

# Testing security policies: going beyond functional testing

Yves Le Traon, Tejedine Mouelhi,  
Nora Cuppens  
GET-ENST Bretagne  
2, rue de la Châtaigneraie  
CS 17607  
35576 Cesson Sévigné Cedex, France  
yves.letraon@enst-bretagne.fr

Benoit Baudry  
IRISA/INRIA 35042 Rennes, France  
bbaudry@irisa.fr

## Abstract.

*While important efforts are dedicated to system functional testing, very few works study how to test specifically security mechanisms, implementing a security policy. This paper introduces security policy testing as a specific target for testing. We propose two strategies for producing security policy test cases, depending if they are built in complement of existing functional test cases or independently from them. Indeed, any security policy is strongly connected to system functionality: testing functions includes exercising many security mechanisms. However, testing functionality does not intend at putting to the test security aspects. We thus propose test selection criteria to produce tests from a security policy. To quantify the effectiveness of a set of test cases to detect security policy flaws, we adapt mutation analysis and define security policy mutation operators. A library case study, a 3-tiers architecture, is used to obtain experimental trends. Results confirm that security must become a specific target of testing to reach a satisfying level of confidence in security mechanisms.*

## 1 Introduction

For an organization, the security policy addresses constraints on access to data or functions by external users (or programs) and by people belonging to the organization. Such constraints are expressed using one of several access control models (DAC[1], MAC[2, 3], RBAC[4-6], TBAC[7], OrBAC [8, 9]) For a system to be developed, all these models describe the permissions or prohibitions for people to any of the resources of the system (it may apply to configure a firewall as well as to define who can access a given service or data in a database). The most advanced models ([4, 8, 9]) express rules that specify permissions or prohibitions that apply only to specific circumstances, called contexts. For instance, in the health care domain, physicians have special permissions in specific contexts, such as the context of urgency. Also, some models provide means to specify

the different security policies applicable to the various parts of an organization (sub-organizations). At the end of this specification process, the security policy specifies what the permissions and prohibitions should be in the system, in function of contexts, roles and views.

The software development process then consists in building a system according both to the functional requirements and to the extra-functional ones, including the security policy. This process is a sequence of activities, from high-level analysis to design and implementation. The output of the development process is a software system that should meet the requirements, including security ones. In most cases, the deployment of a security policy is not automated and the correctness of its implementation has to be verified or tested. In that context, testing can be applied as a technique to get some evidence that the security policy implementation is correct with respect to the requirements (security policy).

In this paper, we introduce security policy testing as a necessary but hard step ([10]) that takes place at the end of the development process. We build test cases from the security policy specification and run those tests on the system to reveal security flaws or inconsistencies between the policy and the implementation. In practice, the difficulty is that any security policy is strongly connected to system functionality: testing functions includes exercising many security mechanisms. The issue is to determine:

- whether testing functions provide enough confidence in the security mechanisms,
- how to improve this confidence by selecting test cases specific to security.

We study how to select security policy test cases to reveal security flaws in the security policy implementation. We thus propose two strategies for producing security policy test cases, depending if they are built in complement of existing functional test cases or independently from them.

The first step when applying testing to a specific context, such as security testing, is to analyze which faults (in our case security flaws) can occur in this

specific context, and to define test criteria, to qualify a set of test cases.

In this paper, we propose and discuss what SP testing is and propose test criteria to generate test cases from an access control model. To obtain an experimental basis and compare the criteria and the strategies, a fault/flow model for security policies is defined. It leads to an adaptation of mutation analysis [11] and thus provides a security testing qualification method. This flow model is expressed in the form of mutation operators to be applied on the security policy mechanisms. Using mutation analysis on a case study (a library system implemented with a 3-tiers architecture), we compare and discuss the criteria and strategies and highlight the specific issues related to SP testing. The most important contribution of the paper consists of demonstrating that testing a security policy is a task distinct (but tightly related) from functional testing. The corollary contributions are a first definition of this testing task, the security flaws model, and first SP test criteria and strategies that still need to be improved, as shown by the case study results.

Section 2 introduces the context of the paper, the chosen security policy model called OrBac, and the method to obtain a consistent security policy on one hand and to qualify test cases for testing this security policy on the other hand. Section 3 goes into the details of the possible mutation operators. In section 4, an empirical study reveals the first results on the proposed criteria and strategies, and shows various issues which arise when testing security.

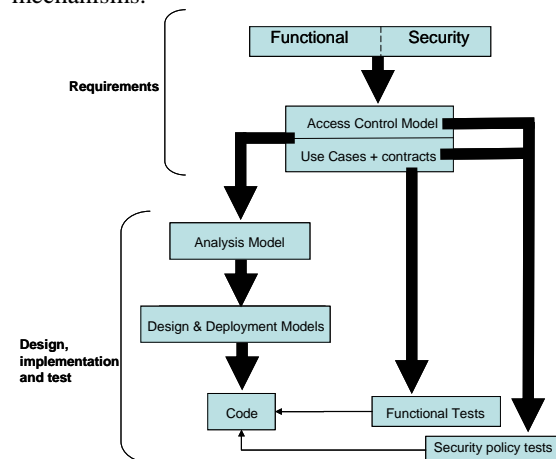
## 2 Process and definitions

In this section we present the overall process to derive test cases from requirements. In this process we identify the main artefacts that must be produced for testing and propose precise definitions for the notions used in this paper.

