInfo{
8     Method_Identifier MethodIdentifier;
9     MethodIntegrity HashValue;
10    PermittedExecutionPath Path[jumps_count]}
```

Listing 1.1. Property file structure of a Java Card application.

The property file contains security and reliability information concerning an application that the runtime environment can utilise to execute an application. The structure of the property file is illustrated in listing 1.1, which includes information regarding the permitted execution-paths, and integrity matrix (hash values of the non-mutable part of the individual methods in a class).

The `ApplicationInfo` data structure includes the application identifier (e.g. AID) and an array of classes that are part of the respective application. For each class in the application, we have a `ClassInfo` structure that contains the `MethodInformation`

array that contains information regarding all methods associated with the given class. Each method is represented by the `MethodInfo` structure that includes the permitted execution-paths that are generated for each method. In the permitted execution-paths, child nodes represent jumps to other methods, irrespectively of whether they are from the same application or from a different application. In a way, combining the method paths of all classes can give the complete permitted execution-path of the respective application. In addition to the permitted execution-paths, a `MethodInfo` also contains the hash value (of non-mutable code) of the respective method. This hash value can be generated at the compile time and added to the property file, or at the time of the application installation: the RSM calculates the hash value and stores it in the property file.

3.5 Execution Environment

The runtime environment is modified to support the inclusion of the RSM that is shown in figure 3. At the time of application installation, the application bytecode is stored in the respective SP's domain along with the associated property file. The property file is sealed¹ so that neither the application nor an off-card entity (e.g. an SP or/and adversary) can modify it without detection. At the time of execution, the RSM will retrieve the file, verify the integrity of the file, and then decrypt it. If an SP wants to update its application then it will proceed with the update command² that will notify the RSM of the update. At the completion of the update, the RSM will verify the application security certificate (if available), and update the property file – if required.

3.6 Runtime Security Manager

The purpose of the RSM is to enforce the security counter-measures (section 3.7) defined by the respective platform. To enforce the security counter-measures, the RSM has access to the heap area (e.g. method area, Java stacks) and it can be implemented as either a serial or a parallel mode.

A serial RSM will rely on the execution engine of the JCVM (figure 1) to perform the required tasks. This means that when an execution engine encounters instructions that require an enforcement of the security policy, it will invoke the RSM that will then perform the checks. If successful the execution engine continues with execution, otherwise, it will terminate. A parallel RSM will have its own dedicated hardware (i.e. processor) support that enables it to perform checks simultaneously while the execution engine is executing an application. Note that having multiple processors on a smart card is technically possible [50]. The main question regarding the choice is not the hardware, but the balance between the performance and latency.

Performance, as the name suggests is concerned with the computational speed. Whereas, latency deals with the number of instructions executed between an injected-error to the point it is detected. We will return to this discussion later in section 4 where we provide test (simulated) implementation results.

¹ Sealed: The data is encrypted (authenticated encryption) by the RSM storage key.

² Update Command: We do not propose any update command in this paper but similar commands are defined as part of the GlobalPlatform card specification. The update command enables an authorised entity (e.g. SP) to modify an application.

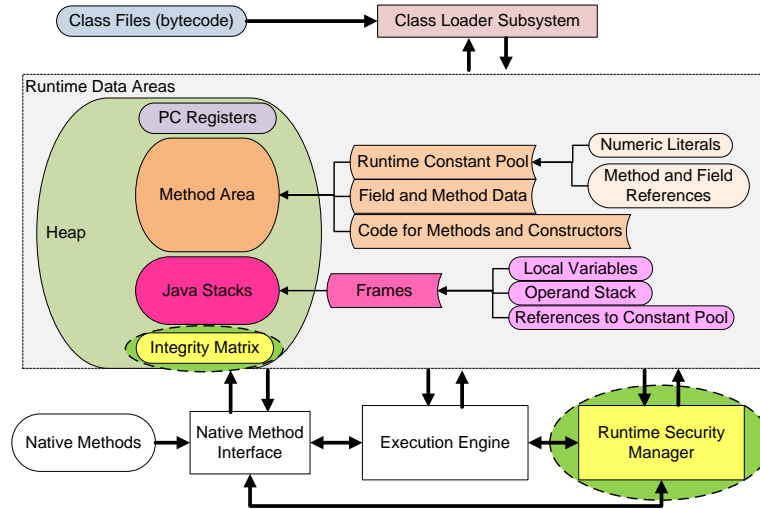


Fig. 3. Architecture of the Proposed Runtime Environment for COM Devices.

