Test Generation for Network Security Rules

Vianney Darmaillacq¹, Jean-Claude Fernandez², Roland Groz¹,
Laurent Mounier², and Jean-Luc Richier^{1,*}

¹ Laboratoire LSR-IMAG,
BP 72, 38402 St Martin d'Hères, France

² Laboratoire Vérimag,
Centre Equation - 2, avenue de Vignate, 38610 Gières, France
{Vianney.Darmaillacq, Jean-Claude.Fernandez, Roland.Groz,
Laurent.Mounier, Jean-Luc.Richier}@imag.fr

Abstract. Checking that a security policy has been correctly deployed over a network is a key issue for system administrators. Since policies are usually expressed by rules, we propose a method to derive tests from a set of rules with a single modality. For each element of our language and each type of rule, we propose a pattern of test, which we call a tile. We then combine those tiles into a test for the whole rule.

1 Introduction

Network and system administrators are in charge of implementing and controlling the security policies of their organisations. Enforcing a policy typically relies on the adequate configuration of Policy Enforcement Points (PEP) in dedicated equipment (such as firewalls, ciphering chips) or specific software (e.g. account managers, mail scanners). Since most networks and systems would undergo daily changes, maintaining the consistency of the rules actually implemented by the PEPs and their conformance to the specified security objectives and the rules expressed in the policy is not an easy task.

One way to guarantee a correct policy enforcement is to derive configurations from a description of the policy (top-down approach). In order to do that systematically, the description must be formal enough to provide unambiguous rules and translations of them into configurations of security devices. Since most policies consist of combination of rules with various scopes and potentially conflicting modalities (e.g. restricting access for generic subjects but authorising it for specific categories), the descriptions usually include constructs or semantic rules to solve conflicts between policy statements. Ponder [5] and OrBAC [1] are typical such description languages with associated methods to allow deployment on networks.

In this paper, we investigate another approach. We consider that testing a given network configuration for compliance with a stated policy is a kind of conformance testing. Therefore, we aim at deriving tests from a formal specification

* This work has been partly funded by the POTESTAT project of the national research programme on security (ACI Sécurité) and by the IMAG project Modeste.

of the security policy to check whether the implementation is correct. These tests could be used either after some initial deployment of a policy to check whether it can be and actually is well implemented, or typically on a more regular basis to see whether any update on the configurations might have breached security rules. Some sort of testing is actually performed by system administrators to check for known vulnerabilities: portscans and password crackers fall into that category. However, such tests are usually limited to just a single security mechanism. With existing techniques, it is difficult to address the issue of consistency of configurations on distributed devices. Although some work has been published on the analysis of the consistency of firewall rules (typically to minimise them or to detect conflicts), this is still limited to specific points of security policies. [12] is an application of the idea of generating conformance tests for a single firewall.

In our work, we address conformance with respect to a global specification of a security policy for a network of interconnected systems. Although at first view this might appear as just another instance of conformance testing of an implementation w.r.t. a given formal specification, there are a number of significant differences with the framework of such standards as IS9646 and its formal counterpart FMCT [8]. We investigated a few of these differences in [6]: typically, protocol conformance testing is done on a well defined protocol level, whereas security policies are often implemented with a mixture of mechanisms at various levels of communication and O/S interfaces. One major issue for test generation is that we cannot expect to start from an updated comprehensive model of the system: this is why we adopt a method based on (security) requirements.

In this paper, we focus on a method to derive tests from a policy expressed as a set of rules. More precisely, we consider an LTL-like specification language that can be used to formalise a reasonable subset of the rules usually found in network security policies. This restricted language allows us to design a “tile-based” generation method. For each element of our language and each type of rule, we propose a pattern of test, which we call a tile. The simple form of rules makes it possible to compose a global test by simply combining the tiles associated to the elements.

This paper is organised as follows. In section 2, we present our approach, in particular the types of rules that we cover and their relation to proposed formal methods for the description of security policies. Our notation for rules is presented in section 3. In section 4, we define the test generation tiles and the algorithm to derive tests. The whole approach is illustrated in section 5 on a typical example taken from an e-mail security architecture. Finally, in section 6 we present our perspectives for development from this basis.

2 Approach

Our aim is to automatically generate test cases from network policy rules. We intend to adapt the approach used in protocol conformance testing. We first started by looking for a formal description of security policies that would make it possible to generate tests. In order to identify typical requirements and the corresponding tests, we worked from a case study.

2.1 A Case Study

We carried out during the summer of 2004 a case study to identify the security policies used in the IMAG network, which connects our laboratories. This study included the analysis of documents provided to the administrators and the users, and interviews. We collected information at different levels of management, from the laboratory to the access supplier. The analysis resulted in a rather broad and detailed description of a typical network security policy in a university environment. In this paper, we shall present a small subset of rules, centred on electronic mail. This set was selected to be rich enough to show the majority of the concepts to be studied: several levels of policies and inter-connected organisations, variety of the services and of the access methods.

No	Requirement
1	Mail relays accepting messages from the exterior should be located at the entry of the network, in the DMZ if possible.
2	There should be no user account on hosts located in the DMZ.
3	Mailbox servers containing user accounts should be in the private zone. There can be as many of these servers as necessary.
4	Relays in the DMZ are the only machines allowed to communicate with the exterior world using the SMTP protocol. Relay of inbound mails (to mailboxes) and of outbound mails (to exterior) is done using these relays.
5	At the entry of the site, a filtering default policy is applied, which forbid all traffic not explicitly authorised.
6	It is forbidden to hosts of the network to relay mails from an external host to an external host.
7	All messages coming from the exterior are redirected to mail relays located at the entry of the site (using DNS MX records), probably in the DMZ.
8	It is forbidden to communicate with a host belonging to the blacklist updated daily provided by the MAPS (Mail Abuse Prevention System) partner.
9	Antivirus and spam filters shall be installed on hosts acting as relays or mailboxes.
10	A mail shall be checked by antivirus software before being opened.
11	All mails entering the network infected by a virus shall be disinfected.
12	All mail shall be modified if it contains a potentially dangerous attached file. The original file is kept somewhere in the network. Recipients are notified.
13	All mail analysed as a possible spam shall be flagged.