The software lifecycle may vary, but the input information is always a requirement document which includes the functional description of the system as well as many extra-functional concerns (performances, real-time, availability, technical and architectural constraints ...). Among these concerns, the security aspects are often mixed with the functional ones, as it will be illustrated with the running example throughout the paper. The requirement analysts have to extract these aspects and express, on one hand the uses cases and the business model and on the other hand explicit the security policy in the form of an access control model. It is composed of a set of security policy rules, each of them specifying rights and restrictions of actors on parts and resources of the system. The business

model (a class diagram + simple dynamic models) focuses on the concepts which are needed to derive functional aspects, in a nominal use of the system. The security policy may introduce specific concepts, and reuse most of the concepts and functions identified in the uses cases and business model. As a result, the security policy makes reference to the business and use case models, but includes new concepts which are taken into account in the refinement process, either during design (SecureUML [12]) or at deployment/coding steps. Section 3 will detail these differences.

The Figure 1 highlights the fact that both the code and the tests, which exercise this code, are produced from the requirements by independent ways. The important point is that the security policy test cases are obtained using the access control model, while the functional test cases are derived only from the uses cases (and business model). Security policy test cases are not only dependent on the security policy but also refer to the use cases and the business model. In this paper, the functional test cases (or system tests in the sense of Briand's work [12]) are generated using the approach presented in [13], based on the use cases improved with pre- and post-conditions, called contracts. The functional test cases cover all the code implementing the functions of the system. We will study how functional test cases can be reused for testing security mechanisms.



**Figure 1 – Process to generate security policy tests from an access control model**

The following definitions help clarifying the differences between functional and security policy testing.

*System/functional testing:* the activity which consists of generating and executing test cases

which are produced based on the uses cases and the business models (e.g. analysis class diagram and dynamic views) of the system [13, 14]. By opposition with security tests, we call these tests *functional*.

*Security policy testing (SP testing)*: it denotes the activity of generating and executing test cases which are derived specifically from a security policy. The objective of SP testing is to reveal as many security flaws as possible.

*Security flaw*: A security flaw is the equivalent of a fault for functional testing. It corresponds to an inconsistency between the security policy (the specification) and a security mechanism (the implementation) which is revealed at runtime.

*Test case*: In the paper, we define a test case as a triplet: *intent*, input test sequence, oracle function. The *intent* may be either to test functional or security policy aspects.

*Intent of a test case*: The intent of a test case is the reason why an input test sequence and an oracle function are associated to test a specific aspect of a system. It includes at least the following information: (functional, names of the tested functions) for functional test cases or (security policy, names of the tested security rules) for SP ones.

*SP oracle function*: The oracle function for a SP test case is a specific assertion which interrogates the security mechanism. There are two different oracle functions:

- For permission, the oracle function checks that the service is activated.
- For a prohibition, the oracle checks that the service is not activated.

In our case, services are methods in the business application, thus the oracle checks the absence or presence of a method call. It has to be noted that some data are constrained by access control rules, but the oracle function does not check any data values (in the database) because they are always accessed through a service.

The intent of the *functional* tests is not to observe that a security mechanism is executed correctly. For instance, for an actor of the system who is allowed to access a given service, the functional test intent consists of making this actor execute this service. Indirectly, the permission check mechanism has been executed, but a specific oracle function must be added to transform this functional test into a security policy test.

A security flaw may occur:

- when an actor (person or program) of the system has (has not) access to a resource while the security policy stipulates it should not (should).
- when an actor (person or program) of the system has (not) access to a resource while no security policy exists for this particular actor and when this resource is the object of other security rules.

The second point corresponds to security test cases highlighting the incompleteness in the security policy which leads to incorrect default security mechanisms. We call such test cases advanced security test cases in the following.

Examples:

**Functional** – a functional test case will make a borrower borrow and return a book.

*Intent*: functional, test borrow and return for a unique borrower

*Oracle*: check that the book is available after it has been returned

**Security Policy** – a SP test case will check that a borrower can borrow a book to the library the working days.

*Intent*: security, permission for a borrower to borrow a book the working days.

*Test sequence*: create the context of a working day, make a borrower borrow a book.

*Oracle*: interrogate the security mechanism and check that the permission has been computed and given to the borrower.

### 3 Running the process on an example

This section introduces an example based on a library management system. In this work we use OrBAC [8, 9] as a specification language to define the access control rules (a set of rules specifies a security policy). Based on the simplified requirements of this system, it is possible to derive a set of access control rules using the OrBAC model. We use these rules to illustrate the main features of the OrBAC language.

#### 3.1 Library management system

The purpose of the library management system (LMS) is to offer services to manage books in a public library. The books can be borrowed and returned by the users of the library on working days. When the library is closed, users can not borrow books. When a book is already borrowed, a user can make a reservation for this book. When the book is available, the user can borrow it. The LMS distinguishes three types of users:

public users who can borrow 5 books for 3 weeks, students who can borrow 10 books for 3 weeks and teachers who can borrow 10 books for 2 months. The library management system is managed by an administrator who can create, modify and remove accounts for new users. Books in the library are managed by a secretary who can order books, add them in the LMS when they are delivered. The secretary can also fix the damaged books n certain days dedicated to maintenance. When a book is damaged, it must be fixed. While it is not fixed, this book can not be borrowed but it can be reserved by a user. The director of the library has the same accesses than the secretary and he can also consult the accounts of the employees. The administrator and the secretary can consult all accounts of users. All users can consult the list of books in the library.