3.7 Runtime Security Counter-Measures

The RSM along with the runtime environment would apply the required security counter-measures (as part of the runtime protection mechanism) that are discussed in subsequent sections.

Operand Stack Integrity As discussed in section 2.1, an operand stack is part of the Java stacks and they are associated with individual Java frames (methods). During the execution of an application, the runtime environment pushes and pops local variables, constant fields, and object references to the operand stack. The instructions specified in an application can then process the values at the top of the stack. Barbu et al. [36] showed that a fault injection that changes the values stored on the operand stack could have adverse effect on an application's security. Furthermore, they also provided three different counter-measures to the proposed attack.

The proposed countermeasure (second-refined method) of Barbu et al. [36] is based on the idea of operand stack integrity. They define a variable α , and all values that are pushed on or popped from the operand stack are XORed with the α . On every jump instruction beyond the scope of the current frame (method), the runtime environment XORs all the values stored on the operand and compares the result with α . If they match then the integrity of the operand stack is verified. Their proposal does not measure the integrity of the operand stack on instructions like if-else or loops, which could be the target of the malicious user. In their proposed counter-measures they sacrificed security and (to some extent) performance for the sake of memory use, whereas our proposal focuses on security rather than saving the memory.

In our proposal, we use a Last In First Out (LIFO) stack referred as integrity stack. One thing to note is that JCVm knows the size of the operand stack when it loads a frame (section 2.1); therefore, the RSM just creates an integrity stack of the

size n where n is the size of the respective operand stack. We refer to the integrity stack as “InS” in listing 1.2.

When a frame is loaded, the JCVM and the RSM will create an operand and integrity stack, respectively. Furthermore, the RSM will also generate a random number and stores it as S_r . The rationale for using the random number will become apparent in the subsequent discussion.

```

1 // Executed by RSM when a value is pushed onto an integrity stack.
2 On_Stack_Push(pushValue){
3   push(InS[top] XOR pushValue);}
4 // Executed by RSM when a value is popped from an operand stack.
5 On_Stack_Pop(poppedValue){
6   if(pop(InS) XOR poppedValue := InS[top]){
7   } else {
8     terminateExecution();
9   }}

```

Listing 1.2. Operand stack integrity operations.

When a value V_i is pushed to the operand stack, if it is the first value on the stack then the value pushed on the InS will be $I_i = V_i \oplus S_r$. For all subsequent values (where $i > 1$) the values pushed on the InS will be $I_i = V_i \oplus V_{i-1}$.

The rationale for using a random number is to avoid parallel fault injections that try to change the values on both operand and integrity stack simultaneously. Such a parallel fault injection will become difficult if an adversary cannot predict the values stored on the integrity stack, as each value on the integrity stack will be chained with the generated random number.

When a value is popped out of the operand stack, we also pop the integrity value from the integrity stack, XOR it with the popped value from the operand stack and compare it with the new top value on the integrity stack. If the values match then the integrity of the popped value from the operand stack is verified; otherwise, it has been corrupted and the RSM requests the JCVM to terminate the execution as shown in listing 1.2.

The RSM will continuously monitor the integrity of the operand stack, in comparison to the Barbu’s proposal. Furthermore, in this proposal the validation does not require the calculation of integrity value over the entire operand stack. If we take the Barbu’s proposal then for an operand stack of length ‘ n ’, we have to perform “ $n-1$ ” XOR operations every time we need to verify the state of the operand stack. However, in our proposal we only need to perform one XOR operation. We sacrifice the memory for the sake of performance in our proposal. We consider that operand stacks are not large data structures so even if we double the memory used by them, it will not have an adverse effect on the overall memory usage.

Permitted Execution Path (PEP) Analysis In our proposal, we are concerned with jumps that refer to external resources. The term external resources in the context of PEP analysis means any jump that goes beyond the scope of the current Java frame (i.e. method) while it is still on the Java stack. Once a method completes its execution, the JCVM will remove the associated Java frame from the Java stacks (figure 1). Examples of such jumps defined in Java virtual machine specification [51] are `invokeinterface`, `invokestatic`, `invokevirtual`, `areturn`, etc.

```

1 byte B(byte inputValue){
2   byte a = 1;
3   if(inputValue != a){ C(inputValue);
4   }else{D(inputValue)}
5   return SG(inputValue);}

```

Listing 1.3. Code for an example method B.

