

---

## DISCOVERING, QUANTIFYING, AND DISPLAYING ATTACKS \*

ROBERTO VIGO, FLEMMING NIELSON, AND HANNE RIIS NIELSON

DTU Compute, Technical University of Denmark, Richard Petersens Plads Bldg. 324, DK-2800  
Kongens Lyngby, Denmark  
*e-mail address:* {rvig, fnie, hrni}@dtu.dk

---

**ABSTRACT.** In the design of software and cyber-physical systems, security is often perceived as a qualitative need, but can only be attained quantitatively. Especially when distributed components are involved, it is hard to predict and confront all possible attacks. A main challenge in the development of complex systems is therefore to discover attacks, quantify them to comprehend their likelihood, and communicate them to non-experts for facilitating the decision process.

To address this three-sided challenge we propose a protection analysis over the Quality Calculus that (i) computes all the sets of data required by an attacker to reach a given location in a system, (ii) determines the cheapest set of such attacks for a given notion of cost, and (iii) derives an attack tree that displays the attacks graphically.

The protection analysis is first developed in a qualitative setting, and then extended to quantitative settings following an approach applicable to a great many contexts. The quantitative formulation is implemented as an optimisation problem encoded into Satisfiability Modulo Theories, allowing us to deal with complex cost structures. The usefulness of the framework is demonstrated on a national-scale authentication system, studied through a Java implementation of the framework.

---

2012 ACM CCS: [Theory of computation]: Semantics and reasoning.

*Key words and phrases:* Attack tree, protection analysis, Quality Calculus, Satisfiability Modulo Theories, security cost.

\* This is a revised and extended version of [44].

This work is supported by the IDEA4CPS project, granted by the Danish Research Foundations for Basic Research (DNRF86-10).

## CONTENTS

1. Introduction	3
Contribution	4
Organisation of the paper	5
2. The Modelling Framework	6
2.1. Syntax	6
2.2. Semantics	8
2.3. Security model	10
3. The Danish NemID System	11
4. Discovering Attacks	12
4.1. From processes to flow constraints	13
4.2. Modelling the attacker	14
4.3. A SAT-based solution technique	15
4.4. Translating NemID	16
4.5. Modularity and refinement	17
5. Quantifying Attacks	18
5.1. Security labels	18
5.2. From qualitative to quantitative considerations	19
5.3. Optimisation Modulo Theories	20
5.4. Attacking NemID	22
5.5. Complex cost structures	24
5.6. A semantic interpretation of guessing	24
6. Displaying Attacks	25
6.1. Synthesising attack trees	26
6.2. The attack tree for NemID	27
6.3. Global stopping criterion	28
7. Implementation of the Analysis	29
7.1. Discussion	29
7.2. The Quality Protection Tool	30
8. Related Work	31
8.1. Protection analysis	31
8.2. Attack trees	31
9. Conclusion	32
Acknowledgement	33
References	33
Appendix A. Correctness of the Protection Analysis	35
Appendix B. Properties of Attack Trees	37
Appendix C. First-Order Attack Trees	38

## 1. INTRODUCTION

Our daily life is increasingly governed by software and cyber-physical systems, which are exploited in the realisation of critical infrastructure, whose security is a public concern. In the design of such systems, security is often perceived as a qualitative need, but experience tells that it can only be attained quantitatively. Especially when distributed components are involved, it is hard to predict and confront all possible attacks. Even if it were possible, it would be unrealistic to have an unlimited budget to implement all security mechanisms.

A main challenge in the development of complex systems is therefore to *discover* attacks, *quantify* them to comprehend their likelihood, and *communicate* them to decision-makers for facilitating deciding what countermeasures should be undertaken.

In order to address this three-sided challenge we propose a *protection analysis* over the Quality Calculus [35]. The calculus offers a succinct modelling language for specifying distributed systems with a high degree of interaction among components, which thus enjoy a highly-branching control flow. Moreover, process-algebraic languages have proven useful for describing software systems, organisations, and physical infrastructure in a uniform manner.

In the study of security, the compromise for obtaining such broad domain coverage while retaining a reasonable expressive power takes place at the level of attack definition. At a high level of abstraction, an attack can be defined as a sequence of actions undertaken by an adversary in order to make unauthorised use of a target asset. In a process-algebraic world, this necessary *interaction* between the adversary and the target system is construed in terms of communicating processes. Input actions on the system side can be thought of as *security checks* that require some information to be fulfilled. In particular, the capability of communicating over a given channel requires the knowledge of the channel itself and of the communication standard.

Hence, we shift the semantic load of secure communication onto channels, assuming that they provide given degrees of security. Whilst the correctness of such *secure channels* [27] is investigated in the realm of protocol verification, any system with a substantial need for security is likely to have standardised mechanisms to achieve various degrees of protection, and modelling them with secure channels is a coarse yet reasonable abstraction. An attack can be thus construed as a set of channels that fulfil the security checks (input actions) on a path leading to the target, i.e., a program point in the process under study.

The first task of the protection analysis, *discovering* all the attacks leading to a given location  $l$  in a process  $P$ , in this setting reduces to compute all sets of channels over which communication must take place in order for  $P$  to evolve to  $l$ . Technically, a process is translated into a set of propositional constraints that relate the knowledge of channels to communication actions and to reachability of program points. Each model of such set of constraints contains an attack. This perspective corresponds to a *qualitative* formulation of the analysis, in which all attacks are deemed equally interesting.

Nonetheless, secure channels allow modelling a great many domains. In IT systems, a channel can be thought of as a wired or wireless communication link over which messages are encrypted, and its knowledge mimics the knowledge of a suitable encryption key. In the physical world, a channel can represent a door, and its knowledge the ownership of the key, a pin code, or the capability of bypassing a retinal scan. This urges to acknowledge the variety of protection mechanisms and devise *attack metrics* by assigning costs to channels, thus quantifying the effort required to an attacker for obtaining a channel or, equivalently, the protection ensured by the security mechanism represented by a channel. As a consequence,

attacks can be ordered according to their cost. Finally, a conservative approach to security demands to look for attacks that are *optimal* in the cost ordering (that is, minimal or maximal, according to the notion of cost we adopt – in the following we shall focus without loss of generality on minimality).

Quantifying attacks allows to rule out those we deem unrealistic because exceeding the resources available to, or the skills of, a given attacker profile. Moreover, such a quantification allows to contrast the *desired security architecture* of a system, defined as a map from locations to security levels, with the *actual implementation of security*, defined as the protection deployed to guard a location  $l$ , which the analysis computes as the minimal cost required to reach  $l$ .

The second task of the protection analysis, *quantifying* the attacks leading to  $l$ , hence reduces to compute the costs of the attacks detected by the qualitative analysis. For the sake of performance, however, the naive approach that first computes all attacks (models) and then sorts them is not advisable. We instead implement the quest for minimal models on top of a Satisfiability Modulo Theories (SMT) solver, where minimality is sought with respect to an objective function defined on a liberal cost structure in which both symbolic and non-linearly-ordered cost sets can be represented.

The third task of the analysis, *displaying* attacks graphically, is achieved by backward chaining the set of constraints into which a process is translated, thus obtaining an *attack tree* whose root is the target location, leaves are channels, and internal nodes show how to combine sub-trees to attain the goal. Such a tree is denoted by a propositional formula whose minimal models can be computed with the SMT procedure mentioned above. Hence, our attack trees encompass both the qualitative and quantitative developments at once: a tree contains all the attacks leading to a given target, and the quest for the minimal ones can be carried out on its propositional denotation.

A Java implementation of the framework is available, which takes as input a process in the Value-Passing Quality Calculus, the location of interest, and the map from channels to costs. The tool computes the cheapest sets of atomic attacks leading to the goal, and produces an attack tree in which all attacks are summarised. The usefulness of the framework is demonstrated on the study of the NemID system, a national-scale authentication system used in Denmark to provide secure Internet communication between citizens and public institutions as well as private companies.

This work is a revised and extended version of [44]. The analysis has been reformulated for highlighting its dual nature in terms of attack discovery and quantification, and has been combined with the insights of [43] on generating attack trees. The main novelty of the paper lies in the harmonisation of two distinct results, offering a coherent framework where the quantitative study of security and the generation of graphical models inform each other.

**Contribution.** The novelty of the protection analysis is many-fold. First of all, the problem of inferring attacks and quantifying the protection offered by security checks is interesting *per se* and poorly addressed in the literature. In particular, we devise a comprehensive approach, from modelling systems and their security architecture to mechanising the verification of how this architecture has been realised in the implementation.

The extension of the analysis from qualitative to quantitative settings can be mimicked in a great many contexts, and it is seamlessly applicable to formal specifications that resort to encoding into logic formulae. In connection with this, our SMT-based solution procedure

can be applied to all problems requiring to rank models of a logic formula according to given criteria.

The quest for optimal attacks can cope with arbitrary cost structures and objective functions, whose shape is only limited by the expressiveness of modern SMT solvers. Whilst our SMT approach to optimisation is not entirely novel, non-linearly-ordered and symbolic cost structures have not been addressed so far in connection with this technique.

Finally, we overcome a number of shortcomings of existing methods for the automated generation of attack trees. These techniques typically suffer from focusing on computer networks, and thus suggest modelling languages tailored to this domain (cf. § 8.2). Moreover, the model-checking techniques that have been proposed recently for generating attack trees lead to an exponential explosion of the state space, limiting the applicability of automated search procedures.

In order to overcome these drawbacks, the protection analysis derives attack trees from process-algebraic specifications in a syntax-directed fashion. As we have already put forward, process calculi have proven useful in describing a variety of domains, well-beyond computer networks. On the complexity side, being syntax-driven, static analysis often enjoys better scalability than model-checking approaches. Even though the theoretical complexity of our analysis is still exponential in the worst case, such a price depends on the shape of the process under study, and is not incurred systematically. Interestingly enough, the generation of attack trees seems to boost the performance of the approach, though it is not necessary to computing minimal attacks. In general, the interplay between static analysis and model checking had a key role in advancing the community’s knowledge on the foundations of formal verification, and therefore it is promising to complement the existing studies on the attack tree generation problem with static analysis tools.

**Organisation of the paper.** First, we present our modelling framework based on the Value-Passing Quality Calculus in § 2. Besides syntax (§ 2.1) and semantics (§ 2.2) of the calculus, here we introduce our security interpretation of communication actions and the induced attacker model (§ 2.3). The NemID system, which will be our running example throughout the paper, is described in § 3.

The development of the protection analysis spans §§ 4, 5, 6. First, in § 4 we tackle the problem of discovering attacks leading to a location  $l$  in a process  $P$ . To this end, we translate  $P$  into a set of propositional constraints (§ 4.1), which we then extend to encompass the attacker model (§ 4.2). The solution procedure (§ 4.3) computes all the models of the constraints and from these extracts the attacks. This qualitative formulation of the analysis is demonstrated on the NemID system in § 4.4, while in § 4.5 we comment upon integrating the analysis in a refinement cycle.

In § 5 we extend the analysis to quantitative settings. First, we present how to formalise the intended security architecture of a system in terms of a security lattice (§ 5.1). Then, we introduce costs to channels so as to quantify attacks, and we relate the minimal costs for attacking a location to its desired security level, thus checking whether the desired security architecture is fulfilled by the implementation (§ 5.2). Our novel SMT-based approach to computing models of minimal cost is discussed in § 5.3 and demonstrated on the NemID system in § 5.4. The need for the SMT approach is motivated by the complex cost structures explained in § 5.5. We give a semantic interpretation of the attacker model in § 5.6.

Finally, the problem of generating attack trees is tackled in § 6. First, we show how to generate trees from our sets of constraints describing how to combine channels to reach the

target  $l$  (§ 6.1). The procedure leads to generate a formula describing all attacks to  $l$ , whose optimal models can be computed thanks to the SMT procedure. We discuss the attack tree for the NemID system in § 6.2.

In § 7 we discuss the implementation of the quantitative analysis. The solution approaches of §§ 5 and 6 are compared in § 7.1, and the proof-of-concept implementation is briefly presented in § 7.2.

Finally, § 8 surveys related work, organising them with respect to the analysis in general (§ 8.1) and to attack tree generation in particular (§ 8.2). We conclude and present our outlook on future work in § 9.

## 2. THE MODELLING FRAMEWORK

As specification formalism we adopt the Quality Calculus [35], a process calculus in the  $\pi$ -calculus family. In particular, the Quality Calculus introduces a new kind of input binders, termed *quality binders*, where a number of inputs are *simultaneously* active, and we can proceed as soon as some of them have been received, as dictated by a guard instrumenting the binder. Whilst these binders enhance the succinctness of highly-branching process-algebraic models, thus increasing their readability, they do not increase the expressiveness of the language with respect to  $\pi$ , and thus the protection analysis carries seamlessly to a variety of process calculi without such binders. Therefore, as far as this work is concerned, the reader can think of the calculus as a broadcast value-passing  $\pi$ -calculus enriched with quality binders.

**2.1. Syntax.** The syntax of the Value-Passing Quality Calculus is displayed in Table 1. The calculus consists of four syntactic categories: processes  $P$ , input binders  $b$ , terms  $t$ , and expressions  $e$ . A process can be prefixed by the restriction  $(\nu c)$  of a name  $c$ , by an input binder  $b$ , by an output  $c!t$  of term  $t$  on channel  $c$ , or be the terminated process  $0$ , the parallel composition of two processes, a replicated process, or a **case** clause.

A process  $c!t.P$  broadcasts term  $t$  over channel  $c$  and evolves to  $P$ : all processes ready to make an input on  $c$  will synchronise and receive the message, but we proceed to  $P$  even if there exists no such process (i.e., broadcast is non-blocking).

An input binder can either be a simple input  $c?x$  binding variable  $x$  to a message received over channel  $c$ , or a *quality binder*  $\&_q(b_1, \dots, b_n)$ , where the  $n$  sub-binders are *simultaneously* active. A quality binder is consumed as soon as enough sub-binders have been satisfied, as dictated by the quality guard  $q$ . A quality guard can be any Boolean predicate over the sub-binders: we use the abbreviations  $\forall$  and  $\exists$  for the predicates specifying that all the sub-binders or at least one sub-binder have to be satisfied before proceeding, respectively. For instance, the process  $\&\exists(c_1?x_1, c_2?x_2).P$  evolves to  $P$  as soon as an input is received on  $c_1$ , or an input is received on  $c_2$ .

When a quality binder is consumed, some input variables occurring in its sub-binders might have not been bound to any value, if this is allowed by the quality guard. Building on the example  $\&\exists(c_1?x_1, c_2?x_2).P$ , we can write:

$$(\nu c_1)(\nu c_2)(\nu c_3)(\&\exists(c_1?x_1, c_2?x_2).P|c_1!c_3.P')$$

Such a process evolves to  $P|P'$  even if no input is received on  $c_2$ . In order to record which inputs have been received and which have not, we always bind an input variable  $x$  to an expression  $e$ , which is  $\text{some}(c)$  if  $c$  is the name received by the input binding  $x$ , or  $\text{none}$  if the

Table 1: The syntax of the Value-Passing Quality Calculus.

$P$	$::=$	$0 \mid (\nu c)P \mid P_1 \mid P_2 \mid {}^l b.P \mid {}^l c!t.P$ $\mid !P \mid {}^l \text{case } x \text{ of some}(y): P_1 \text{ else } P_2$	(Process)
$b$	$::=$	$c?x \mid \&_q(b_1, \dots, b_n)$	(Binder)
$t$	$::=$	$c \mid y$	(Term)
$e$	$::=$	$x \mid \text{some}(t) \mid \text{none}$	(Expression)

input is not received but we are continuing anyway. In this sense, we say that expressions represent *optional data*, while terms represent *data*, in the wake of programming languages like Standard ML [26]. In order to insist on this distinction, variables  $x$  are used to mark places where expressions are expected, whereas variables  $y$  stand for terms, which are names (channels). Hence, in the example above variable  $x_1$  is replaced by  $\text{some}(c_3)$  and  $x_2$  is replaced by  $\text{none}$ .

The **case** clause in the syntax of processes is then used to inspect whether or not an input variable, bound to an expression, is indeed carrying data. The process  $\text{case } x \text{ of some}(y): P_1 \text{ else } P_2$  evolves into  $P_1[c/y]$  if  $x$  is bound to  $\text{some}(c)$ ; otherwise, if  $x$  is bound to  $\text{none}$ ,  $P_2$  is executed. Therefore, **case** clauses allow detecting which condition triggered passing a binder, by detecting on which channels communication took place.

In order to identify locations unambiguously, the syntax instruments each program point with a *unique* label  $l \in \mathcal{L}$ . This is the case of input binders, outputs, and case clauses in Table 1. In the following, we will denote labels with the numerals  $1, 2, \dots$ . We say that a label  $l$  is reached in an actual execution when the sub-process following  $l$  is ready to execute.

As usual, in the following we consider closed processes (no free variable), and we make the assumption that variables and names are bound exactly once, so as to simplify the technicalities without impairing the expressiveness of the framework.