### 3.2 Business functional model

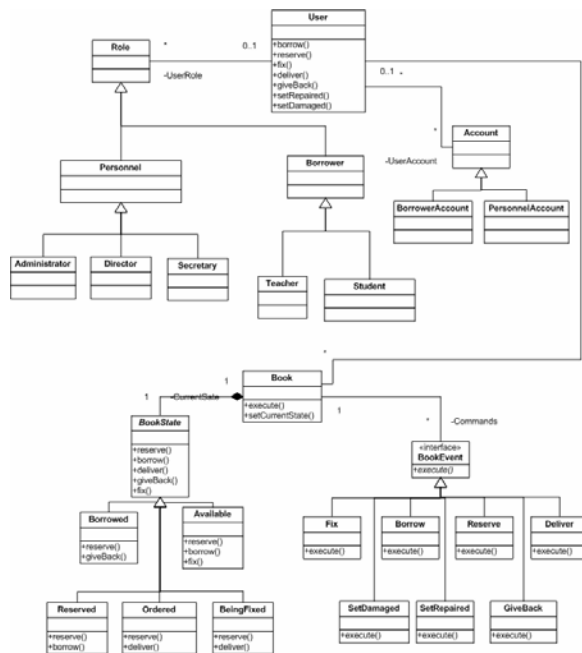


Figure 2 - Business model for the LMS

From the requirements of the library management system, we identify the business concepts that form the design model displayed in Figure 2. The main concepts to describe the functionalities are the user of the system that can have an account, the book and the roles. We have used a Command design pattern [15] to reify the different actions that can be performed on a book (borrow, return, etc.). We also use a State design for the book.

We have also introduced the notion of role as a business concept. It is interesting to notice that during

the design of a secured application, it is necessary to identify in the requirements the concepts that are relevant for the business application and the ones that are related mostly to security. As mentioned in previous section, the functional/business and security concerns are mixed and tightly coupled in the requirements. The way to separate these elements is very much linked to a design process and to design decisions. The notion of role could be present only in the security model. However, it has been introduced it in the business design because we it is an important functional concept that is necessary to execute the basic services of the system. We can notice that the concepts of working day and holiday are not present in this model. These are only related to security issues: they are necessary only to define access rules but are not required to execute the service.

### 3.3 Modeling a security policy with OrBAC

In parallel of the design model, it is possible to model the security policy from the requirements. We distinguish 5 different roles: public users, students, teachers, administrator and secretary. It is also possible to identify several rules that authorize or forbid access to data or services of the LMS.

OrBAC allows defining a set of security rules. A rule can be a permission, prohibition or obligation. Each rule has 5 parameters (called entities): an organization, a role, an activity, a view and a context. To increase modularity for the definition of security rules, OrBAC enables the definition of hierarchies of entities. In that case, rules defined on high level entities are inherited by the sub-entities. From the primary rules, secondary rules are derived, as illustrated in the following.

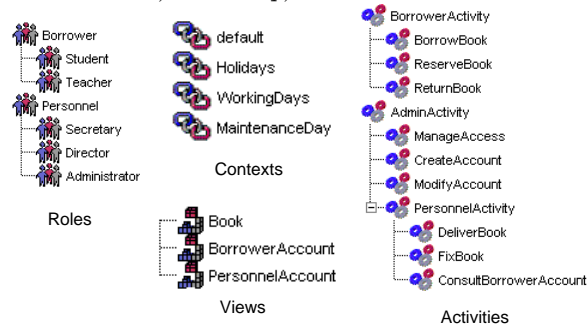
From the LMS requirements, we identify the entities displayed in Figure 3. The graphical representation is the one from MotOrBAC, the tool implementing the OrBAC model. First, we identify the services which are constrained by security rules. They are called activities in OrBAC. All users can perform three activities: borrow, reserve, return a book. Since all these activities are associated to the same rules, we define a high-level activity called `BorrowerActivity`. This means that all rules defined on the super activity will apply on sub activities. There are also a number of administrative tasks that can be executed. Since the administrator is allowed to execute all these tasks we define a super activity `AdminActivity`. The activities that inherit from `PersonnelActivity` are the activities that are permitted for the director and the secretary. It is interesting to notice that the service for consulting the list of books is not modeled as an activity in the access control model. Since everyone can access this

service, there is no restriction for this access, thus there is no rule related to this service.

Concerning the data (called views in OrBAC) that are restricted with access control, we identify the notion of book and of account. There are two types of accounts: borrower and personnel accounts.

The notion of context corresponds to a temporal (or spatial) dimension that appears in access control rules. In the requirements, we clearly identify `Holidays`, `WorkingDays`, `MaintenanceDay` and a default context which is used to define rules that are related to no specific time.

The last type of entities that needs to be defined concerns the roles. In the requirements we distinguish two main categories of roles: the users and the administrative personnel. We define the `Borrower` role which captures the role of user. Since `Teacher` and `Student` are two specific types of users, we model them as two sub-roles for `Borrower`. Concerning the personnel, we distinguish three categories: `Administrator`, `Secretary`, and `Director`.



**Figure 3 - OrBAC entities for the LMS**

Once all the entities are defined, they can be used to specify access control rules. Again, these rules are derived from the requirements. For example, requirements specify that users are allowed to borrow books only when the library is opened. This rule is defined as follows:

```
Permission(Library, Borrower,
BorrowerActivity, Book, WorkingDays)
```

The first parameter for this permission rule is the organization. Since there is only one organization in our example we give a very generic name `Library`. The other parameters are entities that have been defined previously.

