



Attribute evaluation on attack trees with incomplete information

Ahto Buldas^a, Olga Gadyatskaya^b, Aleksandr Lenin^a, Sjouke Mauw^b,
Rolando Trujillo-Rasua^{c,*}

^a Tallinn University of Technology and Cybernetica AS, Estonia

^b CSC/SnT, University of Luxembourg, Luxembourg

^c Centre for Cyber Security Research and Innovation, School of Information Technology, Deakin University, Australia

ARTICLE INFO

Article history:

Received 30 April 2019

Revised 21 August 2019

Accepted 29 September 2019

Available online 30 September 2019

Keywords:

Attack trees

Constraint programming

Historical data

Security risk assessment

Quantitative security

Decoration problem

ABSTRACT

Attack trees are considered a useful tool for security modelling because they support qualitative as well as quantitative analysis. The quantitative approach is based on values associated to each node in the tree, expressing, for instance, the minimal cost or probability of an attack. Current quantitative methods for attack trees allow the analyst to, based on an initial assignment of values to the leaf nodes, derive the values of the higher nodes in the tree. In practice, however, it shows to be very difficult to obtain reliable values for all leaf nodes. The main reasons are that data is only available for some of the nodes, that data is available for intermediate nodes rather than for the leaf nodes, or even that the available data is inconsistent. We address these problems by developing a generalisation of the standard bottom-up calculation method in three ways. First, we allow initial attributions of non-leaf nodes. Second, we admit additional relations between attack steps beyond those provided by the underlying attack tree semantics. Third, we support the calculation of an approximative solution in case of inconsistencies. We illustrate our method, which is based on constraint programming, by a comprehensive case study.

© 2019 Elsevier Ltd. All rights reserved.

1. Introduction

Attack trees are a useful and intuitive graphical modelling language. Since the introduction by Amoroso (as threat trees) (Amoroso, 1994), and Salter et al. (1998), it has enjoyed popularity in the security industry, as well as in the research community. Attack trees have been equipped with various semantics (Horne et al., 2017; Jhavar et al., 2015; Mauw and Oostdijk, 2006) and supported by tools (Amenaza, 2017; Gadyatskaya et al., 2016c). They have also been enhanced with various methods for quantitative analysis (Aslanyan and Nielson, 2015; Aslanyan et al., 2016; Bistarelli et al., 2006; Buldas and Lenin, 2013; Kordy et al., 2013b; Kumar et al., 2015; Lenin, 2015; Lenin and Buldas, 2014; Lenin et al., 2015; Roy et al., 2012), which allow determining for a given attack tree, for example, an organisation's losses due to an attack, the probability that such an attack succeeds, or the cost of a successful attack (Hong et al., 2017).

The underlying assumption upon which all these quantification methods are based is similar to the popular *divide and conquer*

paradigm, in which a problem is recursively broken down into smaller problems that are theoretically simpler to reason about and solve.

A quantification method for attack trees often reduces to the assignment of attribute values to basic attack steps (leaf nodes in the tree). Such assignments are used in a bottom-up propagation manner to determine the value at the root node, which is a quantification measure for the scenario expressed in the tree (Kordy et al., 2013b). It is largely believed that it is relatively easy to assign a reliable attribute value to a basic attack step, which is precise and refined enough. Popular tools operating with attack trees, such as ADTool (Gadyatskaya et al., 2016c; Kordy et al., 2013a) and SecurITree (Amenaza, 2017), work exactly under this premise.

In practice, the assumption that attribute values for more concrete attack steps are easier to obtain has proven incorrect. Indeed, most companies manage to obtain statistical data for abstract attacks, e.g. frequency of skimming attacks, while they might struggle to come up with similar data for more refined attacks, e.g. frequency of stereo skimming attacks based on audio technology. For security consultants, it might be feasible to obtain reliable estimations for (at least some) abstract attacks in relevant domains, but precise historical data for low-level attack steps might be out of reach. Thus, we observe that there is a tension between the limited availability of data and the requirement to provide data values

* Corresponding author.

E-mail addresses: ahto.buldas@taltech.ee (A. Buldas), olga.gadyatskaya@uni.lu (O. Gadyatskaya), aleksandr.lenin@cyber.ee (A. Lenin), sjouke.mauw@uni.lu (S. Mauw), rolando.trujillo@deakin.edu.au (R. Trujillo-Rasua).

for all leaves in an attack tree before proceeding with a quantification method.

Today, existing quantitative approaches for attack trees cannot handle values of intermediate nodes in the tree that may become available from historical data. Moreover, they do not support the use of additional constraints over nodes in the tree, which are obtained from external sources of information rather than from the attack tree model itself. For example, the analysts may be confident that card skimming attacks are more frequent than physical attacks on card holders. Such a relation cannot be captured in an attack tree model, because it is not a hierarchical relation, hence it is ignored in current quantitative approaches for attack trees.