Fig. 1. Security rules for e-mail

In a network, security distinguishes inside and outside. The inside of the network is the set of machines under the responsibility of the organisation. The outside consists of uncontrolled machines considered dangerous for security purposes. This can be refined by considering other zones, differing by degrees of administration, reliability and trust. The sub-zones of the internal zone correspond in general to architectural criteria, for example separations between physical sub-networks, whereas the outside sub-zones correspond to different levels of

trust. One distinguishes moreover among the internal hosts those which provide services accessible from the outside. Due to their visibility, these hosts are more often subject to attacks. Therefore the standard approach is to define for these hosts a strongly controlled buffer zone, often called DMZ (demilitarized zone). Only the hosts in this zone can communicate with the outside world, and all the traffic between the DMZ and the rest of the internal network is controlled.

We give in figure 1 a sample set of rules for the electronic mail drawn from our case study. This sample illustrates: flows of information, separation in zones, possibility, obligation or prohibition of certain actions, at different levels of details. To simplify the problem, we suppose that certain global hypotheses are true and do not have to be clarified: correct routing, systems up to date w.r.t vulnerability patches ...

2.2 Description Techniques for Network Security

Most of the rules in the case study (which actually covers much more than e-mail) express some constraints about the possible behaviour of the system. More specifically, they are of the form “*Mod P*”, where *Mod* is a modality among obligation, permission or interdiction, and *P* is either a predicate on the system or a behaviour.

The deontic logic of von Wright [14] is a modal logic whose modal operator semantic is that of obligation. With only this operator authorisation of a formula is defined as the negation of the obligation of the negation of the formula, while interdiction is defined as the negation of the obligation of the formula. However deontic logic raises a number of paradoxes that have hindered its wider use [10]. Nevertheless modalities are a key issues in formal models of security policies, in particular in formal description techniques such as PDL, Ponder and Or-BAC, which have been proposed to address network security policies.

PDL is a language created to model network management policies, including security requirements. It is based on the idea that a policy is a specification of the behaviour the network should have according to what happened in the network [9]. A rule states that an action is triggered by an event, provided a condition holds. PDL was used to monitor switches in a network, to guarantee quality of service [13].

Ponder is a language created to specify network security policies [5] and proposes a full choice of different modalities: conditional obligation, authorisation, interdiction, delegation and refrain policies could be specified.

Or-BAC [1] is another access control model, loosely based on RBAC. While RBAC abstracts subjects into roles, Or-BAC abstracts moreover objects into views and actions into activities, and introduces the notion of organisation. Or-BAC has been used to model a network security policy, and to generate firewall rules from it [4, 2].

All these formalisms include structuring and typing constructs, resolution of conflict mechanisms and various aspects which are important. In this paper, we go for a simpler course, since we want to derive tests from the rules. These rules could be extracted from description in the above formalisms. All we need is a

description with just enough expressive power to represent the modalities used in PDL, Ponder or Or-BAC, at least as they can be tested through network events.

2.3 Approach for Test Generation

Compared to classical conformance testing for communication protocols, test generation in our case exhibits two major differences.

- The policy is not described by a comprehensive model such as a global LTS or state machine, but by a collection of rules. This is similar to deriving tests from requirements. Much work has been done in test generation from test purposes which are confronted to a formal model of the protocol. Here, we do not assume any formal model of the network, we derive tests from the rules. Our approach is rule based: it generates separate tests for each rule.
- The policy is described at a much higher level than the actual events that can be observed or controlled in the network; whereas the specification of a communication protocol would refer to PDUs or SDUs even though some of them might occur at non-observable interfaces. We need to establish a correspondence between the basic predicates appearing in a rule and sequences of events that can represent a test or instances of such predicates.

To each high-level predicate (such as *externRelay(h)* meaning that machine *h* can relay mails sent from outside the domain considered) we associate a test pattern which we call a *tile*. For instance, in the case of external relaying of mails, a tester would have to establish a SMTP connection to the machine from an external machine and try to send a mail. However, there are different types of predicates. Some may have to be tested dynamically through interaction with the system. Others might be checked without PDU exchanges, for instance if we have access to the configuration files of the system when the test is set up. In the case of *externRelay(h)*, this could be checked in the configuration of the mail system on *h*. The choice of one method or another may depend on accessibility to the system, but on trust as well: typically, information on configurations might not be reliable.

In [7], we investigated a refinement approach to derive control at PDU level from security rules. In this paper, we will consider that the tiles are provided. A policy would be described by combining elementary predicates which would be well-known elements for security policies, so that the corresponding tiles would be readily available.

In this paper, we concentrate on the combination of tiles, based on the structure of the formula in the rule. We address formulas of a restricted form with just one modality, as this corresponds to the usual style of security policy rules; Ponder, PDL and OrBAC also propose a single modality on each rule. We first propose a formal description of the rules in the next section, then we describe the combination of tiles in the following one.

3 Rules Formalisation

We first give the syntax and semantics of a rather general formalism and then we restrict it to a smaller subset used in this work to generate test cases.

Preliminaries. A *labelled transition system* (LTS, for short) is a quadruplet (Q, A, T, q^0) where Q is a set of states, A a set of labels, T the transition relation ($T \subseteq Q \times A \times Q$) and q^0 the initial state ($q^0 \in Q$). We will use the following definitions and notations: $(p, a, q) \in T$ is noted $p \xrightarrow{a}_T q$ (or simply $p \xrightarrow{a} q$). An *execution sequence* ρ is a composition of transitions: $q^0 \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \cdots \xrightarrow{a_n} q_n$. We denote by σ^ρ (resp. α^ρ) the sequence of states (resp. observables actions) associated with ρ . The sequence of actions α^ρ is called a *trace*. We note by Σ_S , the set of finite execution sequences starting from the initial state q^0 of S . For any sequence λ of length n , λ_i or $\lambda(i)$ denotes the i -th element and $\lambda_{[i \dots n]}$ denotes the suffix $\lambda_i \cdots \lambda_n$.