The syntax presented above defines a fragment of the full calculus developed in [35]. First, we have described a value-passing calculus, as channels are syntactically restricted to names as opposed to terms; in particular, no variable  $y$  is allowed in a channel position. Second, terms are pure, that is, only names and variables are allowed (no function application). Last, the calculus is limited to test whether or not a given input variable carries data, but cannot test which data it is possibly carrying. We deem that these simplifications match with the abstraction level of the analysis, where secure channels are assumed that have to be established by executing given security protocols, and thus enjoy known properties. For the very same reason we opted for pure terms, since equational reasoning is usually exploited in process calculi to model cryptographic primitives (a Quality Calculus with such a feature is presented in [42]).

Ultimately, resorting to a value-passing calculus is to be understood as a choice that allows to produce attack trees of reasonable size, and not as a technical shortcoming. The comparison here is to be made between security protocols, often studied with expressive calculi (cf. § 8.1), and distributed systems. Security protocols can be understood as basic building blocks that we would like to prove flawless; to this end, modelling the structure of messages and the functions operating over them is necessary (and often not enough). On the other hand, such a low-level representation seems instead not suitable to modelling complex

Table 2: The structural congruence of the Value-Passing Quality-Calculus.

$P \equiv P$	$P   0 \equiv P$	$P_1   P_2 \equiv P_2   P_1$
$P_1   (P_2   P_3) \equiv (P_1   P_2)   P_3$	$!P \equiv P ! P$	$(\nu c_1) (\nu c_2) P \equiv (\nu c_2) (\nu c_1) P$
$(\nu c) P \equiv P \quad \text{if } c \notin \text{fc}(P)$	$(\nu c) (P_1   P_2) \equiv ((\nu c) P_1)   P_2 \quad \text{if } c \notin \text{fc}(P_2)$	
$\frac{P_1 \equiv P_2}{P_2 \equiv P_1}$	$\frac{P_1 \equiv P_2 \quad P_2 \equiv P_3}{P_1 \equiv P_3}$	$\frac{P_1 \equiv P_2}{C[P_1] \equiv C[P_2]}$

systems and to support analyses whose results can be computable in practice and presented in a human-readable format. In this sense, the coarse abstraction we adopt with flat channels is the price to pay for scaling up from analysing protocols to analysing complex systems, where a flat name can represent an entire sub-system (software, physical, cyber-physical) whose analysis is condensed in the price/security guarantee attached to it.

In this light, a reinforced gate can be represented with a channel  $g$ , and its cost be the minimal grams of dynamite required to blow it up, or the average price of the tools needed to force it, or a synthesis of both, depending on the attacker we want to consider. Similarly, a secure communication link providing secrecy can be represented with a channel  $s$ , and its cost be the length of the cryptographic key used for encrypting information sent over  $s$ , or the trust we put in the encryption protocol used to implement  $s$ . More complex cost notions can be used to describe security properties of cyber-physical devices, where for example the battery level can dictate the security degree a sensor can afford. While it is true that part of the burden is moved onto the definition of suitable cost structures (some solutions particularly useful to facilitate the task in case of cyber-physical systems are advanced in § 5.5), the framework offers a clear separation of concerns between modelling a complex system at an appropriate detail level, specifying the security contracts of its component, and computing a weak-path analysis.

An extension of the framework that encompasses the key features of the full calculus of [35] is briefly commented upon in Appendix C, and is a starting point to delve into more specialised notions of costs, where the analysis on channels is combined with an analysis on the structure of messages.

**2.2. Semantics.** The semantics of the Value-Passing Quality Calculus consists of a set of transition rules parametrised on the structural congruence relation of Table 2. In the congruence, we denote the free names in a process  $P$  as  $\text{fc}(P)$ , where  $(\nu c)P$  binds the name  $c$  in process  $P$ . Moreover, the congruence holds for contexts  $C$  defined as

$$C ::= [] \mid (\nu c)C \mid C|P \mid P|C$$

As usual in  $\pi$ -like calculi, processes are congruent under  $\alpha$ -conversion, and renaming is enforced whenever needed in order to avoid accidental capture of names during substitution.

The semantics of Table 3 is based on the transition relation  $P \Longrightarrow P'$ , which is enabled by combining a local transition  $\xrightarrow{\lambda}$  with the structural congruence.

The semantics models asynchronous broadcast communication, and makes use of a label  $\lambda ::= \tau \mid c_1!c_2$  to record whether or not a broadcast  $c_1!c_2$  is available in the system. If not, label  $\tau$  is used to denote a silent action. Rule (Brd) states that an output is a non-blocking



action, and when it is performed the broadcast is recorded on the arrow. When a process guarded by a binder receives an output, the broadcast remains available to other processes able to input. This behaviour is encoded in rules (In-ff) and (In-tt), where we distinguish the case in which a binder has not received enough inputs, and thus keeps waiting, from the case in which a binder is satisfied and thus the computation may proceed, applying the substitution induced by the received communication to the continuation process. These rules rely on two auxiliary relations, one defining how an output affects a binder, and one describing when a binder is satisfied (enough inputs have been received), displayed in the third and fourth section of Table 3, respectively.

We write  $c_1!c_2 \vdash b \rightarrow b'$  to denote that the availability of a broadcast makes binder  $b$  evolve into binder  $b'$ . When a broadcast name  $c_2$  is available on the channel over which an input is listening, the input variable  $x$  is bound to the expression  $\text{some}(c_2)$ , marking that something has been received. Otherwise, the input is left syntactically unchanged, that is, it keeps waiting. Technically, the syntax of binders is extended to include substitutions:

$$b ::= \dots \mid [\text{some}(c)/x]$$

This behaviour is seamlessly embedded into quality binders, where a single output can affect a number of sub-binders, due to the broadcast paradigm and to the intended semantics of the quality binder, according to which the sub-binders are *simultaneously* active.

As for evaluating whether or not enough inputs have been received, the relation  $b::_r\theta$  defines a Boolean interpretation of binders. An input evaluates to  $\theta = [\text{none}/x]$  with  $r = \text{ff}$ , for it must be performed before continuing with the computation; a substitution evaluates to itself with  $r = \text{tt}$ , since it stands for a received input. The substitution related to a quality binder is obtained by composing the substitutions given by its sub-binders, while its Boolean interpretation is obtained applying the quality guard, which is a Boolean predicate, to the Boolean values representing the status of the sub-binders. This is denoted  $\llbracket q \rrbracket(r_1, \dots, r_n)$ , and corresponds to the Boolean  $\vee$  for the existential guard and to the Boolean  $\wedge$  for the universal guard.

It is worthwhile observing that substitutions are applied directly by the semantics, and since we consider closed processes, whenever an output  $c_1!t$  is ready to execute,  $t$  must have been replaced by a name  $c_2$  (cf. rule (Brd)). For this very same reason the evaluation of case  $x$  of  $\text{some}(y): P_1 \text{ else } P_2$  is straightforward, as when the clause is reached in an actual execution the variable  $x$  must have been replaced with a constant expression.

It is worthwhile observing that no rule for transition under restriction is provided, therefore rule (Sys) must be applied to pull restrictions to the outer-most level. Here  $(\nu \vec{c})$  denotes the restriction of a list of names.

Rules (Par-) take care of interleaving. In particular, rule (Par-brd) takes care of interleaving broadcast, and applies in two situations.

First, if the parallel component  $P_2$  is a replicated process, the broadcast has precedence over any action of instances of  $P_2$ . This means that when  $P_2$  has the form  $!c_1?x.P$ , the number of input performed by its replicated instances will depend on the unfolding performed previously using the structural congruence as per rule (Sys). Note that this “freedom” issue would not arise if the semantics made use of guarded recursion. Mind to observe, however, that the encoding of replication into recursion would lead to the same problem, since replication is essentially an unguarded recursion.

Second, if the parallel component  $P_2$  is broadcast-prefixed, its broadcast is delayed until the current one has exhausted its synchronisation opportunities. If none of the above applies,

Table 3: A broadcast value-passing semantics with replication.

---


$$\frac{P_1 \equiv (\nu \vec{c}) P_2 \quad P_2 \xrightarrow{\lambda} P_3}{P_1 \Longrightarrow P_3} \quad (\text{Sys})$$


---

${}^l c_1!c_2.P \xrightarrow{c_1!c_2} P$	(Brd)
$\frac{P_1 \xrightarrow{c_1!c_2} P'_1 \quad c_1!c_2 \vdash b \rightarrow b' \quad b' ::_{\text{ff}} \theta}{P_1 \mid {}^l b.P_2 \xrightarrow{c_1!c_2} P'_1 \mid {}^l b'.P_2}$	(In-ff)
$\frac{P_1 \xrightarrow{c_1!c_2} P'_1 \quad c_1!c_2 \vdash b \rightarrow b' \quad b' ::_{\text{tt}} \theta}{P_1 \mid {}^l b.P_2 \xrightarrow{c_1!c_2} P'_1 \mid P_2\theta}$	(In-tt)
${}^l \text{case some}(c) \text{ of some}(y): P_1 \text{ else } P_2 \xrightarrow{\tau} P_1[c/y]$	(Then)
${}^l \text{case none of some}(y): P_1 \text{ else } P_2 \xrightarrow{\tau} P_2$	(Else)
$\frac{P_1 \xrightarrow{\tau} P'_1}{P_1 \mid P_2 \xrightarrow{\tau} P'_1 \mid P_2}$	(Par-tau)
$\frac{P_1 \xrightarrow{c_1!c_2} P'_1}{P_1 \mid P_2 \xrightarrow{c_1!c_2} P'_1 \mid P_2} \quad \text{if } P_2 = !P'_2 \vee P_2 = {}^l c'_1!c'_2 P'_2$	(Par-brd)

---


$$c_1!c_2 \vdash c_1?x \rightarrow [\text{some}(c_2)/x] \quad c_1!c_2 \vdash c_3?x \rightarrow c_3?x \quad \text{if } c_1 \neq c_3$$

$$\frac{c_1!c_2 \vdash b_1 \rightarrow b'_1 \quad \cdots \quad c_1!c_2 \vdash b_n \rightarrow b'_n}{c_1!c_2 \vdash \&_q(b_1, \dots, b_n) \rightarrow \&_q(b'_1, \dots, b'_n)}$$


---


$$t?x ::_{\text{ff}} [\text{none}/x] \quad [\text{some}(c)/x] ::_{\text{tt}} [\text{some}(c)/x]$$

$$\frac{b_1 ::_{r_1} \theta_1 \quad \cdots \quad b_n ::_{r_n} \theta_n}{\&_q(b_1, \dots, b_n) ::_r \theta_n \circ \cdots \circ \theta_1} \quad \text{where } r = \{[q]\}(r_1, \dots, r_n)$$


---

then either a synchronisation rule (In-) or an interleaving with a silent action (Par-tau) must take place.

**2.3. Security model.** On top of the standard operational semantics of the calculus, our process-algebraic specifications rely on a security interpretation of communication actions. As binders are blocking actions, they can be thought of as security checks that require the knowledge of some information to be fulfilled. This knowledge is abstracted here by resorting to the notion of *secure channel* and disregarding the messages actually communicated, in line with the value-passing nature of the calculus. This idea seamlessly applies to simple

inputs and is refined by quality binders: the existential quality guard  $\exists$  describes scenarios in which different ways of fulfilling a check are available, e.g., different ways of proving one’s identity. In contrast to this, the universal quality guard  $\forall$  describes checks that require a number of sub-conditions to be met at the same time, and can be used to refine a security mechanism in terms of sub-checks.

On the other hand, an output represents the satisfaction of the security check specified by the corresponding channel. As the semantics is broadcast, all the input waiting on the given channel will be satisfied, that is, all the pending instances of the security check will be fulfilled and the system will proceed until another blocking check is met. As a consequence, if the adversary can trigger a system component to make an output on a given channel  $c$ , it is as if  $c$  were under the control of the attacker, for all inputs on  $c$  in other components would be satisfied.

Finally, case clauses allow to inspect how a given security check, that is, a preceding binder, has been satisfied, by inspecting which input variables are bound to some values, that is, on which channels the communication took place.

As a result, an attack can be construed as a set of channels, namely, those channels needed to fulfil the security checks on a path to the target. We assume that a system  $P$  is deployed in a hostile environment, simulated by an adversary process  $Q$  running in parallel with  $P$ . The ultimate aim of the protection analysis is to compute what channels  $Q$  has to communicate over in order to drive  $P$  to a given location  $l$ , i.e.,  $P|Q \Longrightarrow^* C[lP']$ , where  $C[lP']$  denotes a sub-process of  $P|Q$  that has reached label  $l$  and  $\Longrightarrow^*$  denotes the reflexive and transitive closure of  $\Longrightarrow$ .

In order to compute the channels that  $Q$  needs to reach a program point  $l$ , this security model is translated into an attacker model where  $Q$  can *guess* every required channel, that is, whenever a security check on the way to  $l$  cannot be avoided,  $Q$  can fulfil it.

### 3. THE DANISH NEMID SYSTEM

Let us introduce now an example that will be developed throughout the paper. NemID (literally: EasyID)<sup>1</sup> is an asymmetric cryptography-based log-in solution used by public institutions as well as private companies in Denmark to provide on-line services, used by virtually every person who resides in the country. Most service providers rely on a browser-based log-in application through which their customers can be authenticated (a Java applet or a JavaScript program – in the following simply “the applet”). For technological and historical reasons, the applet allows proving one’s identity with various sets of credentials. In particular, private citizens can log-in with their social security number, password, and a one-time password, or by exhibiting an X.509-based certificate. Moreover, on mobile platforms that do not support Java, a user is authenticated through a classic id-password scheme.

We focus on a formalisation of the system such that authentication does not take place unless some interactions with the environment take place. In other words, we specify the security checks but not their fulfilment by the user who is supposed to be logging in, so as to rule out the legal way to authenticating and focus on malicious behaviours.

The system is modelled in the Value-Passing Quality Calculus as follows:

$$NemID \triangleq (\nu cert) \dots (\nu access)(!Login \mid !Applet \mid !Mobile)$$

---

<sup>1</sup><https://www.nemid.nu/dk-en/>

$$\begin{aligned}
\mathit{Applet} &\triangleq \\
&^1 \&\exists(\mathit{cert}?x_{\mathit{cert}}, \&\forall(\mathit{id}?x_{\mathit{id}}, \mathit{pwd}?x_{\mathit{pwd}}, \mathit{otp}?x_{\mathit{otp}})). \\
&^2 \mathit{case } x_{\mathit{cert}} \mathit{ of some}(y_{\mathit{cert}}): ^3 \mathit{login!ok} \mathit{ else} \\
&^4 \mathit{case } x_{\mathit{id}} \mathit{ of some}(y_{\mathit{id}}): \\
&^5 \mathit{case } x_{\mathit{pwd}} \mathit{ of some}(y_{\mathit{pwd}}): \\
&^6 \mathit{case } x_{\mathit{otp}} \mathit{ of some}(y_{\mathit{otp}}): ^7 \mathit{login!ok} \mathit{ else } 0 \\
&\mathit{ else } 0 \\
&\mathit{ else } 0 \\
\mathit{Mobile} &\triangleq ^8 \&\forall(\mathit{id}?x'_{\mathit{id}}, \mathit{pin}?x_{\mathit{pin}}). \\
&^9 \mathit{case } x'_{\mathit{id}} \mathit{ of some}(y'_{\mathit{id}}): \\
&^{10} \mathit{case } x_{\mathit{pin}} \mathit{ of some}(y_{\mathit{pin}}): ^{11} \mathit{login!ok} \mathit{ else } 0 \\
&\mathit{ else } 0 \\
\mathit{Login} &\triangleq ^{12} \mathit{login}?x.^{13} \mathit{access!ok}
\end{aligned}$$

The system consists of three processes running in parallel an unbounded number of times. For the sake of brevity, we have omitted to list all the restrictions in front of the parallel components, that involve all the names occurring in the three processes.

Process *Login* is in charge of granting access to the system: whenever a user is authenticated via the applet or a mobile app, an output on channel *login* is triggered that can be received at label 12 leading to the output at label 13, which simulates a successful authentication.

Process *Applet* models the applet-based login solution, where login is granted (simulated by the outputs at labels 3 and 7) whenever the user exhibits a valid certificate or the required triple of credentials. The quality binder at label 1 implements such a security check: in order to pass the binder, either ( $\exists$ ) a certificate has to be provided, simulated by the first sub-binder, or three inputs have to be received ( $\forall$ ), mimicking an id (*id*), a password (*pwd*), and a one-time password (*otp*). Observe how *case* constructs are used to determine what combination allowed passing the binder: at label 2 we check whether the certificate is received, and if this is not the case then we check that the other condition is fulfilled. The main abstraction of our approach takes place at this level, as we can only test whether something is received on a given channel, but we cannot inspect the content of what is received. In other words, the knowledge of channel *cert* mimics the capability of producing a valid certificate, and thus we say that the semantic load of the communication protocol is shifted onto the notion of secure channel. Observe that this perspective seamlessly allows reasoning about the cost of attacking the system: to communicate over *cert*, an adversary has to get hold of a valid certificate, e.g., bribing someone or brute-forcing a cryptographic scheme, and this might prove more expensive than guessing the triple of credentials necessary to achieve authentication along the alternative path, as we shall conclude in § 5.4.