There is clearly a need for novel computation methods on attack trees that account for available historical data and domain-specific knowledge. In this paper, we formulate a general *attack-tree decoration problem* that treats assignment of values to tree nodes as a problem of finding a set of data values satisfying a set of predicates. These predicates arise from the attack tree structure and the target attribute to be computed (i.e. semantics) and from the attainable historical values and domain knowledge (i.e. available data). Our methodology to solve the attack-tree decoration problem accounts for scenarios in which the set of predicates cannot be jointly satisfied, due to inconsistencies or possible noise in the data.

Contributions. In this work

- We transform an attack tree semantics together with an attribute interpretation into a constraint satisfaction problem (Section 3). If the attack tree semantics is consistent with the attribute interpretation, this allows us to determine appropriate attribute values for all nodes in the tree.
- Because confidence in the available historical data and domain-specific knowledge may vary, we provide a methodology to deal with inconsistencies (Section 4). The usefulness of our approach is that any consistent valuation is better than no valuation, as it will enable the follow up process of using the attack tree for what-if analysis. The standard bottom-up approach would result in absence of any valuation until all leaf node values can be assigned.
- We introduce two concrete approaches to deal with inconsistencies (Section 5). The first one determines the smallest subset of constraints that makes the decoration problem inconsistent, which is useful to find contradictory or wrong assumptions. The second one is suitable for constraints that are expressed in the form of inequalities. In this approach constraints are regarded consistent and an optimal decoration is always found. The proposed methodology has been implemented as proof-of-concept software tools¹.
- We validate our methodology and the implementations through a comprehensive case study on the security of Automatic Teller Machines (Sections 6 and 7).

2. Related work

Research articles on quantitative security analysis with attack trees in all their flavors (attack-defense trees, defense trees, etc. (Kordy et al., 2014b)) often focus on *providing extensions to attack trees* enabling more complex scenarios (Arnold et al., 2014; Aslanyan et al., 2016; Gadyatskaya et al., 2016a; Jhawar et al., 2016; Kordy et al., 2012) and *defining metrics for evaluating scenarios* captured as attack trees. Various metrics have been considered in the literature, for instance, the probability/likelihood of an attack (Bagnato et al., 2012; Lenin, 2015), expected time until a successful attack (Bagnato et al., 2012), attacker's utility (Buldas

and Lenin, 2013; Lenin, 2015; Lenin and Buldas, 2014; Lenin et al., 2015; 2014), return on security investment (Bistarelli et al., 2006), or assessment of risks (Gadyatskaya et al., 2016b; Potteiger et al., 2016). Bagnato et al. (2012) present a list of metrics found in security literature that can be computed on attack trees. Yet, all these approaches assume that the data values to perform quantitative analysis of system security are readily available. Indeed, to the best of our knowledge, no methodologies have been developed for integrating historical data and domain-specific knowledge in quantitative analysis of attack trees. At the same time, even Bruce Schneier, one of the inventors of attack trees, has acknowledged the painstaking work for data collection that is a prerequisite for quantitative analysis on the trees (Schneier, 2000, Chap. 21).

Benini and Sicari (2008) have proposed a framework for attack tree-based security risk assessment. The approach relies on identification of security vulnerabilities, that are placed in the leaves of an attack tree. Quantitative parameters of the vulnerabilities, such as exploitability and damage potential, allow to estimate security risks to a system. In Benini and Sicari (2008), the exploitability parameters are initially evaluated based on the CVSS scores² of the respective vulnerabilities, and then they are adjusted based on the expert judgement about the system context and mutual effect of vulnerabilities on each other expressed in a vulnerability dependency graph. While this methodology offers more precise quantitative risk assessment with attack trees, it assumes that the attack tree is constructed in a bottom-up manner. All system vulnerabilities have to be identified using a suitable technology, and accommodated in an attack tree. In complex environments it will likely be impractical to apply the bottom-up approach due to the huge amount of potential vulnerability combinations that can be exploited in various attacks. Indeed, in practice attack trees are typically designed in a top-down manner, when the analyst starts by conjecturing the main attacker's goal and iteratively breaks it down into smaller subgoals (Fraille et al., 2016; Mauw and Oostdijk, 2006; Schneier, 1999; 2000).

Recently, de Bijl (2017) studied the use of historical data values to obtain attribute values for attack tree nodes. He proposed several heuristics to deal with missing data values, including the standard bottom-up algorithm to infer parent node values, the reuse of values for recurring nodes, and the use of various data sources to estimate certain attributes. For example, the paper refers to *distance to the police station* as a hidden variable influencing *probability* of attack. Yet, de Bijl does not define a precise methodology to perform computations on attack trees with missing leaf node values.

2.1. Quantitative analysis of fault trees

Fault trees are close relatives of attack trees that are widely used in the reliability domain. There exists a large body of literature dedicated to quantitative analysis of fault trees. Yet, the standard bottom-up approach in attack trees is not the standard approach in fault trees, where min-cut set analysis and translation into more intricate models are common (Ruijters and Stoelinga, 2015).