We also consider in the sequel Boolean Labelled Transition Systems in order to obtain a more compact representation of test cases. A Boolean Labelled Transition System (BLTS for short) is a tuple (X_b, Q, A, T, q^0) where X_b is a set of Boolean variables, Q is a set of states, A is a set of actions, q^0 is the initial state and $T \subseteq Q \times (Bexp \times A \times Bcmd) \times Q$ is the set of transitions, where:

- *Bexp* is a *guard*, i.e. a Boolean expression constructed using the following grammar $\mathbf{b} ::= \mathbf{True} \mid \mathbf{x} \mid \mathbf{False} \mid \mathbf{b} \wedge \mathbf{b} \mid \neg \mathbf{b}$ (where $\mathbf{x} \in X_b$);
- *Bcmd* is either an assignment $\mathbf{x} := \mathbf{b}$ (where $\mathbf{x} \in X_b$, $\mathbf{b} \in Bexp$) or the null command **skip**.

As usual, we note $p \xrightarrow{(b,a,c)} q$ for $(p, (b, a, c), q) \in T$. We can omit b (resp. c) when b is **True** (resp. c is **skip**).

The semantics of a BLTS is given by a LTS. We define a notion of configuration (p, γ) , where p is a state of the BLTS and $\gamma : X_b \mapsto Bool$ a valuation, where $Bool = \{True, False\}$, is the set of Boolean values. Valuation γ are extended to *Bexp* in the usual way (i.e., $\gamma : Bexp \mapsto Bool$). Given a BLTS $B = (X_b, Q, A, T, q^0)$ and the initial valuation γ_0 , where $\gamma_0(x) = False$ for all $x \in X_b$, the underlying LTS $S_B = (Q_1, A, T_1, q_1^0)$ is defined as follows:

$$\begin{aligned} Q_1 &\subseteq Q \times Bool^{X_b}, \\ (p, \gamma) &\xrightarrow{a} (p_1, \gamma_1) \text{ iff } (p, (b, a, c), q) \in T \text{ and } \gamma(b) = True, \\ \gamma_1 &= \begin{cases} \gamma[v/x] & \text{if } c \text{ is } \mathbf{x} := \mathbf{e}, \gamma(\mathbf{e}) = v \text{ and } \gamma(b) = true \\ \gamma & \text{if } c \text{ is } \mathbf{skip} \end{cases} \\ q_1^0 &= (q^0, \gamma_0) \end{aligned}$$

3.1 Syntax of Security Rules

The formalism we adopt to express security policy rules is based on a variant of the LTL temporal logic [11], with only the $\mathcal{F}$ and $\mathcal{G}$ modalities, and mixing state-based and event-based atomic predicates. Each rule is expressed by a logical *formula* (φ), built upon *literals*. Each literal can be either a *condition literal* ($p_c \in P_c$), or an *event literal* ($p_e \in P_e$). A conjunction of condition literals is simply called a *condition* (C), whereas a conjunction of a single event literal and a condition is called a (*guarded*) *event* (E). Finally, we also use a modal operator $\mathcal{F}$, the dual one $\mathcal{G}$, and the usual Boolean connectors $\neg$ and $\Rightarrow$.

The abstract syntax of a formula is then given by the following grammar:

$$\begin{aligned}\varphi &::= C \mid E \mid \neg\varphi \mid \varphi \Rightarrow \varphi \mid \mathcal{F}\varphi \mid \mathcal{G}\varphi \\ C &::= p_c^1 \wedge \dots \wedge p_c^n \quad \text{where } p_c^i \in P_c \\ E &::= p_e[C] \quad \text{where } p_e \in P_e\end{aligned}$$

3.2 Semantics

Formulas are interpreted over LTS. Intuitively, an LTS S satisfies a formula φ iff *all* its execution sequences ρ do, where condition literals are interpreted over *states*, event literals are interpreted over *labels* and the modal operator $\mathcal{F}\varphi$ means that *there exists* a suffix $\rho_{[i..|\rho|]}$ of ρ such that φ holds on $\rho_{[i..|\rho|]}$, where $|\rho|$ is the number of elements of ρ . We first introduce two interpretation functions for condition and event literals:

$f_c : P_c \rightarrow 2^Q$, associates to p_c the set of states on which p_c holds;
 $f_e : P_e \rightarrow 2^A$, associates to p_e the set of labels on which p_e holds.

The satisfaction relation of a formula φ on an execution sequence ρ ($\rho \models \varphi$) is then (inductively) defined as follows:

- $\rho \models C$ for $C = p_c^1 \wedge \dots \wedge p_c^n$ iff $\forall i. \sigma^\rho(1) \in f(p_c^i)$
- $\rho \models E$ for $E = p_e[C]$ iff $\alpha^\rho(1) \in g(p_e) \wedge \sigma^\rho(2) \models C$
- $\rho \models \neg\varphi$ iff $\rho \not\models \varphi$
- $\rho \models \varphi_1 \Rightarrow \varphi_2$ iff $((\rho \models \varphi_1) \Rightarrow (\rho \models \varphi_2))$
- $\rho \models \mathcal{F}\varphi$ iff $\exists i \in [1, |\rho|]. \rho_{[i..|\rho|]} \models \varphi$
- $\rho \models \mathcal{G}\varphi$ iff $\forall i \in [1, |\rho|]. \rho_{[i..|\rho|]} \models \varphi$

Finally, $S \models \varphi$ iff $\forall \rho \in \Sigma_S. \rho \models \varphi$.