Finally, process *Mobile* describes the intended behaviour of the mobile login solution developed by some authorities (e.g., banks, public electronic mail system), where an id and a password or pin have to be provided upon login.

#### 4. DISCOVERING ATTACKS

The first challenge we tackle is discovering all the attacks leading to a given target. Given a process *P* modelling the system of interest and a label *l* in *P* identifying the target of the attack, the task of the analysis is to find all the sets of channels fulfilling the inputs

occurring in  $P$  on the paths to  $l$ . In other words, we look for the knowledge that allows an adversary  $Q$  to drive  $P$  to reach  $l$ .

To this end,  $P$  is translated into a set of propositional formulae (§ 4.1), termed *flow constraints*, describing how the knowledge of channels relates to reachability of locations, and how given a set of channels some other channels can be derived by the attacker. Then, the constraints are extended to consider the capabilities of the attacker (§ 4.2). Each model of the final set of constraints  $P \stackrel{l}{\Leftrightarrow}$  contains a set of channels that under-approximates the knowledge required to reach  $l$  (§ 4.3).

**4.1. From processes to flow constraints.** The call  $\llbracket P \rrbracket \text{tt}$  of the recursive function  $\llbracket \cdot \rrbracket \cdot$ , defined in Table 4, translates a process  $P$  into a set of constraints of the form  $\varphi \rightsquigarrow \bar{p}$ , where  $\varphi$  is a propositional formula and  $\bar{p}$  a positive literal. The intended semantics of a constraint states that if  $Q$  knows (enough information to satisfy)  $\varphi$ , then  $Q$  knows  $p$ , i.e.,  $\bar{p} = \text{tt}$ . As we shall see below, the antecedent  $\varphi$  accounts for the checks made on the path leading to disclosing  $p$ , namely input binders and **case** clauses. The consequent  $\bar{p}$  can either stand for a channel literal  $\bar{c}$ , meaning that  $Q$  controls  $c$ , an input-variable literal  $\bar{x}$ , meaning that the attacker can satisfy the related input (i.e.,  $x = \text{some}(c)$ ), or a label literal  $\bar{l}$ , meaning that  $l$  is reached.

At each step of the evaluation, the first parameter of  $\llbracket \cdot \rrbracket \cdot$  corresponds to the sub-process of  $P$  that has still to be translated, while the second parameter is a logic formula, intuitively carrying the hypothesis on the knowledge  $Q$  needs to reach the current point in  $P$ . The translation function is structurally defined over processes as explained below.

If  $P$  is  $0$ , then there is no location to be attained and thus no constraint is produced. If  $P = !P'$  or  $P = (\nu c)P'$ , then it spontaneously evolves to  $P'$ , hence  $Q$  does not need any knowledge to reach  $P'$  and gains no knowledge since no communication is performed. A parallel composition is translated taking the union of the sets into which the components are translated.

Communication-related actions have instead an impact on the knowledge of  $Q$ : inputs represent checks that require knowledge, outputs fulfil those checks, and **case** clauses determine the control flow. Assume that an action  $\pi$  is reached in the translation under hypothesis  $\varphi$ , that is,  $\llbracket \pi.P' \rrbracket \varphi$ . Then, a constraint  $\varphi \rightsquigarrow \bar{l}$  is generated: if the attacker fulfils the security checks on a path to  $l$ , then  $l$  is reached. In the logic interpretation of the constraints, this happens when  $\varphi$  evaluates to **tt** under a model given by the knowledge of the attacker, forcing  $\bar{l}$  to be **tt** (as standard implication would do). Moreover, the nature of action  $\pi$  determines whether or not other constraints are produced and how the translation proceeds.

Consider a simple input  ${}^l c?x.P'$ : whenever the action is consumed, it must be that the attacker controls the communication channel  $c$ , hence we translate  $P'$  under the hypothesis  $\varphi \wedge \bar{c}$ . Moreover, if the input is consumed, then  $x$  must be bound to  $\text{some}(c')$ , hence we produce a constraint  $(\varphi \wedge \bar{c}) \rightsquigarrow \bar{x}$ . These two steps respectively accommodate the hypothesis we need for passing a binder (success condition), and the conclusion we can establish whenever a binder is passed (strongest post-condition). In Table 4 functions **hp** and **th** take care of formalising this intuition, that seamlessly applies to quality binders, where the hypothesis is augmented accounting for the combinations of inputs that satisfy the binder, as dictated by the quality guard  $q$ . The last section of Table 4 shows two cases for  $q$ , but any Boolean predicate can be used.

Table 4: The translation  $\llbracket P \rrbracket \text{tt}$  from processes to flow constraints.

$\llbracket 0 \rrbracket \varphi$	$= \emptyset$	
$\llbracket !P \rrbracket \varphi$	$= \llbracket P \rrbracket \varphi$	
$\llbracket P_1   P_2 \rrbracket \varphi$	$= \llbracket P_1 \rrbracket \varphi \cup \llbracket P_2 \rrbracket \varphi$	
$\llbracket (\nu c) P \rrbracket \varphi$	$= \llbracket P \rrbracket \varphi$	
$\llbracket {}^l b.P \rrbracket \varphi$	$= \llbracket P \rrbracket (\varphi \wedge \text{hp}(b)) \cup \text{th}(\varphi, b) \cup \{\varphi \rightsquigarrow \bar{l}\}$	
$\llbracket {}^l c!t.P \rrbracket \varphi$	$= \llbracket P \rrbracket \varphi \cup \{\varphi \rightsquigarrow \bar{c}\} \cup \{\varphi \rightsquigarrow \bar{l}\}$	
$\llbracket \text{case } x \text{ of some}(y) : P_1 \text{ else } P_2 \rrbracket \varphi$	$= \llbracket P_1 \rrbracket (\varphi \wedge \bar{x}) \cup \llbracket P_2 \rrbracket (\varphi \wedge \neg \bar{x}) \cup \{\varphi \rightsquigarrow \bar{l}\}$	
$\text{hp}(c?x) = \bar{c}$		$\text{hp}(\&_q(b_1, \dots, b_n)) = \llbracket q \rrbracket (\text{hp}(b_1), \dots, \text{hp}(b_n))$
$\text{th}(\varphi, c?x) = \{(\varphi \wedge \bar{c}) \rightsquigarrow \bar{x}\}$		$\text{th}(\varphi, \&_q(b_1, \dots, b_n)) = \bigcup_{i=1}^n \text{th}(\varphi, b_i)$
	$\llbracket \forall \rrbracket (\bar{c}_1, \dots, \bar{c}_n) = \bigwedge_{i=1}^n \bar{c}_i$	$\llbracket \exists \rrbracket (\bar{c}_1, \dots, \bar{c}_n) = \bigvee_{i=1}^n \bar{c}_i$

The execution of an output  ${}^l c!t$  satisfies all the security checks represented by inputs waiting on  $c$ . Therefore, if  $Q$  can trigger such output, it obtains the knowledge related to  $c$  without having to know the channel directly, and thus a constraint  $\varphi \rightsquigarrow \bar{c}$  is generated. It is worthwhile observing that this behaviour is justified by the broadcast semantics, and by the fact that the calculus is limited to testing whether or not something has been received over a given channel, shifting the semantic load on the notion of secure channel. Moreover, note that the asymmetry between input and output is due to the fact that outputs are non-blocking.

A case construct is translated by taking the union of the constraints into which the two branches are translated: as the check is governed by the content of the case variable  $x$ , we record that the then branch is followed only when  $x$  is bound to  $\text{some}(c)$  by adding a literal  $\bar{x}$  to the hypothesis, as we do for inputs, and we add  $\neg \bar{x}$  if the else branch is followed.

The set of constraints  $\llbracket P \rrbracket \text{tt}$  computed according to Table 4 can be normalised so as to produce a compact representation of  $P$ . Whenever two rules  $\varphi \rightsquigarrow \bar{p}$  and  $\varphi' \rightsquigarrow \bar{p}$  are in  $\llbracket P \rrbracket \text{tt}$ , they are replaced with a single rule  $(\varphi \vee \varphi') \rightsquigarrow \bar{p}$ . This simplification is intuitively sound for if  $\varphi$  leads to obtain  $p$  and  $\varphi'$  leads to obtain  $p$ , then  $p$  is available to the attacker under the condition that  $\varphi \vee \varphi'$  is known. In the following, we assume to deal with sets of constraints in such format.

**4.2. Modelling the attacker.** A rule  $\varphi \rightsquigarrow \bar{p}$  in  $\llbracket P \rrbracket \text{tt}$  describes how  $Q$  can attain  $p$  playing according to the rules of the system, namely fulfilling the checks described by  $\varphi$ . Nonetheless, when  $p$  is a channel, an attacker can always try to obtain it directly, for instance guessing some cryptographic keys or bursting a gate. In order to account for this possibility, we enrich each rule  $\varphi \rightsquigarrow \bar{c}$  by replacing the antecedent with the disjunction  $g_c \vee \varphi$ , where literal  $g_c$  (for “guess  $c$ ”) represents the possibility of learning  $c$  directly. For each channel  $c$  such that no rule  $\varphi \rightsquigarrow \bar{c}$  is in  $\llbracket P \rrbracket \text{tt}$ , we add to  $\llbracket P \rrbracket \text{tt}$  a constraint  $g_c \rightsquigarrow \bar{c}$ , expressing that  $Q$  has no option but guessing the channel.

Finally, observe that having added the literals  $g_c$ , which tell how the attacker can get hold of a channel in any other way than those legal in  $P$ , interpreting the relation  $\sim$  as the propositional bi-implication (equivalence)  $\Leftrightarrow$  preserves the minimal models of the system of constraints. A constraint  $\varphi \vee g_c \Leftrightarrow \bar{c}$  states that  $c$  is only obtained by guessing or by making  $P$  disclose it and, on the other hand, that if  $c$  is known to the attacker it must be because they have guessed it or because they made  $P$  disclose it.

In the following, given a label  $l$  of interest, we shall write  $P_{\Leftrightarrow}^l$  to denote the conjunction of constraints  $\llbracket P \rrbracket \text{tt}$  which have undergone the transformations mentioned above. In particular, in  $P_{\Leftrightarrow}^l$

- $\bar{l}$  is a fact, expressing the query we want to study;
- there is exactly one conjunct  $\varphi \Leftrightarrow \bar{l}$ ;
- for each channel  $c$  occurring in  $P$ , there is exactly one conjunct  $\varphi \Leftrightarrow \bar{c}$ , and  $\varphi$  has the form  $g_c \vee \varphi'$  where  $g_c$  does not occur elsewhere in  $P_{\Leftrightarrow}^l$ .

Intuitively, we assume that  $l$  is reached and we look for the consequences in terms of truth values of channel literals (i.e., we look for *implicants* of  $\bar{l}$  [14]). In the following, we present a SAT-based solution to this problem.

**4.3. A SAT-based solution technique.** A model  $\mu$  of  $P_{\Leftrightarrow}^l$  is a propositional assignment to the literals occurring in  $P_{\Leftrightarrow}^l$  such that the formula evaluates to **tt** (true). Such an assignment can be represented as a function mapping the literals in  $P_{\Leftrightarrow}^l$ , denoted  $\text{dom}(\mu)$ , to truth values  $\{\text{ff}, \text{tt}\}$ . Now, denoted by  $Names$  the set of names (channels) occurring in the process  $P$  under study, then the set

$$\text{attack}(\mu) = \{\bar{c} \in \text{dom}(\mu) \mid c \in Names \wedge \mu(g_c) = \text{tt}\}$$

identifies a set of channels that, if guessed, satisfies the constraints in  $P_{\Leftrightarrow}^l$ . In the semantics,  $\text{attack}(\mu)$  under-approximates a set of channels that fulfil the security checks on a path to  $l$  in  $P$ , i.e., a way for  $Q$  to drive  $P$  to  $l$ . Denoted by  $M^l$  the set of all models of  $P_{\Leftrightarrow}^l$ , the corresponding set of sets of channels  $\text{attack}(M^l)$ , obtained by point-wise application of  $\text{attack}$  to the elements of  $M^l$ , contains under-approximations to all the attacks leading to  $l$ .

Hence, in order to solve the analysis we need essentially to compute all the models  $M^l$  of the propositional constraints  $P_{\Leftrightarrow}^l$ , that is, we have to solve the ALL-SAT problem for the input formula  $P_{\Leftrightarrow}^l$ . If no solution is found, i.e.,  $P_{\Leftrightarrow}^l$  is unsatisfiable, then the program point indicated by  $l$  is not reachable.

It is worthwhile observing that the translation into flow constraints over-approximates the behaviour of a process, giving rise to more executions than those actually arising in the semantics. Such spurious executions correspond to attacks that under-approximate the sets of channels required to reach the target  $l$ , and therefore the overall analysis results in an under-approximation. This intuition is formalised in the following correctness statement:

$$\text{if } P|Q \Longrightarrow^* C[lP'] \text{ then } \exists \mathcal{N} \in \text{attack}(M^l) \text{ s.t. } \mathcal{N} \subseteq \text{fc}(Q)$$

(where we slightly abuse the syntax writing  $C[lP']$  for the sake of conciseness), i.e., for all the executions in which  $Q$  drives  $P$  to  $l$ , the analysis computes a set of channels  $\mathcal{N} \in Names$  that under-approximates the knowledge required of  $Q$ . The formulation of the actual theorem requires to establish some additional notation, and therefore is deferred to Appendix A. In the following, we shall focus instead on an example process that pinpoints the imprecision of the analysis.

A main source of over-approximation in the translation to flow constraints is the treatment of replication and restriction, whose interplay is simply disregarded by the analysis. As a matter of fact, a name restricted under replication is a different name in all the instances of the replicated process. Consider the following example:

$$P \triangleq (\nu c) \left( (!(\nu a)^1 a ? x_a . ^2 c ! c) | ^3 c ? x_c . ^4 c ? x'_c . ^5 \dots \right)$$

Let label 5 be the location of interest and consider the following translation of  $P$  into flow constraints, conveniently simplified for the sake of conciseness:

$$(g_a \Leftrightarrow \bar{a}) \wedge (g_c \vee \bar{a} \Leftrightarrow \bar{c}) \wedge (\bar{c} \Leftrightarrow \bar{5}) \wedge (\mathbf{tt} \Leftrightarrow \bar{5})$$

According to the translation, an attacker can reach label 5 either by knowing  $a$  or by knowing  $c$ , that is,  $\text{attack}(M^5) = \{\{a\}, \{c\}, \{a, c\}\}$ .

Whilst it is clear that if the attacker guesses  $c$  then the security checks at labels 3, 4 can be satisfied, it is less obvious what it means for the attacker to guess  $a$ . In fact, if the adversary makes an output on  $a$  twice, then two outputs on  $c$  are triggered, and thus the checks on the path to the goal 5 are fulfilled. According to the semantics this may happen in all traces in which the replication is unfolded at least twice. Nonetheless, the two instances of  $a$  in the two copies of the process are different:  $\alpha$ -renaming applies producing names  $a$  and  $a'$ . Hence, by claiming that 5 can be reached by guessing  $a$ , the analysis under-approximates the knowledge required of the attacker, who needs  $a$  and  $a'$ .