Fuzzy fault trees address evaluation of fault trees under uncertainty (Mahmood et al., 2013), when failure statistics are not fully available. The main difference is that fuzzy fault trees have been developed to serve the needs of the reliability community and fault tree application methods (fuzzy probability functions, error propagation estimates, etc.), while the security community and attack tree application methods have different needs (bottom-up computation for a large variety of attributes). Therefore, solutions

¹ <https://github.com/vilena/at-decorator/>.

² Common Vulnerability Scoring System <https://www.first.org/cvss>.

designed for fault trees do not fully address the problems in the attack tree space.

2.2. Computations on attack graphs

Attack graphs are a popular related graphical model for security assessment. This model establishes possible attack steps through a system and enables quantitative system security evaluation based on a variety of metrics (probability of compromise, time of attack, etc.) (Holm et al., 2015; Kordy et al., 2014b; LeMay et al., 2011; Singhal and Ou, 2017). The attack graph methodology extensively relies on available statistical data and other available metrics (e.g. vulnerability severity scores). Yet, the issue of uncertainty is known also for quantitative analysis of attack graphs (Sommestad et al., 2009). Still, unlike attack trees, attack graphs are usually automatically generated from available system information (network topology, vulnerabilities present) and thus computations on attack graphs are well-defined. For the traditionally manually designed attack trees, quantitative data are not so readily available because there is no well-defined semantics of the leaf nodes (i.e. it is only up to the analyst to decide to stop refining the nodes further).

2.3. Data issues in quantitative risk assessment

Attack trees are typically used for threat modeling and security risk assessment (Shostack, 2014). Thus, it is necessary to evaluate the data availability perspective also in the more general context of security risk assessment. Indeed, the general question of data validity in quantitative security risk assessment (QSRA for short) and the reliability of QSRA results in presence of uncertainty in data values has been raised by many practitioners and researchers (Vose, 2008).

QSRA enables decision making based on quantitative estimations of some relevant variables (e.g. probability of an event, cost, time, vulnerability, etc.). These quantitative estimations are typically aggregated in a model that can then be utilised by a decision maker (Vose, 2008). Many studies, books on security, and industry reports have acknowledged that the quality of quantitative risk analysis, and, correspondingly, the decisions made based on it, heavily depend on the quality of data used (Baker et al., 2007; Jaquith, 2007; Schneier, 2000; Vose, 2008). Notably, it has been established that probabilities of particular loss events and costs associated to security spending can be hard to obtain from historical data (Ahmed et al., 2007; Aven, 2007; Böhme, 2010; Jaquith, 2007; Oppliger, 2015). This body of knowledge serves as evidence of inherent difficulty to obtain meaningful estimates for probability and cost of detailed attack steps, i.e., values for leaf nodes in attack trees.

Nevertheless, it has been acknowledged that, for instance, for insurance companies it might be feasible to get meaningful data, because they have access to an entire population, i.e. they have good statistics (Oppliger, 2015). It has also been demonstrated, for example, that breach statistics can be used to predict future breaches in different segments (Sarabi et al., 2015), and that statistics pertinent to different user profiles can be applied to estimate success rates of intrusions (Dacier et al., 1996).

Furthermore, for security assessment, it has been long established that external data sources, such as threat level indicators (e.g. malware numbers) can be helpful to update quantitative risk assessment models (Böhme, 2010). Therefore, enabling better usage of available historical data, which may not directly correspond to information about low-level attack events (leaf nodes), will be a valuable enhancement for quantitative analysis of attack trees.

From this review of the relevant scientific literature, we can conclude that there is a strong need for an approach to perform




Fig. 1. An attack tree representing stealing money from someone's bank account.

quantitative analysis on attack trees in case the analyst cannot confidently assign values to all leaf nodes. Furthermore, this approach needs to integrate available historical data that can come in form of values for some abstract attacks (intermediate nodes) or constraints (equalities and inequalities) on combinations of attack tree node values. In the remainder of this paper, we propose such an approach.

3. Attack-tree decoration

In this section we give, to the best of our knowledge, the first formulation of the attack-tree decoration problem as a constraint satisfaction problem. We start by introducing the necessary attack tree basics. The interested reader can find more details about the attack tree formalism in the paper by Mauw and Oostdijk (2006).

3.1. Attack trees

In an attack tree the main goal of the attacker is captured by the root node. This goal is then iteratively refined into subgoals, represented by the children of the root node. Leaf nodes in an attack tree are called *atomic subgoals*, as they are not refined any further. Non-leaf nodes, instead, can be of two types: disjunctive (OR) or conjunctive (AND). A conjunctive refinement expresses that all subgoals must be achieved in order to succeed on the main goal, while in a disjunctive refinement the achievement of a single subgoal is already enough.