To explain the Permitted Execution Path Analysis further, we consider an example method B that has three jumps before it reaches the return statement that completes the execution of the method. The method B’s code is shown in listing 1.3. Each invocation of a method (e.g. C, D, and SG) is represented by a symbolic method name (i.e. alphanumeric form that is easily readable/recognisable by humans) that has an associated unique byte sequence referred as method identifier. For example, unique method identifier of methods B, C, D, and SG are 0xF122, 0xF123, 0xF124, and 0xF125, respectively. For explanation we have used method identifiers that consist of two bytes. Along with the method identifier the property file also includes `PermittedExecutionPath`, which is a set of PEPs sanctioned for the given method.

The `PermittedExecutionPath` in the property file (listing 1.1) is simply constructed by taking into account every possible (legal) execution path of a method. Taking the example method B, the first jump can either be to method C or D depending upon the input. The construction of the `PermittedExecutionPath` (set of legal jumps) is constructed by XORing the method identifiers of individual jumps.

The PEP analysis requires that the RSM have a PEP variable “*cfa*” that stores the path taken by an application as $cfa = \sum_{j=1}^n C_j$. Where C_j represents the jumps taken during execution of an application. During the execution of a method, when the JVM encounters a jump to another method the RSM XORs the method identifier with the current value of “*cfa*” and lookup the `PermittedExecutionPath` of the given method in the associated property file. If it finds a matching value, the JVM will proceed with the execution; if not it will terminate the execution.

Our scheme also deals with the loop instructions that contain jumps to multiple methods depending upon the loop condition. For example, for an odd value of ‘i’ jump to method B and for even values jump to method C. The loop iterates through the values of ‘i’ until it meets the condition that might be based on runtime values (i.e. unpredictable at the time of the compilation of the application). Consider a control flow graph of four methods: A, B, C, and D. Methods B and C are part of a loop as discussed before and illustrated in listing 1.4. The `PermittedExecutionPath` set will be $A_{cfa-Set} = \{A \oplus B, A \oplus C, A \oplus D, A \oplus B \oplus D, A \oplus C \oplus D, A \oplus B \oplus C \oplus D\}$. For a potential execution path $A \rightarrow B \rightarrow C \rightarrow B \rightarrow C \rightarrow B \rightarrow C \rightarrow D$, if we compute the “*cfa*” it would be $A \oplus B \oplus C \oplus B \oplus C \oplus B \oplus C \oplus D$ that is effectively $A \oplus B \oplus C \oplus D$, which is a member of the `PermittedExecutionPath` set $A_{cfa-Set}$.

```

1 byte A(byte inputValue){
2   for(byte i=0; i<inputValue; i++){
3     if(i % 2 == 0){C(inputValue);
4     }else{B(inputValue);}}
5   return D(inputValue);}

```

Listing 1.4. Handling loop statements in the Permitted Execution Path Analysis.

Bytecode Integrity The property file associated with an application stores the hash values of individual methods. When the runtime environment fetches an application,

the RSM will measure the integrity value of individual methods of the application and compare them with the hash values in the property file. Therefore, any method that is loaded to the heap goes through the integrity validation. This validation protects against the fault attacks on an application stored while it is stored on a non-volatile memory.

The hashes of the individual methods (code and constant local-data variables) along with the integrity values (hash values) generated on the global persistent data can create a whole application integrity matrix. During the execution of an application, when it jumps from one method to another, it can be assured that the execution path is going to an (potentially) trusted method or the integrity of the called method will be verified. The RSM also tracks the integrity of the global variables and update the hash values if any authorised changes are being performed by the application.

4 Analysis of the Runtime Protection Mechanism

In this section, we evaluate the suitability of countermeasures against the attacks discussed in section 2.2 under the adversary’s capability detailed in section 3.1. Furthermore, we provide the latency and incurred overhead analysis for both serial and parallel RSMs.

4.1 Security Analysis

In this section, we discuss how the proposed counter-measures protect against the combined attacks under the attacker’s capability detailed in section 3.1.

Operand Stack Integrity We proposed a more refined approach to Barbu et al. [36] and removed the need to perform integrity measurement of the entire operand stack on each validation. In addition, we made the validation process continuous thus checking the integrity of the operand stack on each pop and push operation. If a malicious user changes values on the operand stack, the RSM can not only detect the modification but can also provide error correction service by providing the correct value that was stored on the operand stand. Furthermore, by using a random number, our proposal makes it difficult for an adversary to know the values stored on the integrity stack, even if he has the knowledge of all values on the operand stack.