3.3 Expression of Security Rules

Our purpose is to use the specification language defined in the previous paragraph to express *security rules* to be satisfied by a network. In this particular context its semantics should be interpreted as follows:

- The network behaviour is expressed by the LTS S : each state of S represents the global state of the network at a given time (network configuration and topology, transiting PDUs, etc.), and each label of S represents an observable action performed at the network level (modification of the configuration/topology, PDU reception, PDU emission, etc.).
- A condition literal p_c expresses a (*static*) predicate on a network state, at the security policy level. For instance *externRelay(h1)* holds on a state iff machine *h1* is configured as an external mail relay in this state, or *infected(m1)* holds on a state iff message *m1* contains a virus.
- An event literal p_e expresses a (*dynamic*) predicate on a network transition, from a given state, at the security policy level. For instance *enterNetwork(m)* holds iff the current transition corresponds to reception of mail *m* by the network, and *chkVirus(m)* holds iff the current transition corresponds to a virus check on mail *m*.
- The $\mathcal{F}$ operator is used here to express an *obligation*, meaning that a given formula *should eventually* hold later, in a bounded future. For instance

$enterNetwork(m) \Rightarrow \mathcal{F}chkVirus(m)$ means that whenever mail m enters the network then it should be checked.

As a matter of fact, it happens that all the formulas found in the case study could be expressed using only a restricted subset of this language. In particular formulas can be classified into three types, according to the following grammar:

$$\begin{aligned}\varphi &::= \mathcal{G} \mathcal{C}\text{-Rule} \mid \mathcal{G} \mathcal{F}\text{-Rule} \mid \mathcal{G} \mathcal{G}\text{-Rule} \\ \mathcal{C}\text{-Rule} &::= C \Rightarrow C \mid E \Rightarrow C \\ \mathcal{F}\text{-Rule} &::= E \Rightarrow \mathcal{F}E \\ \mathcal{G}\text{-Rule} &::= C \Rightarrow \mathcal{G}\neg E \mid E \Rightarrow \mathcal{G}\neg E\end{aligned}$$

A **$\mathcal{C}$ -Rule** expresses a static conditional implication, an **$\mathcal{F}$ -Rule** expresses a (triggered) obligation and a **$\mathcal{G}$ -Rule** expresses that, when a given condition holds or when a given event happens, then a particular event is always *prohibited*.

4 Test Generation

In this section, we propose a “tile-based” approach to generate *abstract test cases* from a formula expressing a security rule. The test generation principle is the following: assuming that elementary test cases (i.e., *tiles*) t_i are provided for each (condition and event) literals appearing in a formula φ , the test case t associated to φ is obtained by combining test cases t_i with “test operators” (defined below), corresponding to the logical operators appearing in φ . This allows defining a structural correspondence between formulas and test cases.

4.1 Test Cases and Test Execution

We can define the notion of test case as a BLTS extended with two special actions (to deal with *timers*), and three special states (called *verdicts*): action *timerset* means a timer initialisation to a given value, and action *timeout* means the timer expiration; verdict states are “sink states”, indicating the end of a successful (*pass*), unsuccessful (*fail*) or inconclusive (*inconc*) test execution. We denote by A^t the set $A \cup \{\text{timeout}, \text{timerset}\}$, and by V the set $\{\text{pass}, \text{fail}, \text{inconc}\}$. We denote by Σ_S^{pass} (resp. Σ_S^{fail} , Σ_S^{inconc}) the set of execution sequences, starting from the initial state and ending in the state *pass* (resp. *fail*, *inconc*). A *test case* t is then a BLTS $t = (X_p, Q, A^t, T, q^0, V)$, with $V \subseteq Q$.

A test case is supposed to be executed by a tester against a network whose behaviour can be modelled by an LTS $I = (Q^I, A^I, T^I)$ (we ignore initial states for the network). Usually, in “black-box” conformance testing, this behaviour is observed/controlled by the tester only through a restricted interface. For the sake of simplicity we assume here that *any* output action (resp. input action) performed by the network can be observed (resp. controlled) by the tester. Thus, execution of a test t on an IUT I , noted $\text{Exec}(t, I)$, is expressed as a set of common execution sequences of S_t and I , defined by a composition operator $\otimes$:

Let $\rho_I = q_0^I \xrightarrow{a_1} q_1^I \xrightarrow{a_2} q_2^I \cdots \xrightarrow{a_n} q_n^I \cdots \in \Sigma_I$ and $\rho_{S_t} = q_0^{0,t} \xrightarrow{a_1} q_1^t \xrightarrow{a_2} q_2^t \cdots \xrightarrow{a_n} q_n^t \in \Sigma_{S_t}$, then

$$\rho_{S_t} \otimes \rho_I = (q^{0,t}, q_0^I) \xrightarrow{a_1} (q_1^t, q_1^I) \cdots \xrightarrow{a_n} (q_n^t, q_n^I) \in \text{Exec}(t, I).$$

For $\rho \in \text{Exec}(t, I)$, we define the verdict function: $\text{VExec}(\rho) = \text{pass}$ (resp. *fail*, *inconc*) iff there is $\rho_{S_t} \in \Sigma_{S_t}^{\text{pass}}$ (resp. $\Sigma_{S_t}^{\text{fail}}$, $\Sigma_{S_t}^{\text{inconc}}$) and $\rho_I \in \Sigma_I$ such that $\rho_{S_t} \otimes \rho_I = \rho$.