Example 3.1. Consider the simple attack tree in Fig. 1. The root node of this tree represents the main goal of the attack: to steal money from a bank account. This goal is disjunctively refined into two alternative sub-attacks: the attacker may try to get money from an automated teller machine (ATM), or they might attempt to hack the online bank account system. The sub-goal that explores getting money at an ATM is further conjunctively refined into two complementary activities: the attacker must steal the credit card of the victim *and* they also need to obtain the PIN code by shoulder-surfing. Note that a conjunctive refinement is denoted graphically with an arc spanning the child nodes.

Definition 3.2 (Attack tree). Given a set of labels \mathcal{L} , an attack tree (ATree) is constructed according to the following grammar (for $\ell \in \mathcal{L}$):

$$t ::= \ell \mid \text{OR}(t, \dots, t)_{\ell} \mid \text{AND}(t, \dots, t)_{\ell}.$$

Our grammar above slightly differs from the grammar used in other notations to represent attack trees (Kordy et al., 2014a; Mauw and Oostdijk, 2006), as we require every node in the tree to

be annotated with a label ℓ . The reason for this is that, as opposed to standard attack tree semantics that focus on the leaf nodes, we render every node in the tree equally important.

To provide a definition of our running example we will use shorter labels than those in Fig. 3.2. The actual mapping between labels should become clear through a quick visual inspection.

$\text{OR}(\text{AND}(\text{card}, \text{pin})_{\text{money-atm}}, \text{hack-account})_{\text{money-account}}$.

We say that an attack tree has *unique labels* if it does not contain two distinct nodes with the same label. We use \mathbb{T} to denote the universe of attack trees. We also use the auxiliary functions $\text{root} : \mathbb{T} \rightarrow \mathcal{L}$ and $\text{labels} : \mathbb{T} \rightarrow \mathcal{P}(\mathcal{L})$ to obtain, respectively, the root node's label and all labels of a given tree. Formally,

- $\text{root}(t) = \ell \iff t \equiv \ell \vee t \equiv \text{OR}(t_1, \dots, t_n)_\ell \vee t \equiv \text{AND}(t_1, \dots, t_n)_\ell$ for some $t_1, \dots, t_n \in \mathbb{T}$
- $\text{labels}(t) = \{\ell\}$ if $t \equiv \ell$, otherwise $\text{labels}(t) = \{\ell\} \cup \text{labels}(t_1) \cup \dots \cup \text{labels}(t_n)$ when $t \equiv \text{OR}(t_1, \dots, t_n)_\ell \vee t \equiv \text{AND}(t_1, \dots, t_n)_\ell$ for some $t_1, \dots, t_n \in \mathbb{T}$

For example, given the tree from Fig. 1, we have that $\text{root}(\text{AND}(\text{card}, \text{pin})_{\text{money-atm}}) = \text{money-atm}$ and $\text{labels}(\text{AND}(\text{card}, \text{pin})_{\text{money-atm}}) = \{\text{card}, \text{pin}, \text{money-atm}\}$.

3.2. The attack-tree decoration problem

We proceed by formulating the attack-tree decoration problem as a constraint satisfaction problem. Intuitively, we map an attack tree to a set of boolean expressions whose variables are drawn from the set of labels of the tree. Such a set of boolean expressions, defined over a given domain, can be seen as a constraint satisfaction problem whose solutions correspond to solutions of the attack-tree decoration problem. The remainder of this section is dedicated to formalising this intuition.

Decorating an attack tree is a process whereby nodes in the tree are assigned with values. Given an attack tree t , we use a total function $\alpha : \text{labels}(t) \rightarrow D$ from labels of the tree to values in a domain D to represent the decoration process, and $\text{VAL}_{\text{labels}(t) \rightarrow D}$ to denote the universe of such functions. To that effect, we often refer to labels as variables and to α as a *valuation*. The co-domain D of a valuation is determined by the attribute of the tree under evaluation. For example, *minimum time of a successful attack* uses the natural number domain \mathbb{N} to express discrete time, while *required attacker skill to succeed* typically uses a discrete and categorical domain, such as $\{\text{low}, \text{medium}, \text{high}\}$.

Definition 3.3 (Attribute semantics). Given an attack tree t and a domain D , an *attribute semantics* is a set of valuations with domain $\text{labels}(t)$ and co-domain D .

Intuitively, a semantics provides an attribute with the set of valuations that the attribute regards as valid in a given tree. Because defining an attribute semantics by exhaustive enumeration of its valuations might be cumbersome, we consider in this article attributes whose semantics can be derived from a constraint satisfaction problem.