PEP Analysis The PEP analysis performed by the RSM during the execution of an application effectively prevents execution path attacks. If an attacker has the capability of multiple fault injections simultaneously, (which is beyond the capability of our attacker as stated in section 3.1) then he can in theory affect the RSM execution. Nevertheless, even with simultaneous injections the attacker may be able to skip a node in the execution tree but the RSM calculation on the subsequent nodes will reveal an illegal path of execution. Therefore, even in the parallel injection model the RSM will detect the erroneous execution path, unless the attacker will constantly keep on introducing injections for the whole execution of an application.

Bytecode Integrity Our countermeasure prevents an adversary from changing an application while it is stored on a non-volatile memory (capability four of an adversary discussed in section 3.1). To avoid such modifications, the RSM generates a hash of individual methods that are requested by the JCVM. If the hash matches the value stored (`MethodIntegrity` in listing 1.1) in the respective property file, the JCVM will proceed with the execution of the method; otherwise, the RSM will signal the termination of the application (and possibly mark it malicious and up for deletion). Furthermore, this protection mechanism can also safeguard the dynamic loading of applications/classes/routines as part of the web server or other applications, which are stored on off-card storage.

4.2 Evaluation Context

For evaluating the proposed counter-measures, we have selected four sample applications. Two of the applications selected are part of the Java Card development kit distribution: `Wallet` and `Java Purse`. The other two applications are offline attestation algorithm [52] and `STCPSP` protocol [53].

4.3 Latency Analysis

As discussed before, latency is the number of instructions executed after an adversary mounts an attack and the system becomes aware of it. Therefore, in this section we analyse the latency of the proposed counter-measures under the concepts of serial and parallel RSMs that are listed in table 1 and discussed subsequently.

Table 1. Latency measurement of individual countermeasure

Counter-measures	Serial RSM	Parallel RSM
Operand Stack Integrity	$0 + i$	$3 + i$
Permitted Execution Path Analysis	0	$3(C_n)$
Bytecode Integrity	0	0

In case of the operand stack integrity, the serial RSM finds the occurrence of an error (e.g. fault injection) with latency “ $0+i$ ”, where ‘ i ’ is the number of instructions executed before the manipulated value reaches the top of the operand stack. Similarly, the latency value in case of the operand stack integrity for the parallel RSM is “ $3+i$ ”, where ‘3’ is the number of instructions required to perform a comparison on pop operation (`On_Stack_Pop(poppedValue)` in listing 1.2). The latency value of the parallel RSM is higher than the serial. This has to do with the fact that while parallel RSM is applying the security checks the JCVM does not need to stop the execution of subsequent instructions.

Regarding the PEP analysis, the serial RSM has a latency of zero where the parallel RSM has latency value of “ $3(C_n)$ ”, where the value C_n represents the number of legal jumps in the respective `PermittedExecutionPath` set. To explain this further, consider the example in listing 1.3. The `PermittedExecutionPath` of method B has four possible values (`Bcfa-Set` in section 3.7). Thereby, the latency value for a jump in the method B in the worse case is “ $3(4) = 12$ ”. The value ‘3’ represents the number of instructions required to execute individual comparison.

A notable point to mention here is that all latency measurements listed in table 2 are based on the worst-case conditions. Furthermore, it is apparent that it might be difficult to implement a complete parallel RSM. To explain our point, consider two consecutive jump instructions in which the parallel RSM has to perform PEP analysis. In such a situation, there might be a possibility that while the RSM is still evaluating the first jump, the JCVM might initiate the second jump instruction. Therefore, this might create a deadlock between the JCVM and parallel RSM — we consider that either JCVM should wait for the RSM to complete the verification, or for the sake of performance the RSM might skip certain verifications. We opt for the parallel RSM that will switch to the serial RSM mode, restricting the JCVM to proceed with next instruction until the RSM can apply the security checks. This situation will be further explained during the discussion on the performance measurements in the next section.