Other example of rules include the prohibition to borrow a book during holidays, the permission for an administrator to manage the accounts for the personnel and the borrowers or the permission for the secretary to consult borrower accounts. These examples can be expressed as follows:

```
Prohibition(Library, Borrower,
BorrowerActivity, Book, Holidays)
```

```
Permission(Library, Administrator,
ManageAccess, PersonnelAccount, default)
```

```
Permission(Library, Administrator,
CreateAccount, BorrowerAccount, default)
```

```
Permission(Library, Secretary,
ConsultBorrowerAccount, BorrowerAccount,
default)
```

From the *explicit* (or *primary rules*), *secondary rules* are derived based on the parameters hierarchy. For example the rule:

```
Permission(rennesLibraries,Borrower,BorrowerActi
vity,Book,WorkingDays)
```

Based on hierarchies shown in Figure 3, the following rules are automatically derived such as:

```
Permission(Library,Student,BorrowBook,Book,
WorkingDays)
```

```
Permission(Library,Student,GiveBackBook,Book,
WorkingDays)
```

```
Permission(Library,Student,ReserveBook,Book,
WorkingDays)
```

```
Permission(Library,Teacher,BorrowBook,Book,
WorkingDays)
```

```
Permission(Library,Teacher,GiveBackBook,Book,
WorkingDays)
```

```
Permission(Library,Teacher,ReserveBook,Book,
WorkingDays)
```

The primary rule which is derived into these six secondary rules becomes useless from a testing point of view, if all the secondary rules are tested. We call such rules *generic rules* since they have no related security mechanism in the system under test.

One issue when specifying control access rules is that *conflicting rules* may appear. These conflicts can occur when 2 opposite rules have exactly the same parameters.

For example the following two rules are conflicting:

```
Permission(Library, Borrower,
BorrowerActivity, Book, WorkingDays)
```

```
Prohibition(Library, Borrower,
BorrowerActivity, Book, WorkingDays)
```

It is important to point out that conflicts may also occur when the parameters of the rule are not exactly the same. In the case a parameter in one rule inherits from a parameter in an opposite rule, the two rules are conflicting.

For example, in the following two rules, `Teacher` inherits from `Borrower`, thus the rules are conflicting:

```
Permission(Library, Borrower,
BorrowerActivity, Book, WorkingDays)
```

```
Prohibition(Library, Teacher,
BorrowerActivity, Book, WorkingDays)
```

An important benefit of using OrBAC is that the rules can be processed by a tool called `MotOrBAC` that

automatically detects the conflicting rules in the definition of a security policy, using an underlying Prolog motor. One way to solve the conflicts consists in assigning priorities to rules. The rule with the highest priority is executed.

For the LMS, 20 access control primary rules have been expressed. The total number of secondary rules is 22 and 7 rules are generic. So, the total number of rules is 42, 35 corresponding to a security mechanism.

## 4 SP test cases selection

For SP testing, the question is to ensure that security mechanisms are covered not only by input test sequences but are also exercised in every way that may lead to a failure of the policy.

In this section, we propose several test criteria to select test cases from an OrBaC specification. We also consider two strategies to produce efficient SP test cases w.r.t. the criteria.

### 4.1 Test criteria

Two test criteria are studied in this paper to select SP test cases from an OrBAC security policy model.

*CR1* - The criterion 1 (CR1) is satisfied iff a test case is generated for each primary access control rule of the security model.

In the case of the LMS, we have 20 such primary rules. In the case of a generic rule, a test case testing one instance of this rule is considered as sufficient.

*CR2* - The criterion 2 (CR2) is satisfied iff a test case is generated for each primary and secondary rule of the security model, except the generic rules.

In that case, 35 test cases are generated corresponding to the 42 total access control rules minus the 7 generic ones. The CR2 criterion is stronger than CR1, since it forces the coverage of all secondary rules.

*Advanced SP test cases* - *Advanced SP test cases* that exercise the default/non specified aspects of the security policy.

These test cases are selected to kill mutants generated with a specific mutation operator (ANR) presented in the next section.

*Functional test cases*: It corresponds to system tests, in the sense they are generated based on the use cases of the system under test.

In the case study, we used the approach based on use cases + contracts (pre- and post-conditions) to generate the functional test cases. The automated generation is obtained using the UCTS (Use Case Transition System) presented in [14]. The generated test cases cover the nominal code (code implementing the specified use of use cases) and a part of the robustness code of the system under test (unexpected use of a use case and specified situations when the use case execution fails).

Security test cases obtained with CR1 or CR2 should test aspects which are not the explicit objective of functional tests, e.g. that all prohibition rules that are not tested by functional tests.

### 4.2 Test strategies

In this paper, we study whether the functional test cases can be used for SP testing. Reusing functional test cases implies adapting them for explicitly testing the security policy. The intent of the functional test becomes *security* and details the SP rules which are tested by the input test sequence. The test oracle does not check the correctness of the service results, but interrogates the security mechanism and checks if the expected permission/prohibition is executed.

So, we consider two types of strategies depending whether we reuse the existing test cases or not.

*Incremental strategy*: It denotes the strategy for producing security test cases which reuse existing test cases.

An example of incremental test strategy consists of reusing functional test cases, then completing them with one of the CR1 or CR2 criteria, and finally completing the resulting test suite with advanced test cases.