An *attribute constraint* is defined as a boolean expression over the set of labels of a tree. To that effect, when we use labels in expressions we will consider them as variables over a given domain D . For example, if the attribute *minimum time taken by an attack* is being computed over an attack tree of the form $t \equiv \text{OR}(t_1, \dots, t_k)_\ell$, it is typically required that $\ell = \min(\text{root}(t_1), \dots, \text{root}(t_k))$ (Kordy et al., 2012). The intuition for such constraint is that, because ℓ is disjunctively refined, the minimum time needed by an attacker to meet the goal ℓ is considered to be the least time required by any of ℓ 's children.

We use predicates as short-hand notations for boolean expressions. For example, $\text{min-time}(\ell, \ell_1, \dots, \ell_n)$ can be used to denote the boolean expression $\ell = \min(\ell_1, \dots, \ell_n)$. We say that a predicate $p(\ell_1, \dots, \ell_n)$ is valid under interpretation α , denoted $p(\ell_1, \dots, \ell_n) \vdash \alpha$, if $p(\alpha(\ell_1), \dots, \alpha(\ell_n))$ evaluates to **true**. Likewise, a set of predicates A is said to be valid under α , denoted $A \vdash \alpha$, if all predicates in A are valid under α . When it does not lead to confusion, we will often refer to a predicate $p(\ell_1, \dots, \ell_n)$ as p .

Definition 3.4 (Attribute constraint-set). Given an attack tree t and a domain D , an *attribute constraint-set* is a set of predicates $\{p_1, \dots, p_n\}$ over $\text{labels}(t)$ whose variables range over D . Its semantics is defined by $\llbracket \{p_1, \dots, p_n\} \rrbracket_t = \{\alpha \in \text{VAL}_{\text{labels}(t) \rightarrow D} \mid \{p_1, \dots, p_n\} \vdash \alpha\}$.

There exist in literature various ways to relate the value of a parent node in an attack tree to the values at its children (Kordy et al., 2012), of which the bottom-up approach is the most common one (Kordy et al., 2014a; Mauw and Oostdijk, 2006). This bottom-up approach starts from an assignment of concrete values to the leaf nodes of the tree and uses two functions (one for disjunctive refinement and one for conjunctive refinement) to recursively calculate the value of a parent node from the values of its children. We will next define how an attribute constraint-set can be recursively derived from two unranked aggregation operators associated with a bottom-up approach. The actual values of the leaf nodes will have to be defined by the analyst through additional constraints.

Definition 3.5 (Bottom-up attribute constraint-set). Let t be an attack tree, and \vee and \wedge two unranked function symbols (symbols without fixed arity) with domain D . We use $\check{p}(\ell_1, \dots, \ell_{n+1})$ to denote the expression $\ell_1 = \vee(\ell_2, \dots, \ell_{n+1})$. Similarly, the predicate $\hat{p}(\ell_1, \dots, \ell_{n+1})$ denotes the boolean expression $\ell_1 = \wedge(\ell_2, \dots, \ell_{n+1})$. The *bottom-up attribute constraint-set* $P(t)$ of t is recursively computed as follows:

- If $t \equiv \ell$ for some label ℓ , then $P(t) = \emptyset$.
- If $t \equiv \text{OR}(t_1, \dots, t_n)_\ell$ with $\text{root}(t_i) = \ell_i$ for $i \in \{1, \dots, n\}$, then $P(t) = P(t_1) \cup \dots \cup P(t_n) \cup \{\check{p}(\ell, \ell_1, \dots, \ell_n)\}$;
- If $t \equiv \text{AND}(t_1, \dots, t_n)_\ell$ with $\text{root}(t_i) = \ell_i$ for $i \in \{1, \dots, n\}$, then $P(t) = P(t_1) \cup \dots \cup P(t_n) \cup \{\hat{p}(\ell, \ell_1, \dots, \ell_n)\}$.

Definition 3.5 is based on the standard bottom-up approach (Kordy et al., 2014a; Mauw and Oostdijk, 2006) for attack trees where child nodes are aggregated together based on two functions: \wedge for children of a conjunctive refinement and \vee for children of a disjunctive refinement. In literature there exist concrete definitions of \wedge and \vee for various attributes. For example, when computing probability of success it is usually considered that $\wedge = \times$ and $\vee = +$, for cost $\wedge = +$ and $\vee = \min$, and for minimum time $\wedge = \max$ and $\vee = \min$.

Definition 3.6 (The attack-tree decoration problem). Given an attack tree t and an attribute constraint-set $\{p_1, \dots, p_n\}$ for t , the *attack-tree decoration problem* consists in finding a valuation in $\llbracket \{p_1, \dots, p_n\} \rrbracket_t$.

The attack-tree decoration problem corresponds to the well-known *Constraint Satisfaction Problem* (CSP) (Tsang, 1993), where a solution is a valuation that satisfies a set of constraints. Finding a solution or even deciding whether there exists a solution for CSP is a well-known and complex computational problem.

We say that the attack-tree decoration problem is:

- *Determined*: If the cardinality of $\llbracket \{p_1, \dots, p_n\} \rrbracket_t$ is one, i.e. there exists a single valid valuation only.
- *Inconsistent*: If $\llbracket \{p_1, \dots, p_n\} \rrbracket_t = \emptyset$, i.e. there does not exist a valid valuation.