**4.4. Translating NemID.** The translation of the NemID system of § 3 returns the following flow constraints, augmented as explained above. For the sake of simplicity we omit the occurrences of  $\mathbf{tt}$  as a conjunct in all left-hand sides.

$$\left. \begin{array}{l}
 \bar{1} \\
 \overline{\text{cert}} \Leftrightarrow \bar{x}_{\text{cert}} \\
 \overline{\text{id}} \Leftrightarrow \bar{x}_{\text{id}} \\
 \overline{\text{pwd}} \Leftrightarrow \bar{x}_{\text{pwd}} \\
 \overline{\text{otp}} \Leftrightarrow \bar{x}_{\text{otp}} \\
 \underbrace{\overline{\text{cert}} \vee (\overline{\text{id}} \wedge \overline{\text{pwd}} \wedge \overline{\text{otp}})}_{\varphi} \Leftrightarrow \bar{2} \\
 \varphi \wedge \bar{x}_{\text{cert}} \Leftrightarrow \bar{3} \\
 \varphi \wedge (\neg \bar{x}_{\text{cert}}) \Leftrightarrow \bar{4} \\
 \varphi \wedge (\neg \bar{x}_{\text{cert}}) \wedge \bar{x}_{\text{id}} \Leftrightarrow \bar{5} \\
 \varphi \wedge (\neg \bar{x}_{\text{cert}}) \wedge \bar{x}_{\text{id}} \wedge \bar{x}_{\text{pwd}} \wedge \bar{x}_{\text{otp}} \Leftrightarrow \bar{7} \\
 g_{\text{id}} \Leftrightarrow \overline{\text{id}} \\
 g_{\text{pwd}} \Leftrightarrow \overline{\text{pwd}} \\
 g_{\text{otp}} \Leftrightarrow \overline{\text{otp}} \\
 g_{\text{cert}} \Leftrightarrow \overline{\text{cert}}
 \end{array} \right\} \begin{array}{l} \text{from} \\ \llbracket \text{Applet} \rrbracket \mathbf{tt} \end{array}$$
  

$$\left. \begin{array}{l}
 \bar{8} \\
 \overline{\text{id}} \Leftrightarrow \bar{x}'_{\text{id}} \\
 \overline{\text{pin}} \Leftrightarrow \bar{x}_{\text{pin}} \\
 \overline{\text{id}} \wedge \overline{\text{pin}} \Leftrightarrow \bar{9} \\
 \overline{\text{id}} \wedge \overline{\text{pin}} \wedge \bar{x}'_{\text{id}} \Leftrightarrow \bar{10} \\
 \overline{\text{id}} \wedge \overline{\text{pin}} \wedge \bar{x}'_{\text{id}} \wedge \bar{x}_{\text{pin}} \Leftrightarrow \bar{11} \\
 g_{\text{pin}} \Leftrightarrow \overline{\text{pin}}
 \end{array} \right\} \begin{array}{l} \text{from} \\ \llbracket \text{Mobile} \rrbracket \mathbf{tt} \end{array}$$
  

$$\left. \begin{array}{l}
 \bar{12} \\
 \overline{\text{login}} \Leftrightarrow \bar{x} \\
 \overline{\text{login}} \Leftrightarrow \bar{13} \\
 g_{\text{access}} \vee \overline{\text{login}} \Leftrightarrow \overline{\text{access}}
 \end{array} \right\} \begin{array}{l} \text{from} \\ \llbracket \text{Login} \rrbracket \mathbf{tt} \end{array}$$



$$\left( g_{\text{login}} \vee \underbrace{(\varphi \wedge \bar{x}_{\text{cert}})}_{\llbracket \text{Applet} \rrbracket \text{tt}} \vee \underbrace{((\varphi \wedge (\neg \bar{x}_{\text{cert}}) \wedge \bar{x}_{\text{id}} \wedge \bar{x}_{\text{pwd}} \wedge \bar{x}_{\text{otp}}))}_{\llbracket \text{Applet} \rrbracket \text{tt}} \vee \underbrace{((\bar{\text{id}} \wedge \bar{\text{pin}} \wedge \bar{x}'_{\text{id}} \wedge \bar{x}_{\text{pin}}))}_{\llbracket \text{Mobile} \rrbracket \text{tt}} \right) \Leftrightarrow \overline{\text{login}}$$

where the last formula combines the constraints that show the various ways to trigger an output on channel `login`. Notice how each formula models the checks on a given path: for being granted access, i.e., reaching label 13, a communicating process has to know channel `login`, as specified by the constraints derived from  $\llbracket \text{Login} \rrbracket \text{tt}$ . In turn, the last formula describes what is needed in order to get hold of `login`, giving rise to a backward search procedure formalised in § 6.

As for the solutions to the set of constraints, restricting to attacks we have:

$$\begin{aligned} \text{attack}(\mu_{\text{triple}}) &= \{\bar{\text{id}}, \bar{\text{pwd}}, \bar{\text{otp}}\} \\ \text{attack}(\mu_{\text{cert}}) &= \{\bar{\text{cert}}\} \\ \text{attack}(\mu_{\text{mobile}}) &= \{\bar{\text{id}}, \bar{\text{pin}}\} \\ \text{attack}(\mu_{\text{login}}) &= \{\overline{\text{login}}\} \end{aligned}$$

where the first two attacks come from the two ways of authenticating in the applet, the third from the mobile app, and the last one from guessing `login`, i.e., obtaining access in a way not encompassed by the system.

**4.5. Modularity and refinement.** Let us discuss now the role the analysis may take in a broader context, where it can be used to progressively refine the modelling of sub-systems which are revealed as potential attack targets, leveraging the modularity of process-algebraic specifications.

Consider the NemID system and assume that a new way to access its services were offered by a service provider, which would authenticate a user via a phone number:

$$\begin{aligned} \text{Phone} &\triangleq \\ &{}^{14}\text{phone?}x_{\text{ph}}.{}^{15}\text{case } x_{\text{ph}} \text{ of some}(y_{\text{ph}}): {}^{16}\text{login!ok else } 0 \\ \text{NemID}' &\triangleq (\nu \text{login}) \dots (\nu \text{phone}) \\ &(!\text{Login} \mid !\text{Applet} \mid !\text{Mobile} \mid !\text{Phone}) \end{aligned}$$

Then we have  $\llbracket \text{NemID}' \rrbracket \text{tt} = \llbracket \text{NemID} \rrbracket \text{tt} \cup \llbracket \text{Phone} \rrbracket \text{tt}$ , that is, the translation of a new top-parallel process is independent from the formulae that have already been generated (but the quest for models has to be re-computed with all the clauses). Obviously, due care has to be paid to names, e.g., the name `login` in *Phone* has to be the same as the one used in process *NemID*. However, as we have seen, while restrictions play a crucial role in the semantics, they are simply ignored by the translation.

Besides being compositional with respect to the analysis of new components, the translation suitably integrates in a refinement cycle, where we start from a coarse abstraction of the system and then progressively refine those components that are revealed as candidates for being attacked, by replacing the corresponding set of formulae with a finer one. The constraint on names translates to a constraint on the interface of the component: if process *A* is replaced by process *B*, then *B* must be activated by the same inputs that activate *A*,

and vice-versa, it must produce the same outputs towards the external environment that  $A$  produces.

## 5. QUANTIFYING ATTACKS

The second challenge we tackle is quantifying attacks with respect to a general notion of cost. In the previous section we have presented a SAT-based solution technique, where each model of  $P_{\Leftrightarrow}^l$  contains a set of channels that are necessary to fulfil an attack. Not all security mechanisms, however, offer the same protection guarantees, that is, not all channels are equal. A retinal scan can prove more difficult to bypass than a pin lock, and thus offer more protection. This is not the case, however, of an insider who is authorised to enter the corresponding room. A cost structure over channels facilitates formalising these considerations, and assigning costs to channels leads naturally to quantify sets of channels, that is, attacks.

The characterisation of attacks in terms of cost allows to order them, and ultimately to focus on those which are deemed the most likely given our understanding of the candidate attacker profiles. Moreover, the *minimal* cost of reaching a given location  $l$  identifies the protection deployed to guard  $l$  in the implementation, which can be contrasted with the desiderata of the specification. The higher the cost for the attacker, the higher the protection guarding the target.

In the following, we extend the qualitative analysis of § 4 to a quantitative setting. In particular, we should follow a modular approach according to which cost considerations are developed on top of the structure of the original analysis. The benefit of a layered strategy is two-fold: on the one hand, we present a technique that can be exploited to transform a great many qualitative analyses into quantitative analyses; on the other hand, whenever quantitative information about the entities in question is not available, we can resort to the qualitative solution.

**5.1. Security labels.** We have already seen how the calculus is instrumented with labels so as to refer to locations of interest. Such labels can be levered to build a *security lattice* ( $\Sigma = \{\sigma_1, \dots, \sigma_n\}, \sqsubseteq_\Sigma$ ) specifying the desired protection deployed to guard each location. The set  $\Sigma$  is equipped with the greatest lower bound operator  $\sqcap_\Sigma$ , and we assume to have a function  $\text{security} : \mathcal{L} \rightarrow \Sigma$  that maps labels into security levels. In particular,  $\text{security}(l_1) \sqsubseteq_\Sigma \text{security}(l_2)$  denotes that the need for protection of the program point indicated by  $l_2$  is greater than or equal to the need for protection of the program point indicated by  $l_1$ .

As an example of a security lattice, consider the *military lattice* given by

$$\Sigma = \{\text{unclassified}, \text{confidential}, \text{secret}, \text{top-secret}\}$$

with the ordering  $\text{unclassified} \sqsubseteq_\Sigma \text{confidential} \sqsubseteq_\Sigma \text{secret} \sqsubseteq_\Sigma \text{top-secret}$ . More complex lattices, in particular non-linearly-ordered ones, are discussed in [3, Ch. 7].

In the NemID example of § 3, we have observed that the system authenticates a user at label 13. In terms of security levels, we can rely on a simple security lattice  $\text{unrestricted} \sqsubseteq_\Sigma \text{restricted}$ , where  $\text{security}(13) = \text{restricted}$ . Similarly, we would like to have the highest security level for labels 2 to 7 in process *Applet* and 9 to 11 in process *Mobile*. In fact, such labels are reached after fulfilling the security checks at labels 1 and 8, respectively, which have instead level *unrestricted*, as label 12 in process *Login*.

**5.2. From qualitative to quantitative considerations.** Let  $\text{cost}$  be a function from channels  $c \in \text{Names}$  to costs  $k \in \mathcal{K}$ . Formally, we require  $(\mathcal{K}, \oplus)$  to be a commutative monoid (also known as Abelian monoid), that is,  $\oplus$  is an associative and commutative binary operation on the set  $\mathcal{K}$  and has an identity element. Moreover, we require  $\mathcal{K}$  to be equipped with a partial order  $\sqsubseteq_{\mathcal{K}}$ , such that  $(\mathcal{K}, \sqsubseteq_{\mathcal{K}})$  is a lattice, and  $\oplus$  to be extensive, that is, the sum of two elements always dominates both the summands:

$$\forall k_1, k_2 \in \mathcal{K} . k_1 \sqsubseteq_{\mathcal{K}} (k_1 \oplus k_2) \wedge k_2 \sqsubseteq_{\mathcal{K}} (k_1 \oplus k_2) \quad (5.1)$$

Finally, we assume that  $\oplus$  is monotone and the least element  $\perp \in \mathcal{K}$  is its identity element, that is,  $\oplus$  is an upper bound operator of the lattice  $(\mathcal{K}, \sqsubseteq_{\mathcal{K}})$ , and therefore satisfies condition (5.1). For the sake of simplicity, we assume that the costs of channels are independent.

For the sake of simplifying the notation, in the following we shall feel free to apply the function  $\text{cost}$  to sets of names, according to the following definition:

$$\begin{aligned} \text{cost} &: \mathcal{P}(\text{Names}) \rightarrow \mathcal{K} \\ \text{cost}(\{c_1, \dots, c_n\}) &= \bigoplus_{i=1}^n \text{cost}(c_i) \end{aligned}$$

Likewise, we extend the function  $\text{cost}$  also to sets of sets of names by point-wise application:

$$\begin{aligned} \text{cost} &: \mathcal{P}(\mathcal{P}(\text{Names})) \rightarrow \mathcal{P}(\mathcal{K}) \\ \text{cost}(\{c_1^1, \dots, c_{n_1}^1\}, \dots, \{c_1^m, \dots, c_{n_m}^m\}) &= \\ &\{\text{cost}(\{c_1^1, \dots, c_{n_1}^1\}), \dots, \text{cost}(\{c_1^m, \dots, c_{n_m}^m\})\} \end{aligned}$$

Therefore, given a model  $\mu$  of  $P_{\Leftrightarrow}^l$ , the corresponding set  $\text{attack}(\mu) = \{c_1, \dots, c_n\}$  can be quantified as

$$\text{cost}(\text{attack}(\mu)) = \bigoplus_{c \in \text{attack}(\mu)} \text{cost}(c)$$

It is worthwhile noticing that this Boolean approach implies that multiple occurrences of the same basic action refer in fact to the same instance of the action, that is, to the same security mechanism, which once bypassed is bypassed “forever”. Multiple instances of the same security mechanism must be represented with different names. This observation is intimately related to the over-approximating nature of the qualitative analysis discussed in § 4.3, where we have already mentioned that restrictions and replications are simply ignored.

As we mentioned above, however, a conservative approach to security would consider the attacks of *minimal* cost. Hence, given two attacks, i.e., two distinct models  $\mu, \mu'$ , we would discard  $\mu'$  in case  $\text{cost}(\text{attack}(\mu)) \sqsubset_{\mathcal{K}} \text{cost}(\text{attack}(\mu'))$ . We can thus restrict the set of models  $M^l$  to the ones bearing attacks of minimal cost:

$$\text{minimal}(M^l) = \left\{ \mu \in M^l \mid \forall \mu' \in M^l . \text{cost}(\text{attack}(\mu')) \not\sqsubset_{\mathcal{K}} \text{cost}(\text{attack}(\mu)) \right\}$$

It is worthwhile noticing that  $\text{minimal}(M^l)$  may contain more than one model, as (i) we consider all the attacks with same cost and (ii) some attacks may have incomparable costs in case the cost set is not linearly ordered.

Now, we can relate an attack to the corresponding security level  $\sigma \in \Sigma$  required to counter it by means of a function  $\text{level}: \mathcal{K} \rightarrow \Sigma$ , compressing cost regions into security levels:

$$\text{level}(k) = \begin{cases} \sigma_1 & \text{if } k \in \{k_1^1, \dots, k_{h_1}^1\} \\ \vdots & \\ \sigma_m & \text{if } k \in \{k_1^m, \dots, k_{h_m}^m\} \end{cases}$$

where `level` is a well-defined function if the sets of costs  $\{k_1^i, \dots, k_{h_i}^i\}$  are pair-wise disjoint and their union is  $\mathcal{K}$ . Moreover, it is natural to require that `level` is monotone. A simple example in the cost set  $(\mathbb{N}, +)$  and security lattice `low`  $\sqsubset$  `medium`  $\sqsubset$  `high` is given by the choice

$$\text{level}(k) = \begin{cases} \text{low} & \text{if } k \leq 1024 \\ \text{medium} & \text{if } 1024 < k \leq 2048 \\ \text{high} & \text{if } 2048 < k \end{cases}$$

where numbers could represent the length of cryptographic keys, and we state for instance that a program point is poorly protected if no more than 1024 bits are necessary to attain it (for a fixed cryptosystem).

Finally, we extend `level` to work on sets of costs so as to encompass all the sets of channels produced by the analysis at once:

$$\begin{aligned} \text{level} : \mathcal{P}(\mathcal{K}) &\rightarrow \mathcal{P}(\Sigma) \\ \text{level}(\{k_1, \dots, k_n\}) &= \{\text{level}(k_1), \dots, \text{level}(k_n)\} \end{aligned}$$

where the input  $\{k_1, \dots, k_n\}$  is the set of costs of all minimal attacks, computed as

$$\text{cost}(\text{attack}(\text{minimal}(M^l)))$$

Finally, the greatest lower bound  $\sqcap_\Sigma$  is used to derive the greatest security level compatible with all attacks in `minimal`( $M^l$ ), that is, the protection of a program point corresponds at most to the cost of the weakest path leading to it.

A graphical illustration of the various components of the analysis is displayed in Fig. 1, where it is apparent how the quantitative analysis is built on top of the qualitative analysis. Intuitively, function `security` is the *specification* expressing the target security architecture of a system with respect to a given security lattice, while `level(cost(attack(minimal( $M^l$ ))))` captures (an under-approximation of) how this architecture has been realised in the *implementation*.

The overall aim of the analysis, i.e., checking whether the deployed protection lives up to the required security level, can thus be expressed by the property

$$\forall l \in \mathcal{L} . \text{security}(l) \sqsubseteq_\Sigma \text{level}(\text{cost}(\text{attack}(\text{minimal}(M^l))))$$

A violation of this condition is referred to as a potential *inversion of protection*. The overall under-approximation of the analysis is the result of minimising over the costs of under-approximating sets of channels (cf. § 4.3).

**5.3. Optimisation Modulo Theories.** In order to compute the set of sets of channels `attack`(`minimal`( $M^l$ )) that allow reaching  $l$  incurring minimal costs, we need to solve an optimisation problem subject to the Boolean constraints  $P_{\Leftrightarrow}^l$ . There exist various techniques to cope with such problems, each suitable for particular choices of cost sets and objective functions. One solution is to first compute  $M^l$  and then minimise it by comparing models as explained above. Nonetheless, cost information can be levered to skip non-optimal models during the search, hence improving the performance. In the following, we show how to exploit an SMT solver to tackle the problem in its most general form. We limit to mention that linear programming techniques such as Pseudo-Boolean optimisation [10] are efficient alternatives for dealing with the monoid  $(\mathbb{Z}, +)$  and linear objective functions.

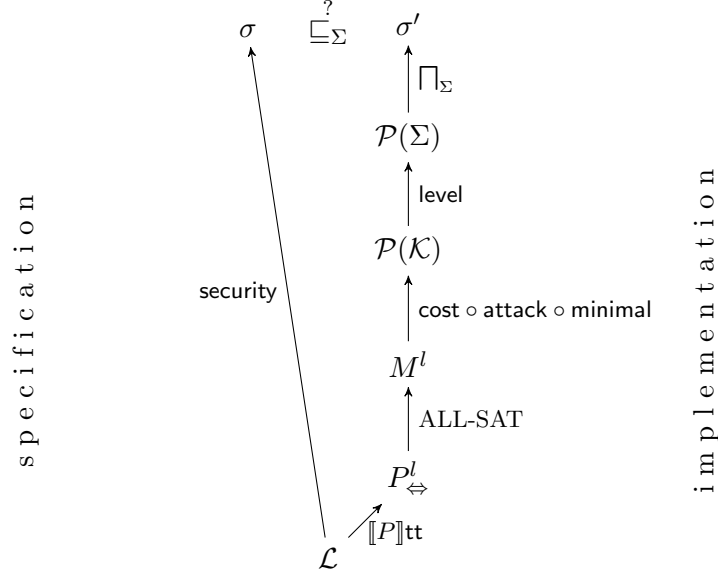


Figure 1: The quantitative protection analysis at a glance for a fixed process  $P$ .