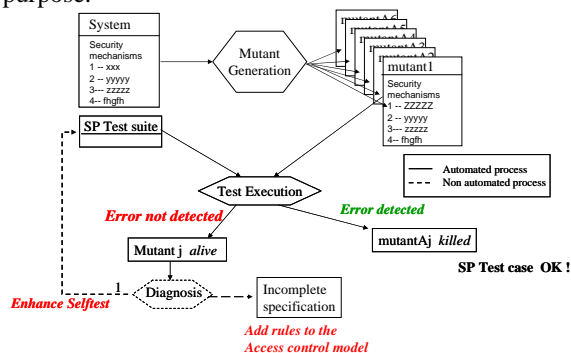
*Independent strategy*: This strategy consists of selecting functional, SP test cases and Advanced SP test cases independently.

We will compare and discuss in the case study several incremental strategies and the independent one. The goal is to highlight the many issues that arise when selecting test cases for the security policy aspect.

### 4.3 Test Qualification

Figure 2 presents the process for improving a SP test case suite so that it reaches a satisfying level, in terms of capacity in detecting security flaws. When flaws have been corrected by initial test cases, the test qualification loop consists of improving the test cases until all security flaws are detected. A security mutant thus differs from the initial code by the introduction of a single flaw in the security policy implementation.

The mutation score corresponds to the proportion of faulty versions of the system which are detected (or “killed”) by the test cases. A set of test cases is satisfying to test a security policy if it kills all the security mutants. The improvement process ends when the 100% mutation score is reached. The improved security policy test cases can then be reapplied on the non-mutated code to check whether actual flaws are detected. The mutation approach may lead to the improvement of the security policy, since to the advanced SP test cases. Exercising the system security policy robustness, these test cases allow revealing incompleteness and lacks in the access control rules. A specific mutation operator has been defined for this purpose.



**Figure 4. The qualification process for security policy tests**

We define mutation operators in a way which is independent from details specific to the implementation. Indeed, there are a lot of possible solutions to implement a security policy mechanism. Because they are expressed at high level, they have a clearer interpretation than if the flaws were injected “blindly” in the code. Injecting classical mutation faults (arithmetic, logical etc.) in security mechanisms may have been possible, but the effect of these faults would be unpredictable w.r.t. a given security policy (and the verdict difficult to define). Conversely to the fact security policy flaws are defined at high level, we assume that traceability matrices exist to locate the security mechanisms corresponding to the security policy specification. In practice, for the case study we present, we designed the software in such a way that these mechanisms were centralized, and inject flaws automatically. The mutants we propose are modifying the security policy in such a way that the mutants cannot be equivalent to the original system (a test sequence exists that produces a different result in the execution of the SP for the mutant in comparison with the original system).

## 4.4 Mutation operators

Mutation analysis is a technique for evaluating the quality of test cases. We propose to apply it in the context of security test cases generation. In order to apply mutation analysis to security tests we need to find suitable mutation operators. We express high level operators that will be applied to security policies and generate faulty versions of a security policy using OrBAC. A mutant is a copy of the original policy that contains one simple flaw.

### a Security mutant operators

We identified 8 mutant operators classified in 4 categories:

- Rule type changing operators
- Rule parameter changing operators
- Rule adding operators
- Hierarchy changing operators.

In this section, we present all of these operators. Then we present a study case in order to illustrate how they can be used and what kind of security flaws are simulated by the resulting policy.

#### Rule type changing operators

The following operators can be used in order to create mutants for security policies by changing predicate type:

Mutation Operator	Description
PRP	Prohibition rule replaced with a permission one.
PPR	Permission rule replaced with a prohibition one.

To illustrate rule type changing operators, we present 2 examples to show how they modify the security policy:

#### PPR:

*Rule used:* Permission (rennesLibraries, Administrator, ModifyAccount, BorrowerAccount, default). *Rule to use instead:* Prohibition (rennesLibraries, Administrator, ModifyAccount, BorrowerAccount, default).

#### PRP:

*Rule used:* Prohibition (rennesLibraries, Secretary, ModifyAccount, BorrowerAccount, default).

*Rule to use instead:* Permission (rennesLibraries, Secretary, ModifyAccount, BorrowerAccount, default).

### Rule parameter changing operators

The following operators replace parameters of predicates:

Mutation Operator	Description
CRD	Rule context is replaced with another context.
RRD	Rule role is replaced with another role

Some examples:

#### CRD:

Rule to use: *Permission(rennesLibraries, Secretary, FixBook, Book, MaintenanceDay)*.

Rule to use instead: *Permission(rennesLibraries, Secretary, FixBook, Book, HolidaysDays)*

#### RRD:

Rule to use: *Permission(rennesLibraries, Administrator, ManageAccess, PersonnelAccount, default)*.

Rule to use instead: *Permission(rennesLibraries, Teacher, ManageAccess, PersonnelAccount, default)*

### Rule adding or removing mutation operators

A mutant with an additional security rule can simulate a security hole. In fact, the security policy will be more permissive or more repudiating. This table presents this mutation operator.

Mutation Operator	Description
ANR	New rule added.

#### Some examples of ANR:

Rule added: *Permission(rennesLibraries, secretary, BorrowBook, Book, WorkingDay)*.

Rule added: *Permission(rennesLibraries, Secretary, DeliverBook, Book, Holidays)*.