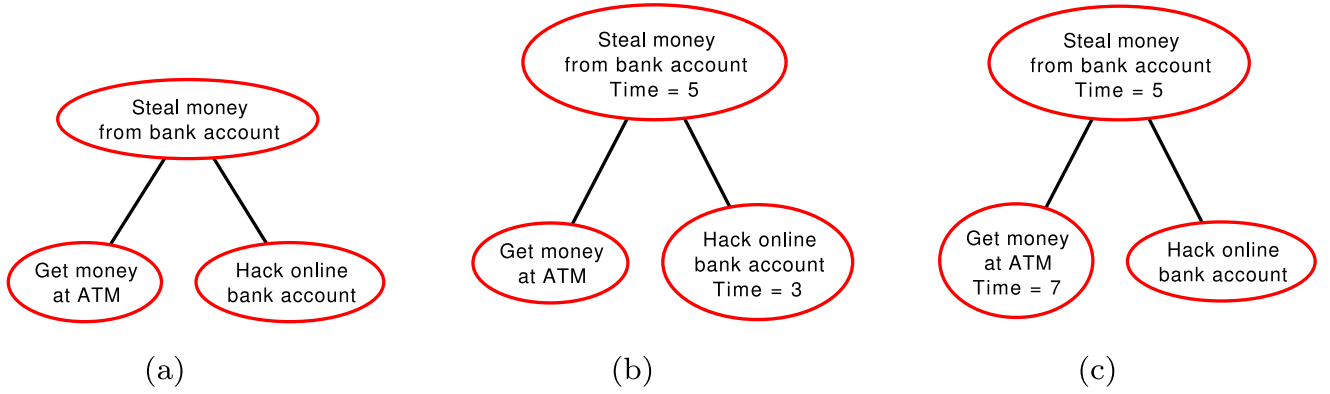


Fig. 2. (a) An example of undeterminism. (b) An example of inconsistency. (c) A determined attribute domain.

- *Undetermined*: If the cardinality of $\llbracket \{p_1, \dots, p_n\} \rrbracket_t$ is larger than one, i.e. the problem is neither inconsistent nor determined.

We illustrate these concepts with the following example that utilises a subtree from our running example. Consider the tree t depicted in Fig. 2a whose set of labels is

$labels(t) = \{money-account, money-atm, hack-account\}$.

Further, we consider subsets of the following set of constraints:

$p_1 := money-account = \min(money-atm, hack-account)$

$p_2 := money-account = 5$

$p_3 := hack-account = 3$

$p_4 := money-atm = 7$

Predicate p_1 follows from the standard interpretation of minimum attack time in an attack tree (Kordy et al., 2013b), where $\wedge = \max$ and $\vee = \min$ defined over the natural numbers \mathbb{N} . This leads to the bottom-up attribute constraint-set $\{p_1\}$. The attack-tree decoration problem with this attribute constraint-set is clearly undetermined given that, for example, the valuations $\{money-account \mapsto 3, money-atm \mapsto 3, hack-account \mapsto 5\}$ and $\{money-account \mapsto 3, money-atm \mapsto 3, hack-account \mapsto 4\}$ both satisfy p_1 .

Now, consider predicates p_2 , p_3 and p_4 . This type of boolean expressions represent variable assignments. We observe that the attribute constraint-set $\{p_1, p_2, p_3\}$ leads to an inconsistent decoration problem (see Fig. 2b), while the decoration problem defined by $\{p_1, p_2, p_4\}$ is determined as there exists a unique valuation satisfying all three predicates, namely $\{money-account \mapsto 5, money-atm \mapsto 7, hack-account \mapsto 5\}$ (see Fig. 2c). Finally, we remark that the attribute constraint-set $\{p_1, p_3, p_4\}$ also leads to a decoration problem that is determined. Moreover, it corresponds to the standard bottom-up calculation in attack trees.

This last example illustrates the general observation that, given an assignment of values to the leaves of an attack tree and a bottom-up attribute constraint-set (Definition 3.5), the attack-tree decoration problem is determined. This is formalized in the following proposition, which can be easily proved by induction on the structure of the tree.

Proposition 3.7. *Let t be an attack tree with unique labels, let $L = \{l_1, \dots, l_n\}$ be the set of labels of its leaf nodes and let D be a domain. Let $P_L = \{l_1 = v_1, \dots, l_n = v_n\}$ be a set of constraints assigning values $v_1, \dots, v_n \in D$ to the leaf nodes and let $P(t)$ be the bottom-up attribute constraint-set of t . Then the attack-tree decoration problem for t and constraint-set $P_L \cup P(t)$ is determined.*

4. A methodology for attack-tree decoration

As indicated above, our approach extends the rather rigid bottom-up way in which attack trees are currently decorated. Our methodology consists of two main steps that complement each other: (1) generation of the attribute constraint-set and (2) analysis of valid valuations. The former boils down to the definition of predicates over the set of labels of a tree. We make a distinction between two types of predicates: hard predicates and soft predicates.