4.2 Test Generation Functions

Let φ be a formula, let p_e , p_c^i its literals, and t_{p_e} , $t_{p_c^i}$ their corresponding elementary test cases. Note that an elementary test case is reduced to a simple verdict state when it corresponds to a literal that can be immediately checked on the network behaviour (without requiring any interaction sequence with a tester). The test generation function $\text{gentest}(\varphi)$ is inductively defined on the syntax of the formula, where X_1 denotes either a condition C or an event E , and test operators $\triangleright_{\text{rrF}}$, $\triangleright_{\text{lrF}}$, $\triangleright_{\text{rrC}}$, $\triangleright_{\text{lrC}}$, $\triangleright_{\text{llX}}$ and Inv are defined below:

$$\begin{aligned} \text{gentest}(X_1 \Rightarrow C_2) &= \text{gentest}_{\text{llX}}(X_1) \triangleright_{\text{lrC}} \text{gentest}_{\text{rrC}}(C_2) \\ \text{gentest}(E_1 \Rightarrow \mathcal{F}E_2) &= \text{gentest}_{\text{llX}}(E_1) \triangleright_{\text{lrF}} \text{gentest}_{\text{rrF}}(E_2) \\ \text{gentest}(X_1 \Rightarrow \mathcal{G}\neg E_2) &= \text{Inv}(\text{gentest}_{\text{llX}}(X_1) \triangleright_{\text{lrF}} \text{gentest}_{\text{rrF}}(E_2)) \\ \text{gentest}_{\text{llX}}(p_c^1 \wedge \cdots \wedge p_c^n) &= (((t_{p_c^1} \triangleright_{\text{llX}} t_{p_c^2}) \triangleright_{\text{llX}} \cdots) \triangleright_{\text{llX}} t_{p_c^n}) \\ \text{gentest}_{\text{llX}}(p_e[C]) &= t_e \triangleright_{\text{llX}} \text{gentest}_{\text{llX}}(C) \\ \text{gentest}_{\text{rrC}}(p_c^1 \wedge \cdots \wedge p_c^n) &= (((t_{p_c^1} \triangleright_{\text{rrC}} t_{p_c^2}) \triangleright_{\text{rrC}} \cdots) \triangleright_{\text{rrC}} t_{p_c^n}) \\ \text{gentest}_{\text{rrC}}(p_e[C]) &= t_e \triangleright_{\text{rrC}} \text{gentest}_{\text{rrC}}(C) \\ \text{gentest}_{\text{rrF}}(p_e[C]) &= t_e \triangleright_{\text{rrF}} \text{gentest}_{\text{rrF}}(C) \end{aligned}$$

4.3 Test Operators

In the following we assume that $t_1 = (X_{b_1}, Q_1, A_1, T_1, q_1^0, \{\text{pass}_1, \text{fail}_1, \text{inconc}_1\})$ and $t_2 = (X_{b_2}, Q_2, A_2, T_2, q_2^0, \{\text{pass}_2, \text{fail}_2, \text{inconc}_2\})$ are two test cases. For each binary test operator $\triangleright$ we define the test case $t = (X_b, Q, A, T, q^0, V)$ such that $t = t_1 \triangleright t_2$ and $V = \{\text{pass}, \text{fail}, \text{inconc}\}$. For each operator, a graphical presentation is proposed on figure 2 and we give hereafter the formal definition of only three operators, the three others being similar.

Operator $\triangleright_{\text{llX}}$ ($t = t_1 \triangleright_{\text{llX}} t_2$). This operator is used to combine two test cases appearing on the left-hand side of an implication. Therefore, t *pass* iff t_1 and t_2 does, and t is *inconclusive* otherwise (the entire formula cannot be tested). More formally:

$$\begin{aligned} X_b &= X_{b_1} \cup X_{b_2}, \\ Q &= (Q_1 \setminus V_1) \cup (Q_2 \setminus V_2) \cup V, \\ A &= A_1 \cup A_2, \\ q^0 &= q_1^0, \\ T &= T_1 \setminus \{(p, a, q) \mid q \in V_1\} \cup T_2 \setminus \{(p, a, q) \mid q \in V_2\} \\ &\quad \cup \{(p, a, q_2^0) \mid (p, a, \text{pass}_1) \in T_1\} \cup \{(p, a, \text{pass}) \mid (p, a, \text{pass}_2) \in T_2\} \\ &\quad \cup \{(p, a, \text{inconc}) \mid (p, a, q) \in T_1 \wedge q \in \{\text{inconc}_1, \text{fail}_1\}\} \\ &\quad \cup \{(p, a, \text{inconc}) \mid (p, a, q) \in T_2 \wedge q \in \{\text{inconc}_2, \text{fail}_2\}\} \end{aligned} \quad (\alpha)$$

Fig. 2. Test operators

Operator $\triangleright_{IRC}$ ($t = t_1 \triangleright_{IRC} t_2$). This operator is used to combine the left-hand side (t_1) and the right-hand side (t_2) of an implication for a $\mathcal{C}$ -Rule. Therefore t can be expressed by a sequential execution of t_1 and t_2 , and it *pass* iff t_1 and t_2 does, it *fails* when t_1 pass and t_2 fails (implication), and it is *inconclusive* otherwise. Formal definition of test t is then simply obtained by replacing the line (α) of the previous definition (operator $\triangleright_{IIX}$) by the following one:

$$\{(p, a, inconc) \mid (p, a, inconc_2) \in T_2\} \cup \{(p, a, fail) \mid (p, a, fail_2) \in T_2\}$$

Operator $\triangleright_{rrC}$ ($t = t_1 \triangleright_{rrC} t_2$). This operator is used to combine two test cases appearing on the right-hand side of an implication, for a $\mathcal{C}$ -Rule. Therefore t *pass* when t_1 and t_2 does, t *fails* when t_1 or t_2 fails, and it is *inconclusive* otherwise. Thus, t can be obtained by executing t_1 first, followed by t_2 (when t_1 does not fail). A Boolean variable $t1pass$ is used to store the verdict of t_1 .