### Hierarchy mutant operators

OR-Bac implements abstract entities inheritance. This means that we can specify an entity, e.g., "borrower" (here as a role) then its sub-entities, e.g., "student". Rules defined for the "borrower" will be automatically applied to "student". If we have this predicate:

*permission(Library, Borrower, borrow, book, WorkingDays)*

Then we have:

*permission(Library, Student, borrow, book, WorkingDays)*

We can create mutants by replacing a parent entity by one of its descendants or by replacing a child by one its ancestor.

Mutation Operator	Description
RPD	Rule role replaced with sub-role
APD	Rule activity replaced with sub-activity.

Some examples:

#### RPD :

Rule used: *Permission(rennesLibraries, Borrower, ReserveBook, Book, WorkingDays)*.

Rule to use instead: *Permission(rennesLibraries, Student, ReserveBook, Book, WorkingDays)*.

#### APD:

Rule used: *Permission(rennesLibraries, Student, BorrowerActivity, Book, WorkingDays)*.

Rule to use instead: *Permission(rennesLibraries, Student, ReserveBook, Book, WorkingDays)*.

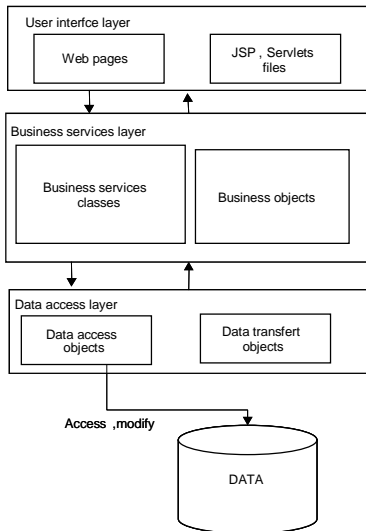
## 5 Case study and results

The case study application has a typical 3-tiers architecture widely used for web applications. Figure 5 presents the main characteristics of this architecture.

The security mechanisms implementing the Or-BAC based security policy are located in the business layer, in the service methods. Before performing the requested action, the method calls the security mechanism that checks if the security policy allows of forbids that action. When the action is prohibited a security policy violation exception is raised and thrown to the caller, otherwise (when the action is allowed) nothing is done and the execution of the action is pursued by the service method. There are two possible behaviours:

- The requested action is allowed: The security mechanism let the application continue normally.
- The requested action is prohibited: The security mechanism raises a security exception that interrupts the execution of the method. This exception is then handled by the caller.





**Figure 5 - The 3-tiers architecture for the LMS**

When we implemented the security policy, we maintained traceability links between the OrBAC rules and the corresponding blocks of code. This means that we are able to modify the security rules used by the application and introduce a mutant rule. The mutation analysis runs as follows:

- Tests are executed against the initial security policy.
- The oracle function associated to the test case checks that the service is activated (permission) or not (prohibition).

### 5.1 Generated mutants

Table 1 gives the number of generated mutants per operator. The ANR operator generates much more mutants since it adds a non specified security rule. This number may vary, depending on the completeness of access control rules. The fewer rules are specified, the more mutants this operator generates. The number of ANR mutants thus reflects the fact that all the possible cases have not been specified in the security policy. To our own experience, this is quite usual: the specification focuses on the most critical and important case, and often considers that default behaviour is acceptable. Testing these cases allow the default policy to be exercised and allow highlighting lack in the specification. On the other hand, they are few mutants generated from hierarchy changing operators because the specification doesn't introduce many hierarchical entities. The basic mutation operators (i.e. all operators except ANR) will generate more mutants when more rules are added. In a general case, there is thus a balance in the number of generated mutants between basic operators and ANR.

Operator category		Operator name	Number of mutants
Basic Mutation operators	Rule type changing operator	PPR	22
		PRP	19
	Rule parameter changing operator	RRD	60
		CRD	60
	Hierarchy changing operators	RPD	5
		APD	5
Rule adding operator		ANR	200
<b>Total</b>			<b>371</b>

**Table 1. Number of mutants per mutation type and operator**

### 5.2 Issue 1: Relationships/differences between functional tests and security ones

The first experiments we performed aim at studying whether the faults detected by functional test cases differ from (are included in or intersect with) the ones detected by test cases dedicated to the security policy. While the functional tests, which cover 100% of the code, necessarily execute security mechanisms, they should only focus on the success/failure of method/service sequence executions.

When reusing functional test cases with the objective of testing security, the intent changes and thus it is necessary to modify the associated test oracle. This task may be costly in the general case, depending on the difficulty to relate the functional test sequence to the security mechanisms it should exercise.

Table 3 summarizes the number of functional test cases. The table shows that some test cases do not trigger the security mechanisms. It corresponds to the execution of code which is not related to the security policy. It confirms the intuition that functional test objectives differ from the objective of explicitly testing security mechanisms.

Table 2 second column shows that the functional test cases, adapted to security, can kill most (78%), but not all, basic operators' mutants. Concerning, the ANR operator, the functional test cases are not efficient (11%): it is due to the fact that the ANR operator generates security flaws which are outside the scope of the specification. In conclusion, the overlap of functional aspect and security policy is high, but the functional test cases do not kill all security mutants. It appears as a meaningful task to generate test cases with the explicit objective of testing the security policy.