4.1. Hard predicates

Hard predicates are derived from the attack tree refinement structure rather than from knowledge databases or an expert's opinion. This choice establishes that all predicates derived from the attack tree structure should be satisfied, as otherwise the attribute semantics and the tree contradict each other. The term *hard* stems from the notion of hard and soft constraints in satisfaction problems. Soft constraints represent desirable properties, while hard constraints are a must.

In this article, we consider hard predicates those contained in the bottom-up attribute constraint-set of the tree (see Definition 3.5). This is a conservative choice that allows us to extend existing bottom-up quantification methods based on the refinement relation of the tree (Kordy et al., 2014a; Mauw and Oostdijk, 2006). In fact, it follows from Proposition 3.7 that, if all labels in a tree are unique, then the resulting attack-tree decoration problem based on this bottom-up attribute constraint-set is either undetermined or determined. We remark, nonetheless, that our methodology can also be used to model other computational approaches such as the Bayesian reasoning proposed by Kordy et al. (2014c). It is ultimately the analyst who decides what constitutes a set of hard predicates, although we require the analyst to come up with hard predicates that are satisfiable; as we do in this article.

4.2. Soft predicates

Statistical data and constraints extracted from industry-relevant knowledge-bases and experts are too valuable to ignore. And our methodology treats them as first-class citizens. As usual, we encode this information in predicates. For example, assume that comprehensive empirical data indicates that the probability of a bank account being hacked is less than 0.01. The semantics of such attribute in our running example tree can be defined by a set containing the predicate $hack-prob(hack-account) := hack-account \leq 0.01$.

In our methodology, predicates obtained from experts and knowledge-bases are regarded as soft. The reason is that, when it comes to opinion and empirical data, inconsistencies are common. Hence we do allow these predicates to be violated up to some extent. For example, consider that for a particular attack tree we obtain that the probability that an account is hacked is 0.02. Although such an outcome violates the predicate $\text{hack-prob}(\text{hack-account})$, one may find it acceptable and not far from the considered empirical data.

4.3. Analysis of attribute semantics with hard and soft predicates

Given an attack tree t and attribute constraint-set $\{p_1, \dots, p_n\}$ over labels of t and domain D , we use $H(t)$ and $S(t)$ to denote the partition of $\{p_1, \dots, p_n\}$ into hard and soft predicates, respectively. As described in the previous section, we analyse an attribute constraint-set by looking at solutions of the corresponding constraint satisfaction problem. Formally, given an attribute constraint-set $H(t) \cup S(t)$, we aim at finding a valuation $\alpha \in \text{VAL}_{\text{labels}(t) \rightarrow D}$ such that $H(t) \cup S(t) \models \alpha$. However, such formulation makes no distinction between hard and soft predicates, which is a feature we regard important in our methodology. For example, it may be the case that no valuation α satisfying $H(t) \cup S(t) \models \alpha$ exists, while we can still find $\alpha' \in \text{VAL}_{\text{labels}(t) \rightarrow D}$ such that $H(t) \models \alpha'$. Note that, although the constraint-set $H(t)$ is an oversimplification of the original attribute constraint-set with all soft constraints being removed, α' satisfies all hard constraints and thus may be worth considering.

In our methodology, when the original attack-tree decoration problem has no solution we propose to solve a weaker version: *the relaxed attack-tree decoration problem*. This new problem allows soft predicates to be weakened, which consists in replacing any soft predicate $p \in S(t)$ by a predicate p' that logically follows from p . We define this type of entailment on predicates over an attack tree t by:

$$p \Rightarrow p' \text{ if and only if } \forall \alpha \in \text{VAL}_{\text{labels}(t) \rightarrow D} : p \models \alpha \Rightarrow p' \models \alpha.$$

Using this notation, we can define the notion of a weakening relation on sets of predicates.

Definition 4.1 (Weakening relation). Let \leq be a partial order on sets of predicates. Then we say that \leq is a *weakening relation* if and only if for all sets of predicates P and P' it holds that

$$P' \leq P \Rightarrow (\forall p' \in P' \exists p \in P : p \Rightarrow p').$$

If $P' \leq P$, we say that P' is a weakening of P under the weakening relation \leq . We provide three examples of weakening relations.

1. Set equality ($=$), which is the trivial weakening relation.
2. Set inclusion (\subseteq), which allows one to weaken a set of predicates by deleting one or more of its elements.
3. The maximal weakening relation (\leq_M), which is defined by

$$P' \leq_M P \iff (\forall p' \in P' \exists p \in P : p \Rightarrow p').$$

The proofs that these are indeed weakening relations and that \leq_M is maximal are straightforward.