In a nutshell, our task reduces to compute models of  $P^l_{\Leftrightarrow}$  containing attacks of minimal cost in the lattice  $\mathcal{K}$ . In other words, we are looking for *prime* implicants of  $\bar{l}$  [14], where primality is sought with respect to the given cost set.

Such an optimisation problem can be tackled by computing models for a list  $\Pi_1, \dots, \Pi_n$  of SMT problems, where  $\Pi_i$  is a more constrained version of  $\Pi_{i-1}$  that requires to improve on the cost of the current solution. The initial problem  $\Pi_1$  consists of the propositional constraints  $P^l_{\Leftrightarrow}$  and of the objective function, whose value on the current model is stored in variable  $\text{goal}$ .

The objective function is essentially the cost of the current model. In order to compute  $\text{cost}(\text{attack}(\mu))$  into variable  $\text{goal}$  as part of  $\mu$  itself we define:

$$\text{goal} := \bigoplus_{i=1}^n (\text{if } g_{c_i} \text{ then } \text{cost}(c_i) \text{ else } \perp)$$

where we combine the costs of all the channels that must be guessed, that is, the channels  $c_i$ 's such that the corresponding guessing literal  $g_{c_i}$  is found to be  $\text{tt}$ . Otherwise, if a  $g_{c_i}$  is  $\text{ff}$ , then the corresponding  $c_i$  needs not be guessed and its cost does not contribute to the cost of the attack. Recall that the least element  $\perp$  of the cost lattice  $\mathcal{K}$  does not contribute any cost, for it is the neutral element with respect to the cost combinator  $\oplus$ . Hence, by construction we have  $\mu(\text{goal}) = \text{cost}(\text{attack}(\mu))$ .

Then, while the problem is satisfiable, we improve on the cost of the current model by asserting new constraints which tighten the value of  $\text{goal}$ , until unsatisfiability is reported. Algorithm 1 displays the pseudo-code of the procedure. In particular, observe that when a new problem  $\Pi_i$  is generated, additional constraints are asserted that ask for (i) a different model and (ii) a non-greater cost: the former condition speeds up the search, while the latter explores the cost frontier.

**Data:** The problem  $\Pi \triangleq P_{\Leftrightarrow}^l \wedge (\text{goal} := \bigoplus_{i=1}^n (\text{if } g_{c_i} \text{ then } \text{cost}(c_i) \text{ else } \perp))$   
**Result:** the set  $\mathcal{M}$  of pairs  $(\mu, \text{cost}(\text{attack}(\mu)))$  such that  $\mu \in \text{minimal}(M^l)$   
 $\mathcal{M} \leftarrow \emptyset;$   
**while**  $\Pi$  *satisfiable* **do**  
   $\mu \leftarrow \text{get-model}(\Pi);$   
   $k \leftarrow \mu(\text{goal});$   
  **forall the**  $(\mu', k') \in \mathcal{M} \mid k \sqsubset_{\mathcal{K}} k'$  **do** //  $\mu$  outperforms  $\mu'$   
  |  $\mathcal{M} \leftarrow \mathcal{M} \setminus \{(\mu', k')\}$   
  **end**  
   $\mathcal{M} \leftarrow \mathcal{M} \cup \{(\mu, k)\};$   
   $\Pi \leftarrow \Pi \wedge \neg(\bigwedge_{i=1}^n (\bar{c}_i = \mu(\bar{c}_i))) \wedge \neg(\text{goal} \sqsupset_{\mathcal{K}} k);$   
**end**

**Algorithm 1:** The SMT-based solution procedure.

The termination of the algorithm is ensured by the finiteness of possible models to the propositional variables of the  $\Pi_i$ 's, and by the fact that the same model cannot occur twice as solution due to the new constraints we generate in each iteration. At most, we need to solve as many  $\Pi_i$ 's as there are models of  $P_{\Leftrightarrow}^l$ , which coincide with the qualitative analysis (ALL-SAT). The correctness of the procedure stems from the fact that when unsatisfiability is claimed, by construction of the  $\Pi_i$ 's there cannot exist further models that comply with the cost constraints.

It is worthwhile noticing how resorting to propositional logic integrates with the overall under-approximating nature of the analysis: a channel can either be learnt or not, and its cost contribute or not to the cost of an attack. This means that we do not keep track of the number of attempts made to guess some information, and always assume that guessing  $c$  is successful whenever  $g_c$  is found to be true. In other words, for a given cost set, we are considering the luckiest or cleverest attacker. As for the cost set (comparing and combining costs), SMT solvers offer native support for numeric costs and common mathematical functions, while more complex cost sets have to be encoded manually.

Finally, observe that the procedure above is not dependent on our analysis, but can be generally exploited to find optimal models of arbitrary logic formulae in arbitrary cost sets, and can be extended seamlessly to more complex logics.

**5.4. Attacking NemID.** Consider the NemID system discussed in § 3. There are several techniques for quantifying the cost of guessing secret information. *Quantification of information leakage* [6] is an information theory-based approach for estimating the information an adversary gains about a given secret  $s$  by observing the behaviour of a program parametrised on  $s$ . If  $s$  is quantified in bits, then the corresponding information leaked by the program is quantified as the number of bits learnt by the adversary by observing one execution of the system. For instance, consider a test program  $T$  parametrised on a secret password.  $T$  inputs a string and answers whether or not the password is matched. Under the assumptions that the adversary knows the program and the length of the secret (no security-by-obscurity), we can estimate the knowledge gained by the adversary after one guessing attempt.

We leverage QUAIL [7], a freely-available tool for quantifying information leakage, for determining costs to channels. Denoted  $\text{leak}(T, s)$  the leakage of  $T$  on a secret  $s$  as computed

by QUAAIL, we quantify the strength of a channel  $c$  of  $n$  bits as

$$\text{cost}(c) = \frac{n}{\text{leak}(T, c)}$$

where we assume the security offered by  $c$  to be uniformly distributed over the  $n$  bits. In this settings we are thus working in the cost monoid  $(\mathbb{Q}, +)$ .

In our running example, the secrets to be guessed are `pwd`, `otp`, `cert`, `pin`, while we assume that `id` is known to the attacker and thus has cost 0 (in particular, in the `NemID` system is not difficult to retrieve such `id`, corresponding to the social security number of an individual). Moreover, we know that `pwd` contains between 6 and 40 alphanumeric symbols and it is not case sensitive: assuming an average length of 10 symbols, given that there are 36 such symbols, we need 5.17 bits to represent each symbol, for a total length of 52 bits. Analogously, we determine the length of `otp` as 20 bits, while the length of the `pin` depends on the service provider: in case of a major bank it is just 14 bits. As for the certificate, the authority is following NIST recommendations, using 2048-bit RSA keys for the time being, and for the sake of simplicity we assume that guessing an RSA key cannot be faster than guessing each of the bits individually. Finally, we disregard `login`, `access` by assigning them the least upper bound of the costs of all the other channels. Names used only as messages can be disregarded by assigning them cost 0, as they do not influence an attack. Exploiting QUAAIL and the formula defined above, we obtain the following cost map:

$$\begin{aligned} \text{cost}(\text{pwd}) &= 4.4 \times 10^{15} & \text{cost}(\text{pin}) &= 1.5 \times 10^4 \\ \text{cost}(\text{otp}) &= 10^6 & \text{cost}(\text{cert}) &= 3.4 \times 10^{616} \end{aligned}$$

Thus, limiting our attention to the set *Names* of channels occurring in the process *NemID*, the problem is to minimise

$$\sum_{c \in \text{Names}} (\text{if } g_c \text{ then } \text{cost}(c) \text{ else } 0)$$

under the constraints given by  $P_{\Leftrightarrow}^{13}$ . Instructed with this input, our procedure finds that the formula is satisfiable and the single cheapest model  $\mu$  contains

$$g_{\text{id}} \mapsto \text{tt} \quad g_{\text{pwd}} \mapsto \text{ff} \quad g_{\text{mail}} \mapsto \text{ff} \quad g_{\text{pin}} \mapsto \text{tt} \quad g_{\text{login}} \mapsto \text{ff}$$

entailing  $\text{attack}(\text{minimal}(M^{13})) = \{\{\text{id}, \text{pin}\}\}$  and  $\text{cost}(\text{attack}(\mu)) = \text{cost}(\text{id}) + \text{cost}(\text{pin}) = 1.5 \times 10^4$ .

As for the desired security levels, we observed in § 3 that  $\text{security}_{(13)} = \text{restricted}$  should hold. It is desirable to take as touchstone the protection offered by the applet, for it is standard among all service providers. The protection offered by the applet is the minimum between the cost of guessing a certificate and the cost of guessing the triple of credential, that is,  $4.4 \times 10^{15} + 10^6$ , therefore we should set

$$\text{level}(k) = \begin{cases} \text{restricted} & \text{if } k \geq 4.4 \times 10^{15} + 10^6 \\ \text{unrestricted} & \text{otherwise} \end{cases}$$

We would like to verify that  $\text{restricted} \sqsubseteq_{\Sigma} \text{level}(\text{cost}(\text{attack}(M^{(13)})))$ , which is false. Hence, it is the case that the implementation potentially guarantees less protection than the amount required by the specification, and thus we shall issue a warning to the designer of the system.

Finally, it is worthwhile noticing that the framework allows measuring the *distance* between the implementation and the specification, and not only their qualitative compliance.




Figure 2: The Hasse diagram of a partially-ordered cost structure.

The analysis suggests that the most practicable way to break the authentication protocol is attacking the mobile app, as long as we believe that our cost map is sensible. For instance, a cryptographer would deem our assumptions on breaking RSA utterly unrealistic. In the following, we will discuss an effective alternative to the definition of numeric cost sets.

**5.5. Complex cost structures.** So far we have worked with an example in the cost set  $(\mathbb{Q}, +)$ , for it is natively encoded into SMT solvers and matches a first intuition of the notion of cost. Nonetheless, it is often difficult to provide an absolute estimate of the strength of a protection mechanism: sometimes different mechanisms are even incomparable, as cryptography and physical security might be. In such cases, it is more natural to describe the relative strength of a set of mechanisms with respect to each other. This is achieved by computing the analysis over symbolic and partially-ordered cost sets. Observe that Algorithm 1 is already equipped to cope with the general problem of optimising on such cost sets.

As a basic example, consider the cost lattice displayed in Fig. 2: we could characterise the cost of obtaining given information as **cheap**, if it does not require a specific effort, as **cpu**, if it requires significant computational capabilities (e.g., breaking an encryption scheme), as **enrg**, if it requires to spend a considerable amount of energy (e.g., engaging in the wireless exchange of a number of messages), or as **expensive**, if it requires both computations and energy. In order to combine such costs, a suitable choice is to take as monoid operator  $\oplus$  the least upper bound  $\sqcup$  of two elements in the cost lattice.

An interesting case of non-linear cost sets is offered by the study of security in Cyber-Physical Systems, where components combine both software and physical features [41]. In particular, in such systems an attack could require to assemble cyber actions with physical tampering, whose costs can either be comparable or not depending on the nature of the quantities we are interested in (for instance, energy and memory are not directly comparable).

In conclusion, three elements push independently for the comprehensive SMT-based approach: the non-linearity of the cost set, its symbolic nature, and the non-linearity of the objective function.

**5.6. A semantic interpretation of guessing.** It would be possible to formulate a neat semantic interpretation of the guessing capability of the attacker. Assume to deal with processes of the form  $((\nu \vec{c}) P) | Q$ , where the first component is the system under study, in which all restrictions are at the outer-most level, and  $Q$  is the attacker. Now, for  $Q$  to interact with  $P$ , the attacker needs to move inside the scope of some restrictions so as to share some channel names with  $P$ . Whenever  $Q$  enters the scope of a restriction  $(\nu c)$ ,






Figure 3: How to enter a bank vault, for dummies.

the name  $c$  is guessed. In order to account for the cost  $k \in \mathcal{K}$  of guessing a name, we can instrument each restriction with the corresponding cost, writing  $(\nu^k c)$ , and then augment the scope extension rule of Table 2 so as to accumulate the cost of names that are guessed. The standard semantics of restriction used in security applications of process calculi, according to which a new name  $c$  is only known to legal participants unless leaked ( $P$ , in our case), is encompassed by assigning  $c$  an infinite cost.

Though possible, such an extension of the semantics is not necessary to prove the correctness of the analysis. Every assignment that satisfies the propositional constraints leads  $Q$  to reach the location  $l$  of interest, hence also the ones of minimal costs. Nonetheless, it is worthwhile noticing that such a quantitative point of view on restrictions generalises the distinction between the operators `new` and `hide` introduced in the *secret*  $\pi$ -calculus [16]. The operator `hide`  $c$ , which introduces a name  $c$  inhibiting its scope extension, would correspond to  $(\nu^\infty c)$ , while we would have a more fine-grained view on plain scope extension.

## 6. DISPLAYING ATTACKS

We shall now embark in the last challenge of ours, that is, obtaining graphical representations of possible attacks that foster communicating effectively security information to non-experts, as these often are those in charge of taking decisions. We shall do this by means of attack trees, a widely-recognised tool for showing how a goal is attained in terms of combination of sub-goals.

Our developments on attack trees encompass both the qualitative analysis of § 4 and the quantitative extension of § 5. On the one hand, an attack tree contains *all* the attacks for a given target; on the other hand, a tree is construed as a propositional formula whose optimal models can be computed by means of the procedure of § 5.3.

Attack trees are a widely-used graphical formalism for representing threat scenarios, as they appeal both to scientists, for it is possible to assign them a formal semantics, and to practitioners, for they convey their message in a concise and intuitive way (cf. references in § 8.2). In an attack tree, the root represents a target goal, while the leaves contain basic attacks whose further refinement is impossible or can be neglected. Internal nodes show how the sub-trees have to be combined in order to achieve the overall attack, and to this purpose propositional conjunction and disjunction are usually adopted as combinators. On top of this basic model, a number of extensions and applications of attack trees have been proposed, demonstrating how flexible and effective tool they are in practice. Figure 3 displays a simplistic attack tree, where the overall goal of entering a bank vault is obtained by either bribing a guard or by stealing the combination and neutralising the alarm.

**6.1. Synthesising attack trees.** In the following, we shall rely on the translation  $\llbracket P \rrbracket_{tt}$  devised in Table 4 but embrace a slightly different interpretation so as to allow explicitly generating trees. In particular, we replace bi-implications in  $P_{\Leftrightarrow}^l$  with implications and we ignore literals  $g$ , obtaining a set of constraints denoted  $P_{\Rightarrow}^l$ . In order to obtain attack trees as commonly defined in the literature, in the following we assume that all the quality guards  $q$  in the process under study are linear.

Given a process  $P$  and a label  $l$  occurring in  $P$ , we generate a formula  $\llbracket l \rrbracket$  representing the attacks reaching  $l$  by backward chaining the formulae in  $P_{\Rightarrow}^l$  so as to derive  $\bar{l}$ . It is central to observe that the procedure re-establishes the original system of bi-implications thus guaranteeing the correctness of the analysis in terms of compatibility with the developments of § 4.

Before explaining the algorithm, it is worthwhile discussing the nature of the backward chaining-like procedure defined in the following. Standard backward chaining [37, Ch. 7] combines Horn clauses so as to check whether a given goal follows from the knowledge base. Instead, we are in fact trying to derive all the knowledge bases that allow inferring the goal given the inference rules  $P_{\Rightarrow}^l$ , which are not strict Horn clauses as they can contain more than one positive literal. The backward-chaining point of view stresses the relationship of our problem to the quest for implicants of  $\bar{l}$ , as we have already observed.

The rules for generating  $\llbracket l \rrbracket$  are displayed in Table 5. For our formulae are propositional, there is no unification other than syntactical identity of literals involved in the procedure. Notice that the algorithm only applies valid inference rules.

Rule (Sel) selects the antecedent of the formula leading to the goal  $\bar{l}$ : since there is a unique such rule, in order to derive  $\bar{l}$  we have to derive the antecedent  $\varphi$  of  $\varphi \Rightarrow \bar{l}$ . Observe that we are not interested in deriving  $\bar{l}$  in any other way: for  $\bar{l}$  is derived assuming  $\varphi$ , the original bi-implication format  $\varphi \Leftrightarrow \bar{l}$  is re-established.

Rule (Pone-c) encodes either a tautology (if  $\bar{c}$  has to be inferred then  $\bar{c}$  is in the knowledge base) or applications of modus ponens ( $\bar{c}$  is derived assuming  $\varphi$ , thanks to  $\varphi \Rightarrow \bar{c}$ ): the whole rule is an instance of disjunction introduction.

This is the point where our algorithm differs from plain backward chaining: since we are building the knowledge bases that allow inferring  $\bar{l}$ , whenever we encounter a literal  $\bar{c}$  we need to account for all the ways of deriving  $\bar{c}$ , namely by placing  $\bar{c}$  itself in the knowledge base or by satisfying a rule whose consequent is  $\bar{c}$ . Moreover,  $\bar{c}$  plays now the role of  $g_c$ : for the procedure assumes  $c$  without further deriving it, we can safely re-use the literal, whose semantics coincides now with the one of  $g_c$ . Observe that we could consider literals  $g$  and their relation to channels explicitly, but this would impact the size of the tree, hence its readability.