	<b>functional</b>		<b>CR1</b>		<b>CR2</b>	
<b>#Test Cases</b>	42		20		35	
Rule type changing operator	35/41	85%	38/41	92%	41/41	100%
Rule parameter changing operator	90/120	75%	101/120	84%	120/120	100%
Hierarchy changing operators	10/10	100%	10/10	100%	10/10	100%
<b>Overall score with basic security operators</b>	135/171	<b>78%</b>	149/171	<b>87%</b>	171/171	<b>100%</b>
Rule adding operator (ANR)	22/200	<b>11%</b>	28/200	<b>14%</b>	33/200	<b>17%</b>
<b>Overall score with all operators</b>	157/371	<b>42%</b>	177/371	<b>47%</b>	204/371	<b>55%</b>

**Table 2. Mutation analysis results by test cases category**

	<b># test cases</b>
Tests that do not that trigger sec. mechanisms	7
Tests that trigger security mechanisms	35
All tests	42

**Table 3. Functional tests analysis**

### 5.3 Issue 2: Comparing test criteria

Table 2 third and fourth columns present the mutation results for CR1 and CR2 test criteria. The 20 test cases selected with CR1 are more efficient than the functional test cases, since with fewer test cases the final mutation score is higher. However, this criterion is not efficient to reach a 100% mutation score on basic mutants (generated with all operators except ANR).

In conclusion, the CR2 criterion is necessary to provide a full coverage of basic mutants, and appears as good trade-off between functional and CR1, in terms of efficiency (mutation score) and cost (#test cases) for detecting security flaws. A corollary conclusion is that a large number of basic mutants are quite easy to kill. Since the hard-to-kill mutants are the most interesting ones (because they require the most efficient test cases), another study would consist of ranking the mutation operators so that only hard-to-kill mutants are generated. This study is beyond the objective of this paper but is necessary to offer a realistic mutation-based approach to test security policies. This paper focuses on the test selection problem for security policy and the question of sufficient mutant operator does not impact the conclusions we draw (in the worse case, we consider more mutants than necessary since some mutants are coupled).

We remark that the CR2 criterion allows covering up to 17% of the ANR mutants. While basic security mutants force the tests to cover the specified security rules, the ANR ones are useful to check the robustness of the system in case of default or underspecified policies. Combining both operators provides a good

criterion to guide the tester when generating the test cases.

The *advanced security test cases* are test cases which are explicitly generated in order to kill all the ANR mutants. In this approach, we are using mutant score as the test criterion. We do not use mutation for analysis purpose (to compare test criteria or testing technique) but as a test selection technique. In the case mutants generation cannot be fully automated, it makes this test selection technique impossible to apply. However, the study we propose provides interesting results, in terms of test selection effort and test quality.

### 5.4 Issue 3: Advanced vs. basic security tests

The issue here is to compare the test cases selected to cover the security rules with CR2 (and which kill all mutants except ANR) and the advanced security tests which are generated in order to kill all the ANR mutants. Table 4 presents the overlap between these two approaches. It is interesting to note that the advanced security test cases kill up to 60% of the basic security mutants. On the other hand, the test cases selected with CR2 only kill 17%. The effort to kill the ANR mutants is much more important (154) than for killing the basic security mutants (35).

	<b>#test cases</b>	<b>Basic security mutants</b>	<b>ANR</b>
<b>CR2</b>	35	100%	17%
<b>Advanced sec. tests</b>	154	59.3%	100%

**Table 4. Overlap of CR2 and advanced test cases**

In conclusion, the test selection based on the ANR mutants cannot replace the CR2 criterion. CR2 and advanced security test cases are not comparable, and are both recommended, the first to efficiently test the specification, the second to cover non-specified cases (robustness).

## 5.5 Issue 4: Incremental vs. independent test strategies

The issue now is to study whether we can leverage an incremental approach to save test generation effort. Table 5 recalls the number of test cases generated with the following strategies:

- the independent approach (we do not reuse functional test cases)
- reusing the functional test cases, completing them to reach the CR2 criterion and to kill all ANR mutants (*Incremental from functional strategy*).
- generating test cases to reach the CR2 criterion, completing them to kill all ANR mutants (*Incr. from CR2 strategy*).

The first incremental strategy seeks to take benefit from the existing functional test cases (which have to be adapted for security), the second one starts from the CR2 test criterion.

Even if quantitative results are displayed, the comparison is difficult because the effort to adapt functional test cases to security cannot be easily estimated in a general case. It depends on the system to be tested. It may be neglected: that's the case for our study since the security mechanisms are centralized and can be quite easily observed. In a general case, adapting these test cases may be as costly as generating security test cases. In Table 5 we put two values that correspond to the case when the cost of adaptation can be neglected in parenthesis. This issue is related to the problem of security mechanisms testability (controllability and observability).

#Test cases Strategy	Funct.	Basic Security CR2	Advanced Security	Total
Independent	-	35	154	189
Incremental from functional	(42)	21	133	196 (154)
Incr. from CR2	-	35	133	168

**Table 5. Independent vs. Incremental test generation strategies**

It appears that the independent generation of basic security and advanced test cases is the most costly. With any incremental strategy, we need to generate 133 test cases to kill all advanced security mutants (saving 21 test cases generation). The final ranking between the two incremental strategies depends on the adaptation cost of functional test cases and may vary from a system under test to another. Only the test experts may estimate this adaptation cost. If it can be neglected, reusing functional test cases and completing

them is the most interesting strategy. Another decision has to be taken which is to put an important effort for testing the security policy robustness (killing ANR mutants).

## 6 Related works

As it has been already mentioned, the most used models to express security at high level [1-9] help to specify the application access rules and then implement them into the code. Other models exist, which are not dedicated to access control, such as [16] which help expressing security requirements and model the potential attacks to be prevented in the design stage. In this paper, we only consider the access control models, but the security aspects to be tested are not restricted to them and many other security aspects need to be covered by systematic testing.