4.4 Incurred Overhead Analysis

To evaluate the performance impact of the proposed counter-measures we developed an abstract virtual machine that takes the bytecode of each Java Card applet and then computes the computational overhead for individual countermeasure. When a Java application is compiled the Java compiler (`javac`) produces a class file as discussed in section 2.1. The class file is Java bytecode representation, and we utilise the `javap` tool that comes with Java Development Kit (JDK) as it produces the bytecode representation of a class file in human-readable mnemonics as represented in the JVM specification [51]. The abstract virtual machine takes the mnemonic bytecode representation and searches for push, pop, and jump (e.g. method invokes) opcodes. Subsequently, we calculated the number of extra instructions required to be executed in order to implement the counter-measures discussed in previous sections.

Table 2. Performance measurement (percentage increase in computational cost)

Applications	Serial RSM	Parallel RSM
Wallet	+29%	+22%
Java Purse	+30%	+26%
Offline Attestation [52]	+27%	+23%
STCP _{SP} [53]	+39%	+33%

To compute the incurred overhead as a number of instructions that are going to be executed by an application that our countermeasures verify/validate. After this measurement, we have associated costs based on additional instructions executed for each JCVM instruction and calculated as an (approximate) increase in the percentage of computational overhead and listed in table 2. Furthermore, the computational cost of generating a hash is dependent on individual hardware configuration. The same is also true for the execution of each of the instructions. This is the reason why we opt for the evaluation based on the number of increased instructions rather than the performance as it provides us a hardware agnostic cost.

For each application, the counter-measures have different computational overhead values because they depend upon how many times certain instructions that invoke the

counter-measures are executed. Therefore, the computational overhead measurements in table 2 can only give us a measure of how the performance is affected in individual cases - without generalising for other applications.

4.5 Comparative Analysis

In this section, we will compare our proposed framework with the existing state-of-the-art discussed in section 2.2 in the context of security requirements for a JCRE (listed in section 3.2).

Table 3. Comparison between Proposed and Existing Proposals

Requirement	Barbu et al. [7]	Sere et al. [39]	Dubreuil et al. [54]	Lancia [55]	RSM
Customisable	Limited	Yes	No	No	Yes
Developer Independent	Yes	No	Yes	Yes	Yes
Code Integrity	No	Yes	No	Yes	Yes
Stack Integrity	Yes	No	Yes	No	Yes
Execution Flow Evaluation	No	No	No	No	Yes

Table 3 illustrates that our proposed framework in comparison to other proposals, in the context of our attacker capabilities and security requirements, performs comparatively better. One thing to note is that other proposals compared in table 3 do not provide any matrix for performance degradation due to their countermeasures — similar to one presented in sections 4.3 and 4.4

5 Conclusion

In this paper we discussed the smart card runtime environment by taking the Java Card as a running example. The JCRE was described with its different data structures that it uses during the execution of an application. Subsequently, we discussed various attacks that target the smart card runtime environment and most of these attacks are based on perturbation of the values stored by the runtime environment. These perturbations are called fault injection, which was defined and mapped to an adversary’s capability in this paper. Based on these recent attacks on the smart card runtime environment, we proposed an architecture that includes the provision of a RSM. We also proposed three counter-measures and provided the computational cost imposed by them. The overall protection framework is then compared with the existing frameworks and we showed that our proposal provides comparatively a better protection. No doubt, counter-measures that do not change the core architecture the Java virtual machine, will almost always incur extra computational cost. Therefore, we concluded in this paper that a better way forward would be to change the architecture of the Java virtual machine. Finally, in the context of this paper we showed that current architecture can be hardened at the reasonable cost of a computational penalty.

References

1. R. Anderson and M. Kuhn, “Low Cost Attacks on Tamper Resistant Devices,” in *Security Protocols*, B. Christianson, B. Crispo, M. Lomas, and M. Roe, Eds. Springer, 1998, vol. 1361, pp. 125–136.
2. P. C. Kocher, J. Jaffe, and B. Jun, “Differential Power Analysis,” in *CRYPTO ’99: Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*. London, UK: Springer-Verlag, 1999, pp. 388–397.
3. D. Sauveron, “Multiapplication Smart Card: Towards an Open Smart Card?” *Inf. Secur. Tech. Rep.*, vol. 14, no. 2, pp. 70–78, 2009.
4. R. N. Akram and K. Markantonakis, “Smart Cards: State-of-the-Art to Future Directions, Invited Paper,” in *IEEE International Symposium on Signal Processing and Information Technology (ISSPIT 2013)*, C. Douligeris and D. Serpanos, Eds. Athens, Greece: IEEE CS, December 2013.