Similarly, rule (Pone-x) encodes an application of modus ponens, taking advantage of the uniqueness of  $\varphi \Rightarrow \bar{x}$  (cf. Lemma B.1).

Rules (Tolle-) collect applications of modus tollens (law of contrapositive) in the classic backward fashion (i.e., when considering the derivation from the leaves to the root such steps would encode that modus). Rules (DM-) encode De Morgan's laws. Finally, rules (Comp-) simply state the compositionality of the procedure.

It is worthwhile observing that in classic backward chaining, loops are avoided by checking whether a new sub-goal (i.e., a literal to be derived) is already on the goal stack (i.e., is currently being derived). Component  $\mathcal{D}$  in Table 5 is in charge of keeping track of the current goals, but this is done on a local basis as opposed to the traditional global stack, that would result if  $\mathcal{D}$  were treated as a global variable. As shown in § 6.3, in our

Table 5: Synthesising the propositional formula  $\llbracket l \rrbracket$  for the attack tree  $T_l$ .

$\llbracket l \rrbracket = \llbracket \varphi \rrbracket \emptyset$	where $(\varphi \Rightarrow \bar{l}) \in P_{\Rightarrow}^l$	(Sel)
$\llbracket \bar{c} \rrbracket \mathcal{D} = \bar{c} \vee$	$\begin{cases} \llbracket \varphi \rrbracket (\mathcal{D} \cup \{\bar{c}\}) & \text{if } \bar{c} \notin \mathcal{D}, \text{ where } (\varphi \Rightarrow \bar{c}) \in P_{\Rightarrow}^l \\ \text{ff} & \text{otherwise} \end{cases}$	(Pone-c)
$\llbracket \neg \bar{c} \rrbracket \mathcal{D} =$	$\begin{cases} \llbracket \neg \varphi \rrbracket (\mathcal{D} \cup \{\neg \bar{c}\}) & \text{if } \neg \bar{c} \notin \mathcal{D}, \text{ where } (\varphi \Rightarrow \bar{c}) \in P_{\Rightarrow}^l \\ \text{tt} & \text{otherwise} \end{cases}$	(Tolle-c)
$\llbracket \bar{x} \rrbracket \mathcal{D} = \llbracket \varphi \rrbracket \mathcal{D}$	where $(\varphi \Rightarrow \bar{x}) \in P_{\Rightarrow}^l$	(Pone-x)
$\llbracket \neg \bar{x} \rrbracket \mathcal{D} = \llbracket \neg \varphi \rrbracket \mathcal{D}$	where $(\varphi \Rightarrow \bar{x}) \in P_{\Rightarrow}^l$	(Tolle-x)
$\llbracket \neg(\varphi_1 \wedge \dots \wedge \varphi_n) \rrbracket \mathcal{D} =$	$\llbracket \neg \varphi_1 \rrbracket \mathcal{D} \vee \dots \vee \llbracket \neg \varphi_n \rrbracket \mathcal{D}$	(DM-1)
$\llbracket \neg(\varphi_1 \vee \dots \vee \varphi_n) \rrbracket \mathcal{D} =$	$\llbracket \neg \varphi_1 \rrbracket \mathcal{D} \wedge \dots \wedge \llbracket \neg \varphi_n \rrbracket \mathcal{D}$	(DM-2)
$\llbracket \varphi_1 \wedge \dots \wedge \varphi_n \rrbracket \mathcal{D} =$	$\llbracket \varphi_1 \rrbracket \mathcal{D} \wedge \dots \wedge \llbracket \varphi_n \rrbracket \mathcal{D}$	(Comp-1)
$\llbracket \varphi_1 \vee \dots \vee \varphi_n \rrbracket \mathcal{D} =$	$\llbracket \varphi_1 \rrbracket \mathcal{D} \vee \dots \vee \llbracket \varphi_n \rrbracket \mathcal{D}$	(Comp-2)
$\llbracket \text{tt} \rrbracket \mathcal{D} = \text{tt}$	$\llbracket \text{ff} \rrbracket \mathcal{D} = \text{ff}$	

setting the global stopping criterion would lead to unsound results. Moreover, observe that using the local environment  $\mathcal{D}$  we lose the linear complexity in  $|P_{\Rightarrow}^l|$  typical of backward chaining, and incur an exponential complexity in the worst case. Nonetheless, observe that this theoretical bound is not incurred systematically, but it depends on the shape of the process under study. Finally, a local stack allows to re-use channel literals in rule (Pone-c).

Notice that we do not need to keep track of literals  $\bar{x}$  in  $\mathcal{D}$ , as we cannot meet with a cycle because a variable cannot be used prior to its definition (in virtue of Lemma B.1).

Finally, observe that a parse tree  $T_l$  of  $\llbracket l \rrbracket$  is an attack tree, showing how  $l$  can be reached by combining the knowledge of given channels. The internal nodes of the tree contain a Boolean operator in  $\{\wedge, \vee\}$ , while the leaves contain literals representing the knowledge of channels. As De Morgan's laws are used to push negations to literals of  $\llbracket l \rrbracket$ , negation can only occur in the leaves of  $T_l$ . In the following, we shall manipulate attack trees always at their denotation level, that is, the object under evaluation is  $\llbracket l \rrbracket$  as opposed to its representation  $T_l$  (cf. § 7.1).

For the sake of discussion, it is worthwhile noticing that the procedure for generating  $\llbracket l \rrbracket$  can be used to generate a tree explicitly during the computation, or even an AND-OR graph [37, Ch. 4]. It is unclear to us, however, whether the more compact graph representation would be simpler for non-expert to understand.

Finally, observe that  $\llbracket l \rrbracket$  only contains literals  $\bar{c}$  corresponding to channels, that is, the backward chaining-like procedure described above and formalised in Table 5 eliminates all the literals  $\bar{x}$ . Therefore, the reachability of  $l$  is only expressed in terms of knowledge of channels. This result is formalised in Lemma B.3, and guarantees that a map from channels to cost suffices to quantify an attack.

**6.2. The attack tree for NemID.** Consider the process *NemID* discussed in § 3 and its translation  $NemID_{\Rightarrow}^{13}$ . Figure 4(a) shows the attack tree  $T_{13}$ , as generated by our

implementation, presented in § 7.2. The denotation of  $T_{13}$  is given by the following formula:

$$\begin{aligned} \llbracket 13 \rrbracket &= \overline{\text{login}} \vee \\ &\quad \left( (\overline{\text{cert}} \vee (\overline{\text{id}} \wedge \overline{\text{pwd}} \wedge \overline{\text{otp}})) \wedge \overline{\text{cert}} \right) \vee \\ &\quad \left( (\overline{\text{cert}} \vee (\overline{\text{id}} \wedge \overline{\text{pwd}} \wedge \overline{\text{otp}})) \wedge (\neg \overline{\text{cert}}) \wedge \overline{\text{id}} \wedge \overline{\text{pwd}} \wedge \overline{\text{otp}} \right) \vee \\ &\quad \left( \overline{\text{id}} \wedge \overline{\text{pin}} \right) \end{aligned}$$

As a matter of fact, the algorithm tends to generate simple but redundant formulae, which can be simplified automatically, e.g., via a reduction to a normal form. The following formula, for instance, is equivalent to  $\llbracket 13 \rrbracket$  but highlights more clearly the ways in which an attack can be carried out:

$$\overline{\text{login}} \vee (\overline{\text{id}} \wedge \overline{\text{pin}}) \vee (\overline{\text{id}} \wedge \overline{\text{pwd}} \wedge \overline{\text{otp}}) \vee \overline{\text{cert}}$$

Observe that the formula above is in Disjunctive Normal Form (DNF). Such normal form has the merit of providing an immediate intuition of the alternative conditions that lead to reach the program point under study, as displayed in Fig. 4(b). However, the conversion to DNF may cause an exponential blow-up in the number of literals, and compact translations require to introduce fresh atoms, garbling the relation between the tree and the original system. Therefore, we did not implement such conversion in the tool.

Finally, notice that the disjunct  $\overline{\text{login}}$  encodes the possibility of obtaining a login token in any other way not foreseen in the system, and thus accounts for all the attacks not explicitly related to the shape of our formalisation. Such possibility can be ruled out by assigning the maximum possible cost to the channel, as we have seen in § 5.4,

**6.3. Global stopping criterion.** We conclude this section by showing why the global stopping criterion is unsound for the procedure of Table 5. Consider the following set of formulae:

$$\bar{a} \Rightarrow \bar{b} \quad \bar{b} \Rightarrow \bar{a} \quad \bar{a} \wedge \bar{b} \Rightarrow \bar{\gamma}$$

which stems from a conveniently simplified translation of the process

$$P \triangleq {}^1a?x_a.{}^2b!b \mid {}^3b?x_b.{}^4a!a \mid {}^5a?x'_a.{}^6b?x'_b.{}^7c!c$$

The generation of  $\llbracket 7 \rrbracket$  unfolds as follows:

$$\llbracket 7 \rrbracket = \llbracket \bar{a} \wedge \bar{b} \rrbracket \emptyset = \llbracket \bar{a} \rrbracket \emptyset \wedge \llbracket \bar{b} \rrbracket \emptyset$$

where, in particular, it is

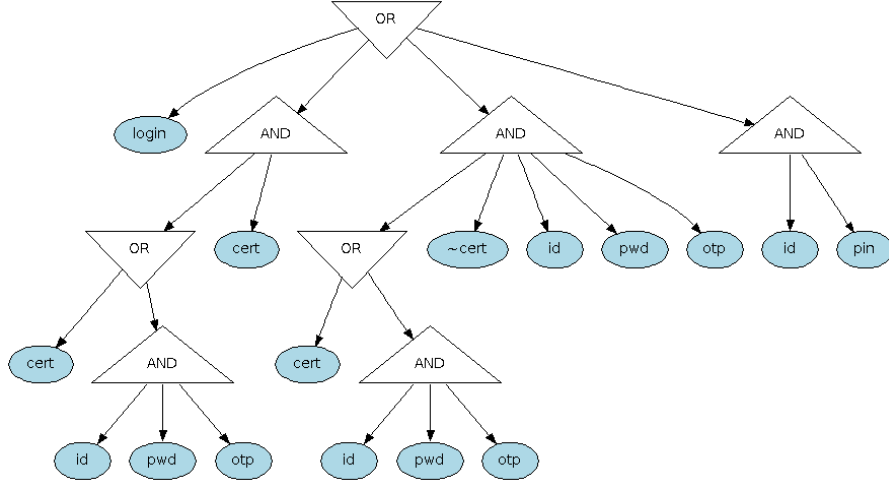

$$\begin{aligned} \llbracket \bar{a} \rrbracket \emptyset &= \bar{a} \vee \llbracket \bar{b} \rrbracket \{\bar{a}\} = \bar{a} \vee \bar{b} \vee \llbracket \bar{a} \rrbracket \{\bar{a}, \bar{b}\} = \bar{a} \vee \bar{b} \vee \text{ff} = \bar{a} \vee \bar{b} \\ \llbracket \bar{b} \rrbracket \emptyset &= \bar{b} \vee \llbracket \bar{a} \rrbracket \{\bar{b}\} = \bar{b} \vee \bar{a} \vee \llbracket \bar{b} \rrbracket \{\bar{b}, \bar{a}\} = \bar{b} \vee \bar{a} \vee \text{ff} = \bar{b} \vee \bar{a} \end{aligned}$$

leading to  $\llbracket 7 \rrbracket = \bar{a} \vee \bar{b}$ , which is consistent with the reachability of label  $\gamma$  in  $P$ .

Assume now to carry out the generation of  $\llbracket 7 \rrbracket$  applying a global stopping criterion, that is, to keep track of derived goals in a global environment, initially empty. We would obtain:

$$\llbracket \bar{a} \rrbracket \emptyset = \bar{a} \vee \llbracket \bar{b} \rrbracket \{\bar{a}\} = \bar{a} \vee \bar{b} \vee \llbracket \bar{a} \rrbracket \{\bar{a}, \bar{b}\} = \bar{a} \vee \bar{b} \vee \text{ff} = \bar{a} \vee \bar{b}$$

at this point, however, the environment contains  $\bar{a}, \bar{b}$ , and thus the generation of  $\llbracket \bar{b} \rrbracket$  leads to  $\bar{b}$ , resulting in  $\llbracket 7 \rrbracket = (\bar{a} \vee \bar{b}) \wedge \bar{b}$ , which is not satisfied by the model where only  $\bar{a}$  is tt, and thus is wrong. Analogously, we would obtain a wrong result if we chose to unfold  $\llbracket \bar{b} \rrbracket$  before  $\llbracket \bar{a} \rrbracket$ .

(a)  $T_{13}$  as displayed by the Quality Protection Tool, presented in § 7.2.

(b) The simplified DNF attack tree.

Figure 4: The attack tree  $T_{13}$  of the NemID example.

## 7. IMPLEMENTATION OF THE ANALYSIS

**7.1. Discussion.** As the backward-chaining procedure on  $P_{\Rightarrow}^l$  re-establishes bi-implications and only applies valid inference rules,  $\llbracket l \rrbracket$  and  $P_{\Leftrightarrow}^l$  are equisatisfiable. They are not equivalent, i.e., in general their models do not coincide, as  $\llbracket l \rrbracket$  only contains channel literals, but they contain the same *attacks*. Hence, we can solve the quantitative version of the protection analysis in either way:

- generate  $P_{\Leftrightarrow}^l$ , compute the models bearing minimal attacks, and extract the corresponding attacks; or,
- generate  $P_{\Rightarrow}^l$ , derive  $\llbracket l \rrbracket$ , and compute its minimal models,

finally relating the result to the security lattice by means of the function `level`.

It is worthwhile observing that while the procedure for generating trees is exponential in the worst case, the size of  $\llbracket l \rrbracket$  is much smaller than the size of  $P_{\Leftrightarrow}^l$ , and therefore it is not necessarily the case that the overall running time would increase when undertaking the tree generation. Though our example set is not extensive enough for supporting any final claim, still it is interesting to comment briefly how the analyses on  $P_{\Leftrightarrow}^l$  and on  $\llbracket l \rrbracket$  behave in terms of running time.

Consider the NemID system. The translation to  $P_{\Leftrightarrow}^l$  takes about one fourth of the time the computation of  $\llbracket l \rrbracket$  takes (in the order of seconds). Solving the optimisation problem on  $P_{\Leftrightarrow}^l$  takes about 1.25 the time it takes on  $\llbracket l \rrbracket$ . Nonetheless, the second step is much more demanding in terms of performance, so that on average the two approaches take the same amount of time. The same applies to the example discussed in [44]. Increasing the size of the process under study it seems that the generation of attack trees, while exponential in general, tends to outperform the overall analysis on  $P_{\Leftrightarrow}^l$ .

It is worthwhile noticing that comparing the two approaches reduces to establishing whether it is faster to find models of  $P_{\Leftrightarrow}^l$  or  $\llbracket l \rrbracket$ , for the translation time is negligible as the size of processes increases. Even limiting to the core propositional structure of the problem, there is no conclusive answer to the question, whose investigation falls outside the scope of this work. For an introduction to the problem of efficiency of satisfiability the reader is referred to [22, Ch. 9,13] and [37, § 7.6.3]; research work on the subject is for instance in [1, 12]. For performance of SMT solvers refer to <http://smtcomp.sourceforge.net/>.

Finally, let us remark that our evaluation of a tree  $T_l$  is, by definition, the evaluation of  $\llbracket l \rrbracket$ . In other words, the minimal cost of a tree is the minimal cost of a model  $\mu$  of  $\llbracket l \rrbracket$  - we evaluate  $\llbracket l \rrbracket$ , as opposed to evaluating  $T_l$ . Another approach would be to evaluate the tree itself by means of traversing its structure, which however would lead to different results, unless a multiset-based model is adopted, such as the one proposed by Mauw and Oostdik [23]. A more precise treatment of the Boolean approach to analyzing attack trees is presented in [4].

**7.2. The Quality Protection Tool.** A proof-of-concept implementation of the framework has been developed in Java and is available at

<http://www.imm.dtu.dk/~rvig/quality-protection.html>

together with the code for the NemID example described in the text.

The *Quality Protection Tool* takes as input an ASCII representation of a Value-Passing Quality Calculus process  $P$  and a label  $l$ , and generates the flow constraints  $P_{\Rightarrow}^l$ .

Furthermore, the tool implements the backward-chaining procedure of § 6.1, relying on our own simple infrastructure for propositional logic, as available libraries tend to avoid the explicit representation of implications, that is instead handy in our case during the backward-chaining computation with non-Horn-like clauses. Once the backward-chaining procedure is executed, and thus  $\llbracket l \rrbracket$  has been derived, the tool can graphically represent the corresponding tree  $T_l$ , thanks to an encoding in DOT<sup>2</sup> and using ZGRViewer<sup>3</sup> for displaying the tree.