In our approach we use mutation analysis in order to improve the security tests. Security tests become stronger because they are capable of detecting security faults. To our knowledge, very few works proposed fault injection to security tests. In [17], Mathur et al. applied fault injection to the application environment. The application environment is perturbed by modifying environment variables, files or process used by the application under test. Then the application has to resist to this perturbation and must not have an insecure behaviour that may lead to security flaw.

In addition, fault injection was applied in another way. Adaptive vulnerability analysis [18] injects faults to the application data flow and internal variables. The objective is to identify parts of application's code that have insecure behaviours when the state of the application is perturbed. In [19], the goal is to simulate the consequences of an attack on the system by simulating the effects of a buffer overflow, thus allowing the comparison of the efficiency of security tools dedicated to these attacks (and impact analysis). Our approach differs from these approaches since we focus on testing the implementation of a security model (the access control model). Our intent is not to create attack scenarios but to check the correctness of the implementation w.r.t. a specification, as any classical testing technique.

## 7 Conclusion

This paper does not answer the many issues security testing tackles but allows underlining them and shows that testing a security policy is a specific task, which is complementary from the functional testing one. It proposes a first methodology to select test cases, with various test selection criteria based on the security

policy (specification) or on mutation (code based). Finally, it distinguishes test selection for testing the “nominal” security policy rules from the advanced test selection that aims at testing the robustness of the security policy. The first test cases can be derived from CR2 test criterion. In this paper, we use ANR mutants as the target and criterion for testing security policy robustness.

The case study highlights the many issues a test expert has to deal with when facing the objective of testing a security policy for a real system. In particular, two qualitative aspects arise that are the possibility, or not, to adapt functional test cases to test a security policy, and the interest of advanced security tests regarding the important additional effort it may require. More fundamentally, the aspect of security mechanism testability in relation to system architecture is critical. The way security mechanisms are spread over the system or centralized, the easiness or difficulty to relate a security rule to a piece of code are major issues to conduct the testing task.

## 8 References

1. Lampson, B. *Protection*. in 5th Princeton Symposium on Information Sciences and Systems., 1971.
2. D. E. Bell and L.J. LaPadula, *Secure computer systems: Unified exposition and multics interpretation*, in Tech. Rep. ESD-TR-73-306, The MITRE Corporation. 1976.
3. Biba, K.J., *Integrity consideration for secure computer systems*, in Tech. Rep. MTR-3153, The MITRE Corporation., 1975.
4. D. F. Ferraiolo, et al., *Proposed NIST standard for role-based access control*. ACM Transactions on Information and System Security, 2001. 4(3): p. 224–274.
5. S. I. Gavrilu and J.F. Barkley. *Formal Specification for Role Based Access Control User/Role and Role/Role Relationship Management*. in Third ACM Workshop on Role-Based Access Control. 1996.
6. R. Sandhu, E.J.C., H. L. Feinstein, and C.E. Youman, *Role-based access control models*. IEEE Computer, 1996. 29(2): p. 38-47.
7. Thomas, R.K. *TMAC: A primitive for Applying RBAC in collaborative environment*. in 2nd ACM, Workshop on RBAC. 1997.
8. A. Abou El Kalam, et al., *Organization Based Access Control*, in IEEE 4th International Workshop on Policies for Distributed Systems and Networks. 2003.
9. F. Cuppens and N. Cuppens-Boulahia and M. Ben Ghorbel, *High-level conflict management strategies in advanced access control models*. Workshop on Information and Computer Security (ICS'06), Timisoara, Roumania, september 2006.
10. Thompson, H.H., *Why security testing is hard*. IEEE Security & Privacy Magazine, 2003. 1(4): p. 83 - 86.
11. R. DeMillo, R. Lipton, and F. Sayward. *Hints on Test Data Selection : Help For The Practicing Programmer*. IEEE Computer, 1978. 11(4): 34 - 41.
12. D. Basin, J. Doser, T. Lodderstedt, *Model driven security: From UML models to access control infrastructures*, ACM Transactions on Software Engineering and Methodology (TOSEM), 15 (1):39-91 January 2006.
13. L. Briand and Y. Labiche, *A UML-based approach to System Testing*. Software and Systems Modeling, 2002. 1(1): p. 10 - 42.
14. C. Nebut, F. Fleurey, J-M Jézéquel, Y. Le Traon, *Automatic Test Generation: A Use Case Driven Approach*, IEEE Transactions on Software Engineering, 32(3):140-155, March 2006.
15. Gamma, E., et al., *Design Patterns: Elements of Reusable Object-Oriented Software*. Professional Computing. 1995: Addison-Wesley.
16. Van Lamsweerden, A. *Elaborating Security Requirements by Construction of Intentional Anti-Models*. in Proceedings of the 26th International Conference on Software Engineering. 2004.
17. Du, W. and A.P. Mathur. *Testing for Software Vulnerability Using Environment Perturbation*. in International Conference on Dependable Systems and Networks. 2000.
18. Ghosh, A., T. O'Connor, and G. McGraw, *An automated approach for identifying potential vulnerabilities in software*, in IEEE Symposium on Security and Privacy. 1998.
19. Ben Breech, Mike Tegtmeier, Lori Pollock, *An Attack Simulator for Systematically Testing Program-based Security Mechanisms*. 17th International Symposium on Software Reliability Engineering, Pages: 136 - 145, (2006).