5. K. Markantonakis, K. Mayes, D. Sauveron, and I. Askoxylakis, “Overview of security threats for smart cards in the public transport industry,” in *Proceedings of the 2008 IEEE International Conference on e-Business Engineering*. IEEE CS, 2008.
6. E. Vétillard and A. Ferrari, “Combined Attacks and Countermeasures,” in *Smart Card Research and Advanced Application, 9th IFIP International Conference, CARDIS 2010*, ser. LNCS, D. Gollmann, J.-L. Lanet, and J. Iguchi-Cartigny, Eds., 2010.
7. G. Barbu, H. Thiebeauld, and V. Guerin, “Attacks on Java Card 3.0 Combining Fault and Logical Attacks,” in *Smart Card Research and Advanced Application, 9th IFIP WG 8.8/11.2 International Conference, CARDIS 2010*, ser. LNCS, D. Gollmann, J.-L. Lanet, and J. Iguchi-Cartigny, Eds., vol. 6035/2010, 2010, pp. 148–163.
8. S. Chaumette and D. Sauveron, “An Efficient and Simple Way to Test the Security of Java Cards,” in *Security in Information Systems*, E. Fernández-Medina, J. C. H. Castro, and L. J. G. Castro, Eds. INSTICC Press, 2005, pp. 331–341.
9. *Java Card Platform Specification*, Oracle Std. v3.0.1, May 2009.
10. *Java Card Platform Specification*, Sun Microsystem Inc Std. v2.2.2, March 2006.
11. G. Barthe, G. Dufay, L. Jakubiec, B. P. Serpette, and S. a. M. d. Sousa, “A Formal Executable Semantics of the JavaCard Platform,” in *Proceedings of the 10th European Symposium on Programming Languages and Systems*, ser. ESOP ’01. London, UK: Springer-Verlag, 2001, pp. 302–319.
12. P. H. Hartel and L. Moreau, “Formalizing the safety of Java, the Java virtual machine, and Java card,” *ACM Comput. Surv.*, vol. 33, no. 4, pp. 517–558, 2001.
13. G. Barthe, G. Dufay, L. Jakubiec, and a. M. d. Sousa, Sim “A Formal Correspondence between Offensive and Defensive JavaCard Virtual Machines,” in *VMCAI ’02: Revised Papers from the Third International Workshop on Verification, Model Checking, and Abstract Interpretation*. London, UK: Springer-Verlag, 2002, pp. 32–45.
14. G. Barthe and S. Stratulat, “Validation of the JavaCard Platform with Implicit Induction Techniques,” in *RTA*, 2003, pp. 337–351.
15. M. Éluard, T. P. Jensen, and E. Denney, “An Operational Semantics of the Java Card Firewall,” in *Proceedings of the International Conference on Research in Smart Cards: Smart Card Programming and Security*, ser. E-SMART ’01. London, UK, UK: Springer-Verlag, 2001, pp. 95–110.
16. M. Éluard and T. Jensen, “Secure Object Flow Analysis for Java Card,” in *CARDIS’02: Proceedings of the 5th conference on Smart Card Research and Advanced Application Conference*. Berkeley, CA, USA: USENIX Association, 2002, pp. 11–11.
17. J. L. Lanet and A. Requet, “Formal Proof of Smart Card Applets Correctness,” in *Proceedings of the The International Conference on Smart Card Research and Applications*. London, UK: Springer-Verlag, 2000, pp. 85–97.

18. H. Meijer and E. Poll, "Towards a Full Formal Specification of the JavaCard API," in *Smart Card Programming and Security*, ser. LNCS, I. Attali and T. Jensen, Eds. Springer Berlin / Heidelberg, 2001, vol. 2140, pp. 165–178, 10.1007/3-540-45418-7_14.
19. V. Almaliotis, A. Loizidis, P. Katsaros, P. Louridas, and D. Spinellis, "Static Program Analysis for Java Card Applets," in *CARDIS '08: Proceedings of the 8th IFIP WG 8.8/11.2 international conference on Smart Card Research and Advanced Applications*. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 17–31.
20. D. A. Basin, S. Friedrich, J. Posegga, and H. Vogt, "Java Bytecode Verification by Model Checking," in *CAV '99: Proceedings of the 11th International Conference on Computer Aided Verification*. London, UK: Springer-Verlag, 1999, pp. 491–494.
21. X. Leroy, "On-Card Bytecode Verification for Java Card," in *E-SMART '01: Proceedings of the International Conference on Research in Smart Cards*. UK, Springer, 2001.