Operator $\triangleright_{lrF}$ ($t = t_1 \triangleright_{lrF} t_2$). This operator is used to combine the left-hand side (t_1) and the right-hand side (t_2) of an implication for an $\mathcal{F}$ -Rule. It is therefore similar to the $\triangleright_{IRC}$ operator, excepted that, due to the $\mathcal{F}$ temporal modality, t *pass* iff t_1 pass and t_2 pass *at some point later in the future* (remember that the right-hand side of an $\mathcal{F}$ -Rule is necessarily an *event*). t is then obtained by executing t_1 first, and then repeatedly executing t_2 until either it

passes or a given timeout is reached (to ensure that the test execution remains finite). A Boolean variable $\mathbf{t2inconc}$ is used to keep track of an occurrence of an inconclusive verdict of t_2 (during its repeated execution), hence leading to an inconclusive verdict of t . More formally:

$$\begin{aligned}
X_b &= X_{b_1} \cup X_{b_2} \cup \{\mathbf{t2inconc}\}, \\
Q &= (Q_1 \setminus V_1) \cup (Q_2 \setminus V_2) \cup V \cup \{q_0\}, \\
A &= A_1 \cup A_2, \\
T &= T_1 \setminus \{(p, a, q) \mid q \in V_1\} \cup T_2 \setminus \{(p, a, q) \mid q \in V_2\} \\
&\quad \cup \{(q_0, (\mathbf{timerset}, \mathbf{tinconc} := \mathbf{False}), q_1^0)\} \\
&\quad \cup \{(p, a, \mathbf{inconc}) \mid (p, a, q) \in T_1 \wedge q \in \{\mathbf{inconc}_1, \mathbf{fail}_1\}\} \\
&\quad \cup \{(p, a, q_2^0) \mid (p, a, \mathbf{pass}_1) \in T_1\} \\
&\quad \cup \{(p, a, q_2^0) \mid (p, a, \mathbf{fail}_2) \in T_2\} \cup \{(p, a, \mathbf{pass}) \mid (p, a, \mathbf{pass}_2) \in T_2\} \\
&\quad \cup \{(p, a, \mathbf{t2inconc} := \mathbf{True}), q_2^0\} \mid (p, a, \mathbf{inconc}_2) \in T_2\} \\
&\quad \cup \{(q_2^0, (\neg \mathbf{t2inconc}, \mathbf{timeout}), \mathbf{fail})\} \cup \{(q_2^0, (\mathbf{t2inconc}, \mathbf{timeout}), \mathbf{inconc})\}
\end{aligned}$$

Operator $\triangleright_{\mathbf{rrF}}$ ($t = t_1 \triangleright_{\mathbf{rrF}} t_2$). This operator is used to combine two test cases appearing on the right-hand side of an implication for an $\mathcal{F}$ -**Rule**, where t_1 tests the occurrence of an *event literal* p_e and t_2 a (static) *condition* C (possibly restricting p_e). Therefore, t *pass* iff both the expected event occurs at some point (t_1 *pass*) *and* condition C holds on the same time. t is then obtained by repeating t_1 followed by t_2 until both pass. Here again, a timeout ensures that execution of t always remains finite, and a Boolean variable $\mathbf{tinconc}$ is used to keep track of an inconclusive verdict of t_1 or t_2 .

Operator $\mathbf{Inv}$ ($t = \mathbf{Inv}(t_1)$). This operator simply “reverts” the *pass* and *fail* verdicts produced by a test case.

4.4 Soundness of the Test Generation Function

It now remains to establish that an abstract test case produced by function $\text{gentest}(\varphi)$ is always *sound*, i.e., it delivers a *fail* verdict when executed on a network behaviour I only if formula φ does not hold on I .

Two hypotheses are required in order to prove this soundness property:

H1. First, for any formula φ , we assume that each elementary test case t_i provided for the (event or condition) literals p_i appearing in φ is *strongly sound* in the following sense:

$$\forall \rho \in \text{Exec}(t_i, I). (\forall \text{Exec}(\rho) = \text{Pass} \Rightarrow \rho \models p_i) \wedge (\forall \text{Exec}(\rho) = \text{Fail} \Rightarrow \rho \not\models p_i)$$

H2. Second, we assume that the whole execution of a (provided or generated) test case t associated to a condition C is *stable* with respect to condition literals: the valuation of these literal does not change during the test execution. This simply means that the network configuration is supposed to remain stable when a condition is tested. Formally:

$$\forall p_i \in P_c. \forall \rho \in \Sigma_I \cdot \rho_{S_t} \otimes \rho \in \text{Exec}(t, I) \Rightarrow (\sigma^\rho \subseteq f_c(p_i) \vee \sigma^\rho \cap f_c(p_i) = \emptyset)$$

where σ^ρ denotes here tacitly a set of states instead of a sequence.

We now formulate the soundness property:

Proposition: Let φ a formula, I an LTS and $t = \text{gentest}(\varphi)$. Then:

$$\rho \in \text{Exec}(t, I) \wedge \text{VExec}(\rho) = \text{fail} \implies I \not\models \varphi.$$

The proof of this proposition relies on the following lemma:

Lemma 1. Test cases generated by auxiliary function gentest_lX are *strongly sound*, and test cases generated by auxiliary functions gentest_rC and gentest_rF are *sound*.

Proof sketch of Lemma 1. Let t a test case generated by function gentest_lX . The proof that t is *strongly sound* is performed by recurrence on the number of elementary tests cases t_i appearing in t (assuming that each t_i itself is *strongly sound* according to hypothesis **H1**). A similar proof can be done for functions gentest_rC and gentest_rF .

Proof of Proposition. By structural induction on the formulas φ (we only detail here some representative induction steps).