Finally, we have implemented the optimisation loop on top of the Z3 SMT solver<sup>4</sup> (Java API). The tool resorts to Z3 for numerical cost sets, optimising the sum of the costs. As for symbolic and non-linearly-ordered cost sets, a finite lattice can be fed into the tool, and the least upper bound is used as monoid operator. Costs can be specified in two ways: numeric costs can be directly fed to the tool, while before specifying symbolic costs the finite lattice  $(\mathcal{K}, \sqsubseteq)$  has to be loaded. In order to specify a lattice, one has to declare  $\top$  and  $\perp$ , and then the operator  $\oplus$  as a list of entries  $x \oplus y = z$ . The names of the elements of the lattice and the partial order  $\sqsubseteq$  are automatically inferred from the graph of  $\oplus$ .

<sup>2</sup><http://www.graphviz.org/>

<sup>3</sup><http://zvtm.sourceforge.net/zgrviewer.html>

<sup>4</sup><http://z3.codeplex.com/>

All these components are glued together thanks to a simple graphical interface. The implementation of the analysis based on  $P_{\Leftrightarrow}^l$  [44] is also available and this, together with improvements on the translation from processes to clauses, is the main difference with respect to the tool of [44].

## 8. RELATED WORK

**8.1. Protection analysis.** Our work is inspired by a successful strand of literature in protocol verification, where a protocol is translated into a set of first-order Horn clauses and resolution-based theorem proving is used to establish security properties [31, 45], and by the flow logic approach to static analysis [34]. In particular, the translation from processes to propositional formulae is inspired by ProVerif [9] translation of protocols into first-order Horn clauses, but can be more formally understood as a flow logic where the carrier logic is not the usual Alternation-free Least Fixed Point Logic, since it cannot express optimisation problems. Moreover, the main problem we discuss is propositional; ideas for a first-order extension are sketched in § C.

In order to formalise the “need for protection” of a location we resort to security lattices, that are widely used for describing levels of security in access control policies. An excellent introductory reference is [3, Chs. 6,7].

As for the solution technique we exploit, different approaches have been presented to solve optimisation problems via SMT. In particular, Nieuwenhuis and Oliveras [29] proposed to modify the DPLL( $T$ ) procedure inherent to SMT solvers so as to look for optimal assignments, while Cimatti et al. [11] developed the search for an optimal assignments on top of an SMT solver, as we do in § 5.3. Nonetheless, both these works focus on numeric weights, which in our settings are represented with linearly ordered cost structures. Our more general notion of weight is modelled after Meadows’s cost sets, formalised as monoids in [24].

Finally, another perspective on the technical developments underpinning the analysis points to computing *prime implicants* of a given formula [15], where in our case primality is sought with respect to the cost set.

**8.2. Attack trees.** Graphical representations of security threats are often used to convey complex information in an intuitive way. Formalisation of such graphical objects are referred to chiefly as *attack graphs* [32, 17, 40, 25] and *attack trees* [38, 39, 23, 33, 18]. In this work we prefer the phrase “attack trees”, but our procedure can be adapted to generate attack graphs. We refer the reader to [21] for a recent survey on the vast literature about attack trees, while in the following we retrace some of historical developments on modelling, generating, and analysing attack trees that inspired our developments.

While different authors have different views on the information that should decorate such objects, instrumental to the analysis that the tree or the graph is supporting, all definitions share the ultimate objective of showing how atomic attacks (i.e., the leaves) can be combined to attain a target goal (i.e., the root). This perspective is enhanced in the seminal work of Schneier [38], that found a great many extensions and applications. In particular, Mauw and Oostdijk [23] lay down formal foundations for attack trees, while Kordy et al. [19] and Roy et al. [36] suggest ways to unify attacks and countermeasures in a single view. Even though Schneier’s work is mostly credited for having introduced attack trees, and it had

certainly a crucial role in making attack trees mainstream in computer security, the origin of this formalism can be traced back to fault trees, expert systems (e.g., Kuang [5]), and privilege graphs [13].

As for the automated generation of attack graphs, the literature is skewed towards the investigation of network-related vulnerabilities: available tools expect as input rich models, including information such as the topology of the network and the set of atomic attacks to be considered. The backward search techniques of Phillips and Swiler [32] and Sheyner et al. [39] have proven useful to cope with the explosion of the state space due to such expressive models. However, the search has to be carried out on a state space that is exponential in the number of system variables, whose construction is the real bottle-neck of these approaches, and the result graph tends to be large even if compact BDD-based representations are used, as argued in [2]. In particular, in [39] a model checking-based approach is developed, where attack graphs are characterised as counter-examples to safety properties; a detailed example is discussed in [40]. Similarly to Phillips and Swiler, we adopt an attacker-centric perspective, which cannot simulate benign system events such as the failure of a component, as in [39]. Directly addressing the exponential blow-up of [39], Ammann et al. [2] propose a polynomial algorithm, but the drop in complexity relies on the assumption of monotonicity of the attacker actions and on the absence of negation. On the same line, Ou et al. [30] present an algorithm which is quadratic in the number of machines in the network under study.

As for the analyses developed on top of attack trees, we present a reachability analysis which computes the cheapest sets of atomic attacks that allow attaining a location of interest in the system, as it is standard in the attack tree literature. This approach seamlessly encompasses the probabilistic analysis of [39, 40] (costs to atomic attacks would represent their likelihood and the objective function would compute the overall probability) and offers a uniform framework to address other quantitative questions [8, 20]. The NP-completeness of our SMT-based approach is in line with the complexity of the minimisation analysis of [39, 40].

Finally, Mehta et al. [25] present a technique for ranking sub-graphs so as to draw attention to the most promising security flaws. Whilst we do not directly tackle this issue, for condensing an entire tree into a formula we gain in performance but we lose the original structure, a post-processing step could be undertaken to compute the value of the internal nodes (sub-formulae).

## 9. CONCLUSION

Discovering attacks is an essential part of investigating security. Quantifying the attacks is necessary when dealing with complex systems and facing budget considerations. Both tasks risk to become a fruitless exercise if their findings cannot be communicated effectively to decision-makers.

Static analysis of process-algebraic specifications offers a unifying framework where these three challenges can be addressed uniformly and by means of modular developments. Our approach can be exploited in a great many context to temper qualitative verification methods with quantitative considerations.

To support this claim, we have developed a protection analysis over the Value-Passing Quality Calculus where attacks can be automatically inferred, quantified, and displayed. At the heart of the analysis lies the abstraction of security mechanisms with secure channels,



which allows to define security checks as input actions and thus attacks as sets of channels over which communication must take place. Moreover, channels lend themselves naturally to support cost considerations, as they represent different security mechanisms.

Starting from this basic assumption, we have developed a qualitative analysis to discover all sets of channels in terms of models of a logic representation of the system under study, resorting to propositional satisfiability (SAT). Levering costs to channels, we have then enhanced the analysis with a quantitative layer, enriching the original SAT problem with the notion of cost of a model and developing an SMT-based optimisation procedure for computing optimal models. Finally, by means of a backward-chaining search on the constraints representing a system, we have shown how to infer an attack tree for a given target, leading to another characterisation of the quantitative analysis, possibly cheaper to compute.

Our SMT-based optimisation allows reasoning with symbolic and non-linearly ordered cost structures, as it is often more natural to describe the relationships between different protection mechanisms instead of assigning them absolute numbers. What is more, this technique is exploitable in all the contexts where models of a formula have to be ranked according to given criteria.

A step necessary to exploit fully the expressiveness of our SMT approach to optimisation is the refactoring of the implementation, so as to produce a stand-alone version of the solution engine, which takes as input an SMT problem (problem, cost lattice, objective function) and returns its optimal models. This would allow to compare our technique to those surveyed above and in particular with the forthcoming optimising version of Z3.

On the modelling side, we present in Appendix C how to lift the developments to the full Quality Calculus. Nonetheless, it is unclear to us whether the resulting attack trees would benefit their intended users, for the additional information may reduce readability drastically. It seems instead promising to investigate further the notion of priced restriction discussed in § 5.6, and to compare its expressiveness to other approaches recently presented in the literature. Finally, as highlighted by Meadows [24], the monoid of the cost set needs not be commutative, as the order in which costs are paid might influence their combination. It would be interesting to investigate mechanisms for re-determining costs dynamically, as a process is evaluated.

#### ACKNOWLEDGEMENT

Special thanks to Zaruhi Aslanyan and Alessandro Bruni for many inspiring and fruitful discussions.

#### REFERENCES

- [1] Dimitris Achlioptas. Random Satisfiability. In *Handbook of Satisfiability*, volume 185, pages 245–270. IOS Press, 2009.
- [2] Paul Ammann, Duminda Wijesekera, and Saket Kaushik. Scalable, graph-based network vulnerability analysis. In *9th ACM conference on Computer and Communications Security (CCS'02)*, pages 217–224, 2002.
- [3] Edward Amoroso. *Fundamentals of Computer Security Technology*. Prentice-Hall, 1994.
- [4] Zaruhi Aslanyan and Flemming Nielson. Pareto efficient solutions of attack-defence trees. In *Principles of Security and Trust - 4th International Conference, POST 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015, Proceedings*, LNCS, pages 95–114. Springer, 2015.

- [5] Robert W. Baldwin. *Rule Based Analysis of Computer Security*. PhD thesis, MIT, 1987.
- [6] Fabrizio Biondi, Axel Legay, Pasquale Malacaria, and Andrzej Wsowski. Quantifying Information Leakage of Randomized Protocols. In *14th International Conference Verification, Model Checking, and Abstract Interpretation (VMCAI'13)*, volume 7737 of *LNCS*, pages 68–87. Springer, 2013.
- [7] Fabrizio Biondi, Axel Legay, and Louis-marie Traonouez. QUAIL : A Quantitative Security Analyzer. In *25th International Conference on Computer Aided Verification (CAV'13)*, volume 8044 of *LNCS*, pages 702–707. Springer, 2013.
- [8] Stefano Bistarelli, Marco Dall'Aglio, and Pamela Peretti. Strategic games on defense trees. In *Formal Aspects in Security and Trust (FAST'06)*, number 3 in *LNCS*, pages 1–15. Springer, 2007.
- [9] Bruno Blanchet. Automatic verification of correspondences for security protocols. *Journal of Computer Security*, 17(4):363–434, 2009.
- [10] Endre Boros and Peter L Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1-3):155–225, 2002.
- [11] Alessandro Cimatti, Anders Franzén, Alberto Griggio, Roberto Sebastiani, and Cristian Stenico. Satisfiability Modulo the Theory of Costs: Foundations and Applications. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *LNCS*, pages 99–113, 2010.
- [12] Amin Coja-Oghla and Konstantinos Panagiotou. Going after the K-SAT Threshold. In *45th ACM symposium on Theory of Computing (STOC'13)*, pages 705–714. ACM, 2013.
- [13] M. Dacier, Y. Deswarte, and M. Kaaniche. Models and tools for quantitative assessment of operational security. In *12th International Information Security Conference (IFIP/SEC'96)*, pages 177–186, 1996.
- [14] Giovanni De Micheli. *Synthesis and Optimization of Digital Circuits*. McGraw-Hill, 1994.
- [15] Isil Dillig, Thomas Dillig, Kenneth L. McMillan, and Alex Aiken. Minimum Satisfying Assignments for SMT. In *Computer Aided Verification (CAV'12)*, volume 7358 of *LNCS*, pages 394–409. Springer, 2012.
- [16] Marco Giunti, Catuscia Palamidessi, and Frank D. Valencia. Hide and New in the Pi-Calculus. In *Proceedings Combined 19th International Workshop on Expressiveness in Concurrency and 9th Workshop on Structured Operational Semantics (EXPRESS/SOS)*, volume 89, pages 65–79, 2012.
- [17] Somesh Jha, Oleg Sheyner, and Jeannette M. Wing. Two formal analyses of attack graphs. In *Proceedings 15th IEEE Computer Security Foundations Workshop CSFW15*, pages 49–63, 2002.
- [18] Aivo Jürgenson and Jan Willemson. Serial Model for Attack Tree Computations. In *Information, Security and Cryptology (ICISC'09)*, volume 5984 of *LNCS*, pages 118–128. Springer, 2010.
- [19] Barbara Kordy, Sjouke Mauw, Sasa Radomirovic, and Patrick Schweitzer. Foundations of Attacks-Defense Trees. In *7th International Workshop on Formal Aspects of Security and Trust (FAST'10)*, volume 6561 of *LNCS*, pages 80–95. Springer, 2010.
- [20] Barbara Kordy, Sjouke Mauw, and Patrick Schweitzer. Quantitative Questions on Attack-Defense Trees. In *15th International Conference on Information Security and Cryptology (ICISC'12)*, volume 7839 of *LNCS*, pages 49–64. Springer, 2012.
- [21] Barbara Kordy, Ludovic Pitre-Cambacds, and Patrick Schweitzer. Dag-based attack and defense modeling: Dont miss the forest for the attack trees. *Computer Science Review*, 1314:1 – 38, 2014.
- [22] R. J. Lipton. *The P=NP Question and Gödel's Lost Letter*. Springer, 2009.
- [23] Sjouke Mauw and Martijn Oostdijk. Foundations of Attack Trees. In *8th International Conference on Information Security and Cryptology (ICISC'05)*, volume 3935 of *LNCS*, pages 186–198. Springer, 2006.
- [24] Catherine Meadows. A cost-based framework for analysis of denial of service in networks. *Journal of Computer Security*, 9(1):143–164, 2001.
- [25] Vaibhav Mehta, Constantinos Bartzis, Haifeng Zhu, Edmund Clarke, and Jeannette Wing. Ranking Attack Graphs. In *9th International Symposium on Recent Advances in Intrusion Detection (RAID'06)*, volume 4219 of *LNCS*, pages 127–144, 2006.
- [26] R. Milner, M. Tofte, R. Harper, and D. MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
- [27] Sebastian Mödersheim and Luca Viganò. Secure Pseudonymous Channels. In *14th European Symposium on Research in Computer Security (ESORICS'09)*, volume 5789 of *LNCS*, pages 337–354. Springer, 2009.
- [28] Flemming Nielson, Hanne Riis Nielson, and René Rydhof Hansen. Validating firewalls using flow logics. *Theoretical Computer Science*, 283(2):381–418, 2002.
- [29] Robert Nieuwenhuis and Albert Oliveras. On SAT Modulo Theories and Optimization Problems. In *Theory and Applications of Satisfiability Testing (SAT'06)*, volume 4121 of *LNCS*, pages 156–169, 2006.

- [30] Xinming Ou, Wayne F. Boyer, and Miles a. McQueen. A scalable approach to attack graph generation. In *Proceedings of the 13th ACM conference on Computer and communications security - CCS '06*, page 336, New York, New York, USA, 2006. ACM Press.
- [31] Lawrence C Paulson. The inductive approach to verifying cryptographic protocols. *Journal of Computer Security*, 6(1-2):85–128, 1998.
- [32] Cynthia Phillips and Laura Painton Swiler. A graph-based system for network-vulnerability analysis. In *Proceedings of the 1998 workshop on New security paradigms NSPW 98*, volume pages, pages 71–79, 1998.
- [33] Martin Reháč, Eugen Staab, Volker Fusenig, Michal Pěchouček, Martin Grill, Jan Stiborek, Karel Bartoš, and Thomas Engel. Runtime Monitoring and Dynamic Reconfiguration for Intrusion Detection Systems. In *Recent Advances in Intrusion Detection (RAID'09)*, volume 5758 of *LNCS*, pages 61–80. Springer, 2009.
- [34] Hanne Riis Nielson, Flemming Nielson, and Henrik Pilegaard. Flow Logic for Process Calculi. *ACM Computing Surveys*, 44(1):1–39, January 2012.
- [35] Hanne Riis Nielson, Flemming Nielson, and Roberto Vigo. A Calculus for Quality. In *9th International Symposium on Formal Aspects of Component Software (FACS'12)*, volume 7684 of *LNCS*, pages 188–204. Springer, 2012.
- [36] Arpan Roy, Dong Seong Kim, and Kishor S. Trivedi. Attack countermeasure trees (ACT): towards unifying the constructs of attack and defense trees. *Security and Communication Networks*, 5(8):929–943, 2012.
- [37] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, 3rd edition, 2009.
- [38] Bruce Schneier. Attack Trees. *Dr. Dobbs's Journal*, 1999.
- [39] Oleg Sheyner, Joshua W. Haines, Somesh Jha, Richard Lippmann, and Jeannette M. Wing. Automated Generation and Analysis of Attack Graphs. In *2002 IEEE Symposium on Security and Privacy*, pages 273–284, 2002.
- [40] Oleg Sheyner and Jeannette M. Wing. Tools for Generating and Analyzing Attack Graphs. In *2nd International Symposium on Formal Methods for Components and Objects (FMCO'03)*, volume 3188 of *LNCS*, pages 344–371, 2004.
- [41] Roberto Vigo. The Cyber-Physical Attacker. In *7th ERCIM/EWICS Workshop on Cyberphysical Systems*, volume 7613 of *LNCS*, pages 347–356. Springer, 2012.
- [42] Roberto Vigo, Flemming Nielson, and Hanne Riis Nielson. Broadcast, Denial-of-Service, and Secure Communication. In *10th International Conference on integrated Formal Methods (iFM'13)*, volume 7940 of *LNCS*, pages 410–427, 2013.
- [43] Roberto Vigo, Flemming Nielson, and Hanne Riis Nielson. Automated Generation of Attack Trees. In *27th Computer Security Foundations Symposium (CSF'14)*, pages 337–350. IEEE, 2014.
- [44] Roberto Vigo, Flemming Nielson, and Hanne Riis Nielson. Uniform Protection for Multi-exposed Targets. In *34th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems (FORTE'14)*, volume 8461 of *LNCS*, pages 182–198. Springer, 2014.
- [45] Christoph Weidenbach. Towards an automatic analysis of security protocols in first-order logic. In *16th International Conference on Automated Deduction (CADE- 16)*, pages 314–328. Springer-Verlag, 1999.