22. D. A. Basin, S. Friedrich, and M. Gawkowski, "Verified Bytecode Model Checkers," in *TPHOLs '02: Proceedings of the 15th International Conference on Theorem Proving in Higher Order Logics*. London, UK: Springer-Verlag, 2002, pp. 47–66.
23. X. Leroy, "Bytecode verification on Java smart cards," *Softw. Pract. Exper.*, vol. 32, no. 4, pp. 319–340, 2002.
24. E. Biham and A. Shamir, "Differential Fault Analysis of Secret Key Cryptosystems," in *Proceedings of the 17th Annual International Cryptology Conference on Advances in Cryptology*. London, UK: Springer-Verlag, 1997, pp. 513–525.
25. T. S. Messerges, E. A. Dabbish, and R. H. Sloan, "Investigations of power analysis attacks on smartcards," in *Proceedings of the USENIX Workshop on Smartcard Technology on USENIX Workshop on Smartcard Technology*. Berkeley, CA, USA: USENIX Association, 1999, pp. 17–17.
26. S. P. Skorobogatov and R. J. Anderson, "Optical Fault Induction Attacks," in *Revised Papers from the 4th International Workshop on Cryptographic Hardware and Embedded Systems*, ser. CHES '02. London, UK, UK: Springer-Verlag, 2003, pp. 2–12.
27. J.-J. Quisquater and D. Samyde, *Eddy current for Magnetic Analysis with Active Sensor*. Springer, 2002.
28. C. Aumiller, P. Bier, W. Fischer, P. Hofreiter, and J.-P. Seifert, "Fault Attacks on RSA with CRT: Concrete Results and Practical Countermeasures," in *Cryptographic Hardware and Embedded Systems - CHES 2002*, Springer, 2003, vol. 2523, pp. 81–95.
29. "Joint Interpretation Library - Application of Attack Potential to Smartcards," Online, Tech. Rep., April 2006.
30. O. Vertanen, "Java Type Confusion and Fault Attacks," in *Fault Diagnosis and Tolerance in Cryptography*, ser. LNCS, L. Breveglieri, I. Koren, D. Naccache, and J.-P. Seifert, Eds. Springer Berlin / Heidelberg, 2006, vol. 4236, pp. 237–251, 10.1007/11889700_21.
31. A. Lemarechal, "Introduction to fault attacks on smartcard," in *On-Line Testing Symposium, 2005. IOLTS 2005. 11th IEEE International*, July 2005, p. 116.
32. W. Mostowski and E. Poll, "Malicious Code on Java Card Smartcards: Attacks and Countermeasures," in *CARDIS '08*. Berlin, Heidelberg: Springer-Verlag, 2008.
33. J. Hogenboom and W. Mostowski, "Full Memory Read Attack on a Java Card," in *4th Benelux Workshop on Information and System Security*, O. Pereira, J.-J. Quisquater, and F.-X. Standaert, Eds. Belgium: Springer, November 2009.
34. J.-L. Lanet and J. Iguchi-Cartigny, "Developing a Trojan applet in a Smart Card," *Journal in Computer Virology*, vol. 6, no. 1, 2009.
35. A. A. Sere, J. Iguchi-Cartigny, and J.-L. Lanet, "Automatic Detection of Fault Attack and Countermeasures," in *Proceedings of the 4th Workshop on Embedded Systems Security*, ser. WESS '09. New York, NY, USA: ACM, 2009, pp. 71–77.
36. G. Barbu, G. Duc, and P. Hoogvorst, "Java Card Operand Stack: Fault Attacks, Combined Attacks and Countermeasures," in *The tenth Smart Card Research and Advanced Application IFIP Conference (CARDIS2011)*, E. Prouff, Ed. Belgium, September 2011.

37. G. Barbu and H. Thiebauld, "Synchronized Attacks on Multithreaded Systems - Application to Java Card 3.0-," in *The tenth Smart Card Research and Advanced Application IFIP Conference (CARDIS2011)*, ser. LNCS, E. Prouff, Ed. Springer, September 2011.
38. G. Bouffard, J. Iguchi-Cartigny, and J.-L. Lanet, "Combined Software and Hardware Attacks on the Java Card Control Flow," in *The tenth Smart Card Research and Advanced Application IFIP Conference*, E. Prouff, Ed. Belgium: Springer, September 2011.
39. A. A. Sere, J. Iguchi-Cartigny, and J.-L. Lanet, "Evaluation of Countermeasures Against Fault Attacks on Smart Cards," in *International Journal of Security and its Applications*, vol. 5, no. 2, April 2011.