- $\varphi = C_1 \Rightarrow C_2$. By definition of function gentest there exists test cases t_1 and t_2 such that $t_1 = \text{gentest_lX}(C_1)$, $t_2 = \text{gentest_rC}(C_2)$, and $t = t_1 \triangleright_{\text{lrc}} t_2$. Let ρ be an execution sequence of $\text{Exec}(t, I)$ such that $\text{VExec}(\rho) = \text{fail}$. Then, by definition of operator $\triangleright_{\text{lrc}}$, there exist ρ_1 and ρ_2 such that: $\rho = \rho_1 \cdot \rho_2$, $\rho_1 \in \text{Exec}(t_1, I)$, $\rho_2 \in \text{Exec}(t_2, I)$, $\text{VExec}(\rho_1) = \text{pass}_1$ and $\text{VExec}(\rho_2) = \text{fail}_2$. Therefore, by Lemma 1, $\rho_1 \models C_1$, hence $\sigma^\rho(1) \models C_1$ and similarly $\rho_2 \not\models C_2$ and $\sigma^\rho(|\rho_1|) \not\models C_2$. By hypothesis **H2** we obtain $\sigma^\rho(1) \not\models C_2$ and then $\sigma^\rho(1) \not\models \varphi$.
- $\varphi = E_1 \Rightarrow \mathcal{F}E_2$. By definition of function gentest there exist test cases t_1 and t_2 such that $t_1 = \text{gentest_lX}(E_1)$, $t_2 = \text{gentest_rF}(E_2)$, and $t = t_1 \triangleright_{\text{lrf}} t_2$. Let ρ be an execution sequence of $\text{Exec}(t, I)$ such that $\text{VExec}(\rho) = \text{fail}$. Then, by definition of operator $\triangleright_{\text{lrf}}$, there exist ρ_1, ρ_2 such that: $\rho = q_0 \xrightarrow{a_0} q_0^1 \xrightarrow{\rho_1} \text{pass}_1 (\xrightarrow{\rho_2} \text{fail}_2)^* \xrightarrow{a_1} q_0^2 \xrightarrow{a_2} \text{fail}$, with $a_0 = (\text{timerset}, \text{t2inconc} := \text{False})$, $a_1 = \text{t2inconc} := \text{False}$, and $a_2 = (\text{timeout}, \neg \text{t2inconc})$. Moreover, $\rho_1 \in \text{Exec}(t_1, I)$ and $\rho_2 \in \text{Exec}(t_2, I)$. Therefore (by Lemma 1) $\sigma^\rho(1) \models E_1$ and, for all i in $[|\rho_1|, |\rho|]$, $\sigma^\rho(i) \not\models E_2$. We conclude that $\sigma^\rho(1) \not\models (E_1 \Rightarrow \mathcal{F}E_2)$.

5 Case Study Application

This section shows how the approach presented above can be applied to generate concrete tests for some examples from the case study of section 2.1.

5.1 C–Rule

Consider the requirement “*External relays shall be in the DMZ*”, which can be modelled by the **C–Rule**:

$$\text{externRelay}(h) \Rightarrow \text{inDMZ}(h)$$

The goal of the test is to verify that each external relay is in the DMZ. As noted in section 4.2, an elementary test case is reduced to a simple verdict state when it corresponds to a literal that can be checked without requiring an interaction sequence. Such a case arises when the value can be checked by an analysis of the configuration of devices in the network and/or administrators' databases. For example, if the value of $externRelay(h)$ is true, that means that h is defined as an external relay in the administrators' database and/or by the configuration of the network. This is known and trusted, not to be tested.

On the other side, the value is unsure if one has no knowledge about the fact that h is an external relay from the analysis of configurations, or if these data are untrusted. In this case the behaviour of the network should be tested to decide whether h acts as an external relay.

The following table shows the different formulas that may be built depending on which literals can be immediately asserted:

	$externRelay(h)=true$	$... = false$	$... unsure$
$inDMZ(h)=true$	t_{pass}	t_{inconc}	$t_{externRelay(h)} \triangleright_{lrc} t_{pass}$
$inDMZ(h)=false$	t_{fail}	t_{inconc}	$t_{externRelay(h)} \triangleright_{lrc} t_{fail}$
$inDMZ(h) unsure$	$t_{pass} \triangleright_{lrc} t_{inDMZ(h)}$	t_{inconc}	$t_{externRelay(h)} \triangleright_{lrc} t_{inDMZ(h)}$

If both values of $externRelay(h)$ and $inDMZ(h)$ are known and trusted, there is nothing to test. Also, no test is needed if $externRelay(h)$ is *false*, as we cannot put the system in the desired state, and the verdict is *inconc*. If the value of $inDMZ(h)$ (resp. $externRelay(h)$) is unsure, then it should be tested whether h behave like a host in the DMZ (resp. an external relay). These tests are then composed as described in section 4 into the formula $t_{externRelay(h)} \triangleright_{lrc} t_{inDMZ(h)}$, as illustrated in figure 3.

5.2 $\mathcal{F}$ -Rule

Consider the requirement “If an electronic mail is infected by a virus, the virus shall be deleted from the mail”. It can be modelled by the $\mathcal{F}$ -Rule rule:

$$\begin{aligned}
 &enterNetwork(m)[infected(m)] \Rightarrow \\
 &\mathcal{F} \ transfer(h_1, h_2, m)[interior(h_2) \wedge \neg infected(m)]
 \end{aligned}$$

The goal of this test is to verify that if a mail infected by a virus is sent to a user in the network, eventually one of the hosts the mail is passing through will suppress the virus, before a certain time elapses.

As always in our approach, a choice is made concerning which predicates are sure or unsure. The formula $t = (t_{enterNetwork(m)} \triangleright_{llx} t_{pass}) \triangleright_{lrf} (t_{transfer(h_1, h_2, m)} \triangleright_{rrf} (t_{pass} \triangleright_{rrf} t_{\neg infected(m)}))$ corresponds to the case when we build a test tile with a parameter m made of an infected message. This is the case because we choose to actively test the conformity of this particular rule against infected messages. One could also use a “passive” mode, checking the $infected(m)$ literal in the left part of the formula on incoming messages.