## APPENDIX A. CORRECTNESS OF THE PROTECTION ANALYSIS

The correctness of the protection analysis with respect to the semantics of the calculus is formalised as follows:

$$\text{if } P|Q \Longrightarrow^* C[lP'] \text{ then } \exists \mathcal{N} \in \text{attack}(M^l) \text{ s.t. } \mathcal{N} \subseteq \text{fc}(Q)$$

i.e., for all the executions in which  $Q$  drives  $P$  to  $l$ , the analysis computes a set of channels  $\mathcal{N} \in \text{Names}$  that under-approximates the knowledge required of  $Q$ .

Technically, it is convenient to organise a formal proof in two steps. First, if  $P|Q$  reaches  $l$  then  $P|H[\text{fc}(Q)]$  reaches  $l$ , where process  $H$  is the hardest attacker possible and is parametrised on the knowledge of  $Q$ .  $H$  can be thought as the (infinite) process executing

all possible actions on  $\text{fc}(Q)$ , and the proof simply argues that whatever  $Q$  can,  $H$  can (Fact A.2). A similar approach is detailed in [28].

Finally, the second step shows that if  $P|H[\mathcal{N}']$  reaches  $l$ , then there must be a set  $\mathcal{N} \in \text{attack}(M^l)$  such that  $\mathcal{N} \subseteq \mathcal{N}'$  (Theorem A.3). Observe that this formulation corresponds to the qualitative analysis of § 4. However, as  $\text{attack}(\text{minimal}(M^l)) \subseteq \text{attack}(M^l)$ , the correctness of the quantitative analysis follows as a particular case.

**Definition A.1** (Hardest attacker). *Let  $\mathcal{N} = \{c_1, \dots, c_n\}$  be a finite set of channels.  $H[\mathcal{N}]$  is the process that does all possible sequence of output actions over channels in  $\mathcal{N}$ :*

$$H[\mathcal{N}] \triangleq (\nu d) (!c_1!d) | \dots | (!c_n!d)$$

(where labels are of no use hence omitted).

In the definition we used a fresh name  $d$  as output term, but any name can be chosen as  $P$  cannot check the content of input variables. Observe that the channels in  $\mathcal{N}$  might be used to trigger necessary outputs on other channels, according to the constraints in  $P_{\Leftrightarrow}^l$ .

**Fact A.2.** Let  $P, P', Q$  be processes and  $C, C'$  contexts. It holds that

$$\text{if } P|Q \Longrightarrow^* C[lP'] \text{ then } P|H[\text{fc}(Q)] \Longrightarrow^* C'[lP']$$

As a matter of fact, the only blocking actions in  $P$  are inputs, and since the calculus is value-passing, the execution of  $P$  is driven exclusively by (i) the number of output actions  $Q$  performs, (ii) the channels over which they are executed, and (iii) their order. Now, for each channel  $c \in \text{fc}(Q)$ , that is, for each channel known to  $Q$ , by construction  $H[\text{fc}(Q)]$  interleaves an arbitrary number of output on  $c$ , thus mimicking all the possible sequence of output actions on  $\text{fc}(Q)$ , among which is the one performed by  $Q$ .

The main correctness result is phrased as follows. Since the semantics is value-passing, the proof does not present any particular obstacle, and therefore we limit to present its structure and major cases.

**Theorem A.3** (Correctness of the protection analysis). *Let  $P, P'$  be processes,  $C$  a context, and  $\mathcal{N} \in \text{Names}$  a set of channels. It holds that*

$$\text{if } P|H[\mathcal{N}] \Longrightarrow^* C[lP'] \text{ then } \left( \exists \mathcal{N}' . \mathcal{N}' \in \text{attack}(M^l) \wedge \mathcal{N}' \subseteq \mathcal{N} \right)$$

*Proof sketch.* The proof is organised by induction on the length  $k$  of the derivation sequence  $P|H[\mathcal{N}] \Longrightarrow^* C[lP']$ .

*Basis.* If  $k = 0$ , then it is  $P = C[lP']$ , from which  $\emptyset \in \text{attack}(M^l)$ , for  $\bar{l}$  is a fact in  $P_{\Leftrightarrow}^l$ , and thus it does not entail any channel literal to be **tt**. Since  $\emptyset \subseteq \mathcal{N}$ , for all set  $\mathcal{N}$ , the thesis follows.

*Step.* Assume  $k = k_0 + 1$ . The derivation sequence can be written as

$$P|H[\mathcal{N}] \Longrightarrow^{k_0} C''[lP''] \Longrightarrow C'[lP']$$

for some context  $C''$  and process  $P''$ . The inductive hypothesis applies to the first  $k_0$  steps of the derivation: there exists  $\mathcal{N}'' \in \text{attack}(M^l)$  such that  $\mathcal{N}'' \subseteq \mathcal{N}$ . Now, it suffices to show that the last step in the derivation sequence, leading to reaching  $l$ , preserves the inclusion relationship.

The last reduction  $C''[P''] \Longrightarrow C'[lP']$  is a short-hand writing that conflates a number of cases, but observe that it must be entailed by combining rule (Sys) with a transition

$P'' \xrightarrow{\lambda} P'$ , where we assume that the contexts  $C'', C'$  take care of hiding restrictions preceding  $P''$  and parallel components of  $P'', P'$  that are not affected by the transition. As for the congruence step in the premise of rule (Sys), observe that the rewrite cannot produce inputs or outputs not already considered by the analysis, as the latter always assumes replications to be unfolded. To conclude, a formal proof requires an induction on the shape of the inference tree for the transition  $P'' \xrightarrow{\lambda} P'$ .

Let us comment upon the case of rule (In-tt), which is the most interesting and the only non-trivial. Assume that the binder  $b$  is a simple input  $c?x$ , passing which the label of interest is attained. Now, it is either  $c \in \mathcal{N}''$ , in which case we conclude  $\mathcal{N}' = \mathcal{N}'' \subseteq \mathcal{N}$ , or  $c \notin \mathcal{N}''$ . Again, we have two cases.

If there exists a subset of  $\mathcal{N}''$  which can trigger another component of  $P$  to make an output on  $c$ , we again conclude  $\mathcal{N}' = \mathcal{N}'' \subseteq \mathcal{N}$ . Otherwise, we are in the case  $\varphi \wedge \bar{c} \Leftrightarrow \bar{l}$  with  $c \notin \mathcal{N}''$  and  $\bar{c}$  not a consequence of the literals corresponding to  $\mathcal{N}''$ . In the set of constraints we have  $g_c \vee \varphi' \Leftrightarrow \bar{c}$ . It must then be either  $c \in \mathcal{N}$ , or  $\mathcal{N}''' \subseteq \mathcal{N}$ , where  $\mathcal{N}'''$  satisfies  $\varphi'$ , otherwise  $\mathcal{H}[N]$  would not pass the input. If we look at models of  $P_{\Leftrightarrow}^l$ , we have that either  $g_c$  is tt or  $\varphi'$  evaluates to tt – in every model. In the first case we conclude  $\mathcal{N}' = \mathcal{N}'' \cup \{c\} \subseteq \mathcal{N}$ . In the latter  $\mathcal{N}' = \mathcal{N}'' \cup \mathcal{N}^{iv} \subseteq \mathcal{N}$ , with  $\mathcal{N}^{iv} \subseteq \mathcal{N}'''$ , because the least way of satisfying  $\varphi'$  by the analysis under-approximates the least way of satisfying  $\varphi'$  by the semantics.

The same reasoning applies to the case in which  $b$  is a quality binder, as formulae are computed according to the semantics of quality binders.

## APPENDIX B. PROPERTIES OF ATTACK TREES

This appendix contains some results that substantiate the procedure for generating attack trees discussed in § 6.1.

**Lemma B.1.** *Let  $P$  be a process. For any variable  $x$  in  $P$ , there exists exactly one formula  $\varphi \Rightarrow \bar{x}$  in the translation  $P_{\Rightarrow}^l$ , and  $\bar{x}$  does not occur in  $\varphi$ .*

*Proof.* By induction on the structure of processes. In particular, observe that in Table 4 a literal  $\bar{x}$  is added to  $\varphi$  only when a case clause is met, and by hypothesis  $\bar{x}$  must previously appear in a binder, for processes are closed. Finally, recall that we assume processes to be renamed apart (cf- § 2.1), hence the same variable or name cannot be bound twice.  $\square$

Let us discuss now the *complexity* of the translation given in Table 4. Let  $\text{size}(\mathcal{C})$  denote the number of literals occurring in a set of formulae  $\mathcal{C}$ , that is,  $\text{size}(\mathcal{C}) = \sum_{\varphi \in \mathcal{C}} (\text{size}(\varphi))$ , where  $\text{size}(\varphi)$  counts the literals in  $\varphi$ .

**Lemma B.2.** *Let  $P$  be a process, and assuming that  $P$  contains  $n$  actions. Then  $\text{size}(\llbracket P \rrbracket \text{tt}) = O(n^2)$ .*

*Proof.* If  $P$  consists of  $n$  actions,  $\llbracket P \rrbracket \text{tt}$  consists of at most  $O(n)$  formulae. More in detail,  $\llbracket P \rrbracket \text{tt}$  consists of  $n_o + n_c + n_i + n_b$  formulae,  $n_o$  being the number of outputs in  $P$ ,  $n_c$  the number of case clauses,  $n_i$  the number of simple inputs (including the ones occurring within quality binders), and  $n_b$  the number of binders. The number of literals in a formula depends linearly on the number of actions preceding the label at which the formula is generated (cf. Table 4), hence the number of literals in  $\llbracket P \rrbracket \text{tt}$  is asymptotically bounded by  $n^2$ .  $\square$

It is interesting to observe that from a theoretical point of view  $O(n^2)$  is a precise bound to  $\text{size}(\llbracket P \rrbracket_{\text{tt}})$ . Consider the process  $IN_n$  that consists of  $n$  sequential inputs  $c_1?x_1 \dots c_n?x_n$ . The number of literals in  $\llbracket IN_n \rrbracket_{\text{tt}}$  grows with

$$\begin{aligned} \sum_{i=1}^n (2(i-1) + 3) &= \sum_{i=1}^n (2i + 1) = \\ &= n + 2 \sum_{i=1}^n i = n + 2 \frac{n(n+1)}{2} = \\ &= n^2 + 2n \end{aligned}$$

where  $i$  records the number of literals in the hypothesis  $\varphi$ , we have omitted counting the  $\text{tt}$  conjuncts, and we leverage the fact that an input generates two formulae whose size is  $\text{size}(\varphi) + 2$  adding 1 literal to the hypothesis, from which the relation  $2(i-1) + 3$  is derived. Similarly, the translation of a process made of alternating inputs and case clauses would grow quadratically (with greater constants than  $IN_n$ ).

**Lemma B.3.** *Let  $P$  be a process. For all labels  $l$  occurring in  $P$ , the formula  $\llbracket l \rrbracket$  built according to the rules of Table 5 contains no literal  $\bar{x}$ . In particular,  $\llbracket l \rrbracket$  only contains literals related to channels  $c$ .*

*Proof.* By induction on the number of steps in the unfolding of the generation of  $\llbracket l \rrbracket$ , according to the rules in Table 5.  $\square$

## APPENDIX C. FIRST-ORDER ATTACK TREES

We present in this section an extension to the framework whose detailed development deserves to be deepened in future work. The ideas discussed in the following have not been implemented in the tool of § 7.2.

The notion of knowledge needed to perform an attack adopted so far shifts the semantics load on the concept of secure channel. Besides its simplicity, this abstraction proves useful to model a great many different domains and lead to a sensible notion of attack tree. Nevertheless, it seems interesting to explore less abstract scenarios, where messages exchanged over channels do enjoy a structure and their content is exploitable in the continuation. There is a substantial corpus of literature on how to extend a process calculus to handle reasoning on terms (e.g., via equational theories or pattern matching, cf. [42]), but at the semantic heart of such calculi lies the capability of testing if what is received matches what was expected.

In order to fully encompass the original Quality Calculus we should introduce both testing capabilities and structured messages. We limit here to show how to deal with the first extension, as it has a wider impact on the technical developments. As a matter of fact, distinguishing between a term  $t$  and an expression  $\text{some}(t)$  we are already dealing with a (very simple) signature, and this gives the necessary insight onto our idea.

The syntax of the Value-Passing Quality Calculus, introduced in § 2.1, is enhanced as follows. First of all, we allow now input and output channels to range over terms  $t$ , writing  $t?x$  and  $t_1!t_2$ . In particular,  $t$  can be a variable  $y$ , realising name-passing. Second, we update the case clause as  ${}^l\text{case } x \text{ of } \text{some}(t) : P_1 \text{ else } P_2$ , allowing to check the data payload (if any)

of an input variable  $x$ . The semantics of § 2.2 is modified accordingly:

$$\begin{aligned}
& \text{}^l\text{case some}(c) \text{ of some}(c): P_1 \text{ else } P_2 \xrightarrow{\tau} P_1 \\
& \text{}^l\text{case some}(c) \text{ of some}(y): P_1 \text{ else } P_2 \xrightarrow{\tau} P_1[c/y] \\
& \text{}^l\text{case some}(c) \text{ of some}(c'): P_1 \text{ else } P_2 \xrightarrow{\tau} P_2 \text{ if } c \neq c' \\
& \text{}^l\text{case none of some}(c): P_1 \text{ else } P_2 \xrightarrow{\tau} P_2 \\
& \text{}^l\text{case none of some}(y): P_1 \text{ else } P_2 \xrightarrow{\tau} P_2
\end{aligned}$$

The translation from processes to formulae of § 4.1 is lifted from propositional to first-order logic, so as to account for the richer expressiveness of the case clause:

$$\begin{aligned}
\llbracket \text{}^l\text{case } x \text{ of some}(t): P_1 \text{ else } P_2 \rrbracket \varphi &= \llbracket P_1 \rrbracket (\varphi \wedge \exists \text{fv}(t).(x = \text{some}(t))) \cup \\
& \llbracket P_2 \rrbracket (\varphi \wedge \neg(\exists \text{fv}(t).(x = \text{some}(t)))) \cup \\
& \{\varphi \Rightarrow \bar{l}\}
\end{aligned}$$

where  $\text{some}(\cdot)$  is a unary predicate,  $\text{fv}(t)$  denotes the variables free in  $t$ , and we write  $x$  instead of  $\bar{x}$  for now  $x$  ranges over a set of optional data. Similarly, the translation of binders has now to record the term to which an input variable is bound when the corresponding binder is satisfied:

$$\text{th}(\varphi, t?x) = \{\exists y.(\varphi \wedge t \Rightarrow (x = \text{some}(y)))\}$$

where  $t$  ranges over a set of data (the translation of output has to be updated similarly).

Finally, for building the tree some unification is needed in the backward-chaining search of § 6.1:

$$\begin{aligned}
\llbracket \exists \text{fv}(t)(x = \text{some}(t)) \rrbracket \mathcal{D} &= \llbracket \varphi \sigma \rrbracket \mathcal{D} \\
&\text{where } (\exists \text{fv}(t')(\varphi \Rightarrow (x = \text{some}(t')))) \in P_{\Rightarrow}^l \wedge \exists \sigma.t = t'\sigma \\
\llbracket \neg \exists \text{fv}(t)(x = \text{some}(t)) \rrbracket \mathcal{D} &= \llbracket \neg \varphi \sigma \rrbracket \mathcal{D} \\
&\text{where } (\exists \text{fv}(t')(\varphi \Rightarrow (x = \text{some}(t')))) \in P_{\Rightarrow}^l \wedge \exists \sigma.t = t'\sigma
\end{aligned}$$

where  $\sigma$  is a most general unifier.

We have thus shown how to lift all levels of the framework to name-passing calculi with full testing capabilities. From a high-level perspective, the extension allows inspecting how security checks are performed, while the basic developments consider checks as atomic entities, distinguishing between them through the cost map. There is the concrete risk, however, that the additional information available in the more detailed “first-order” trees would decrease readability drastically. In addition to this, whenever a finer-grained investigation is needed, we could take advantage of the modularity of the propositional framework, as discussed in § 4.5.

Finally, in order to carry the extension to the Quality Protection Tool of § 7.2, the main extension would consist in introducing the unification of terms in the backward-chaining procedure.