40. O. Derouet. (2007, September) Secure Smartcard Design Against Laser Fault. (Invited Speaker) In 4th Workshop on Fault Diagnosis and Tolerance in Cryptography (FDRC 2007). IEEE-CS. Vienna, Austria.
41. S.-K. Kim, T. H. Kim, D.-G. Han, and S. Hong, "An efficient CRT-RSA algorithm secure against power and fault attacks," *Journal of Systems and Software*, vol. 84, no. 10, pp. 1660–1669, 2011.
42. S. Liu, B. King, and W. Wang, "A CRT-RSA Algorithm Secure against Hardware Fault Attacks," *2006 2nd IEEE International Symposium on Dependable Autonomic and Secure Computing*, pp. 51–60, 2006.
43. E. Trichina and R. Korkikyan, "Multi Fault Laser Attacks on Protected CRT-RSA," *2010 Workshop on Fault Diagnosis and Tolerance in Cryptography*, pp. 75–86, 2010.
44. S. Zhou, B. R. Childers, and N. Kumar, "Profile Guided Management of Code Partitions for Embedded Systems," in *DATE '04: Proceedings of the conference on Design, automation and test in Europe*. Washington, DC, USA: IEEE CS, 2004, p. 21396.
45. T. Zhang, S. Pande, and A. Valverde, "Tamper-resistant Whole Program Partitioning," in *LCTES '03: Proceedings of the 2003 ACM SIGPLAN conference on Language, compiler, and tool for embedded systems*. New York, NY, USA: ACM, 2003, pp. 209–219.
46. X. Zhuang, T. Zhang, H.-H. S. Lee, and S. Pande, "Hardware Assisted Control Flow Obfuscation for Embedded Processors," in *CASES '04*, USA: ACM, 2004.
47. G. Bouffard, J.-L. Lanet, J.-B. Machemie, J.-Y. Poichotte, and J.-P. Wary, "Evaluation of the Ability to Transform SIM Application into Hostile Application," in *the Tenth Smart Card Research and Advanced Application Conference (CARDIS 2011)*, ser. LNCS, E. Prouff, Ed. Leuven, Belgium: Springer, September 2011.
48. J. Loinig, C. Steger, R. Weiss, and E. Haselsteiner, "Identification and Verification of Security Relevant Functions in Embedded Systems Based on Source Code Annotations and Assertions," in *Information Security Theory and Practices. Security and Privacy of Pervasive Systems and Smart Devices*, ser. LNCS, P. Samarati, M. Tunstall, J. Posegga, K. Markantonakis, and D. Sauveron, Eds. Springer, 2010, vol. 6033, pp. 316–323.
49. A.-A.-K. Séré, J. Iguchi-Cartigny, and J.-L. Lanet, "Checking the Paths to Identify Mutant Application on Embedded Systems," in *Future Generation Information Technology - Second International Conference (FGIT 2010)*, ser. LNCS, T.-H. Kim, Y.-H. Lee, B. H. Kang, and D. Slezak, Eds., Korea,: Springer, December 2010, pp. 459–468.
50. W. Rankl and W. Effing, *Smart Card Handbook*, 3rd ed. New York, NY, USA: John Wiley & Sons, Inc., 2003.
51. T. Lindholm and F. Yellin, *The Java Virtual Machine Specification*, 2nd ed. Addison-Wesley Longman, Amsterdam, April 1999.
52. , "Remote Attestation Mechanism for User Centric Smart Cards using Pseudorandom Number Generators," in *5th International Conference on Information and Communications Security*, S. Qing and J. Zhou, Eds. Beijing, China: Springer, November 2013.
53. —, "A Secure and Trusted Channel Protocol for the User Centric Smart Card Ownership Model," in *12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications*. Australia: IEEE CS, July 2013.

54. J. Dubreuil, G. Bouffard, J. Lanet, and J. Cartigny, "Type Classification against Fault Enabled Mutant in Java Based Smart Card," in *Availability, Reliability and Security (ARES), 2012 Seventh International Conference on*, Aug 2012, pp. 551–556.
55. J. Lancia, "Java Card Combined Attacks with Localization-Agnostic Fault Injection," in *Smart Card Research and Advanced Applications*, ser. LNCS, S. Mangard, Ed. Springer Berlin Heidelberg, 2013, vol. 7771, pp. 31–45.