Fig. 3. Composition for example of $\mathcal{C}$ -Rule

On the other hand the literal $\neg \text{infected}(m)$ in the right part shall be tested. The event predicate $\text{enterNetwork}(m)$ and the static predicate $\neg \text{infected}(m)$ are tested by the tiles shown in figure 4. Using these tiles, the formula $t_{\text{enterNetwork}(m)} \triangleright_{\text{IrF}} (t_{\text{transfer}(m)} \triangleright_{\text{rrF}} t_{\neg \text{infected}(m)})$ gives the test on figure 4. The tinconc and t2inconc variables have been suppressed since they cannot be true, and also the corresponding transitions.

6 Perspectives

In this paper, we have proposed a “tile-based” approach to derive test cases from rules expressed using a restricted set of logical operators that can be applied to network security policies. Complete test cases (dedicated to a whole formula) are obtained by combinations of more elementary ones (the tiles), following a syntax driven approach (a test combinator is associated to each logical operator of the formula). Elementary test cases, allowing to test basic events or predicates appearing in the security policy, have to be provided by the user (the way of testing such predicate depends on the network architecture and protocols involved). Our test generation method is based on the fact that security policies are most of the time expressed by rules which can be captured by a restricted logic as the one we described in section 3.

This approach is not limited to security, the idea of combining test tiles into a new one can be applied to other domains, for example testing software architecture, or testing systems built from existing components. However the com-

Fig. 4. Composition for example of $\mathcal{F}$ -Rule

combination operators are based on the form of the rule being tested, and thus may be dependent on the domain. Up to now our examples use general purpose modalities, but we may have to consider more specialised cases.

At this point this work should be viewed essentially as a first step towards a formal approach to (automatically) test the compliance of a network with a given security policy. Therefore it should be extended into several directions.

First of all, the test cases produced are still very *abstract*. Turning them into executable test cases needs to take into consideration a concrete test architecture. Assuming that each elementary test case complies with this architecture, it would remain to ensure that it is also the case for the complete test case (or alternatively to take this architecture into account during the combination process). Moreover, these abstract test cases also need to be *instantiated* with concrete data (e.g. by selecting particular machines of the network). Suitable selection strategies should therefore be investigated (for instance a test could focus on the more recent changes in a network configuration, as in regression testing).

Furthermore, the proposed generation technique itself could be improved. In particular, the test case currently produced to test a condition (i.e., a disjunction of static predicates) consists in executing each corresponding elementary test case in sequence (according to the definition of our test combinators). An alternative way would have been to consider the *parallel* execution of such test cases (when it is compatible with the test architecture).

Another improvement could be to extend the formalism we considered to specify the security rules. This initial choice was motivated by our case study, and it was sufficient to demonstrate the effectiveness of the approach. However, it is clear that this formalism may be not sufficient to deal with arbitrary security rules, and that more specific operator/modalities need to be considered. One can think for instance of a triggered obligation bounded by an event (and not by an arbitrary timeout), or of some of the general operators proposed in the

NOMAD logic [3]. Further work remains to be done in order to check which of these operators could be supported by our tile-based approach.

Finally, we also intend to evaluate this work on other case studies, and to prototype it on a real network.

Acknowledgements. The authors thank Keqin Li for giving valuable comments on this paper.

References

1. A. Abou El Kalam, R. El Baida, P. Balbiani, S. Benferhat, F. Cuppens, Y. Deswarte, A. Miège, C. Saurel, and G. Trouessin. Organization Based Access Control. In *IEEE 4th International Workshop on Policies for Distributed Systems and Networks*, 2003.
2. S. Benferhat, R. E. Baida, and F. Cuppens. A Stratification-Based Approach for Handling Conflicts in Access Control. In *8th ACM Symposium on Access Control Models and Technologies*, 2003.
3. F. Cuppens, N. Cuppens-Boulahia, and T. Sans. Nomad: A security model with non atomic actions and deadlines. In *18th IEEE Computer Security Foundations Workshop, (CSFW-18 2005)*, pages 186–196, Aix-en-Provence, France, 2005.
4. F. Cuppens, N. Cuppens-Boulahia, T. Sans, and A. Miège. A formal approach to specify and deploy a network security policy. In *Second Workshop on Formal Aspects in Security and Trust (FAST)*, 2004.
5. N. Damianou, N. Dulay, E. Lupu, and M. Sloman. The Ponder Policy Specification Language. In *International Workshop on Policies for Distributed Systems and Networks*, 2001.
6. V. Darmaillacq, J.-C. Fernandez, R. Groz, L. Mounier, and J.-L. Richier. Éléments de modélisation pour le test de politiques de sécurité. In *Colloque sur les RISques et la Sécurité d'Internet et des Systèmes, CRiSIS*, Bourges, France, 2005.
7. V. Darmaillacq and N. Stouls. Développement formel d'un moniteur détectant les violations de politiques de sécurité de réseaux. In *AFADL2006 - Approches Formelles dans l'Assistance au Développement de Logiciels*, Paris, March 2006.
8. ITU. Framework on formal methods in conformance testing. ITU-T Recommendation Z.500, ITU, 1997.
9. J. Lobo, R. Bhatia, and S. Naqvi. A Policy Description Language. In *AAAI'99*, 1999.
10. J.-C. Meyer, F. Dignum, and R. Wieringa. The Paradoxes of Deontic Logic Revisited: A Computer Science Perspective. Technical Report UU-CS-1994-38, Utrecht University, 1994.
11. A. Pnueli. The Temporal Logic of Programs. In I. C. S. Press, editor, *18th Annual Symposium on Foundations of Computer Science*, 1977.
12. D. Senn, D. Basin, and G. Caronni. Firewall Conformance Testing. In *TestCom 2005, 17th IFIP TC6/WG6.1 International Conference on Testing of Communicating Systems, Montréal, LNCS 3502*, June 2005.
13. A. Virmani, J. Lobo, and M. Kohli. Netmon: Network Management for the SARAS Softswitch. In *IEEE/IFIP Network Operations and Management Symposium*, 2000.
14. G. H. von Wright. Deontic Logic. *Mind*, 60:1–15, 1951.