Using this notion of a weakening relation we reformulate the attack-tree decoration problem as an optimisation problem in the following way.

Definition 4.2 (The relaxed attack-tree decoration problem). Let t be an attack tree and $H(t) \cup S(t)$ an attribute constraint-set over labels of t and domain D , where $H(t)$ and $S(t)$ are hard and soft predicates, respectively. Let \leq be a weakening relation. The *relaxed attack-tree decoration problem* consists of two stages:

1. Finding a set of predicates S over $\text{labels}(t)$ and domain D such that:
 - $S \leq S(t)$,
 - $\llbracket H(t) \cup S \rrbracket_t \neq \emptyset$, and
 - $\forall S' : S \leq S' \leq S(t) \Rightarrow (\llbracket H(t) \cup S' \rrbracket_t = \emptyset \vee S' \leq S)$.
2. Solving the attack tree decoration problem with constraint-set $H(t) \cup S$.

A solution is a pair (S, α) , such that $\alpha \in \llbracket H(t) \cup S \rrbracket_t$.

The choice of the weakening relation is relevant in an instantiation of the relaxed attack tree decoration problem, as we show in the next section. In particular, we analyse two relevant decoration problems resulting from two concrete weakening relation definitions, namely the set inclusion (\subseteq) and maximal weakening relation (\leq_M).

5. Decoration algorithm for specific classes of predicates

Solving a constraint satisfaction problem is in general NP-hard. Thus, this section is devoted to instantiating each component of the developed theory into concrete predicate languages that can be used in standard solver tools to find solutions for the relaxed attack-tree decoration problem. In Section 7 below we show how those instantiations of the theory can be used to analyse a comprehensive attack tree case study.

5.1. Maximal weakening over inequality relations

Here we address the question of whether there exists a meaningful predicate language and constraint satisfaction solver that can be used to solve the relaxed attack tree decoration problem with respect to the maximal weakening relation. Note that, among all possible weakening relations the maximal one is the less restrictive. Hence it leads to more fine-grained solutions than other weakening relations.

The chosen predicate language defines predicates of three types, all based on comparing one or two labels to a constant value. The three types of predicates are:

1. $\ell \leq a$
2. $\ell \geq a$
3. $\ell \leq \ell' + a$

where ℓ and ℓ' are labels and a is a real number (positive or negative). We denote the set of all such predicates by $\mathcal{P}_{\text{ineq}}$ and we will often use $s_a \in \mathcal{P}_{\text{ineq}}$ to denote a predicate in this set with constant value a . It is easy to verify that $\{p\} \leq_M \{p'\}$ can only hold if predicates p and p' are of the same type.

Lemma 5.1. Let $p, p' \in \mathcal{P}_{\text{ineq}}$ be two predicates, then $\{p\} \leq_M \{p'\}$ implies that, for some labels ℓ_1, ℓ_2 , and some real numbers a, a' ,

$$\begin{aligned} p \equiv \ell_1 \leq a & \quad \wedge \quad p' \equiv \ell_1 \leq a', \text{ or} \\ p \equiv \ell_1 \geq a & \quad \wedge \quad p' \equiv \ell_1 \geq a', \text{ or} \\ p \equiv \ell_1 \leq \ell_2 + a & \quad \wedge \quad p' \equiv \ell_1 \leq \ell_2 + a'. \end{aligned}$$

The three types of predicates have been chosen in such a way that the maximal weakening relation \leq_M on single predicates can be easily characterised by the numerical order of their constant values a . This characterisation will allow us later to define the distance between two predicates as the difference between their constant values a and a' .

Lemma 5.2. Let ℓ_1, ℓ_2 be labels and a, a' be real numbers. Then the following properties hold.

$$\begin{aligned} \{\ell_1 \leq a\} \leq_M \{\ell_1 \leq a'\} & \iff a \geq a' \\ \{\ell_1 \geq a\} \leq_M \{\ell_1 \geq a'\} & \iff a \leq a' \\ \{\ell_1 \leq \ell_2 + a\} \leq_M \{\ell_1 \leq \ell_2 + a'\} & \iff a \geq a' \end{aligned}$$

From this lemma it follows, for instance, that for $S, S' \subseteq \mathcal{P}_{ineq}$, if $S \preceq_M S'$ and $(\ell_1 \leq a) \in S$, then there exists $(\ell_1 \leq a') \in S'$, such that $a \geq a'$. Hence we consider the set $\mathcal{F}(S, S')$ containing all total functions $f: S \rightarrow S'$ such that $\forall p \in S: f(p) \Rightarrow p$. The Euclidean distance between two predicate $s_a, s_{a'} \in \mathcal{P}_{ineq}$ is given by $d(s_a, s_{a'}) = |a - a'|$ if s_a and $s_{a'}$ are of the same type, $d(s_a, s_{a'}) = \infty$ otherwise. Given $f \in \mathcal{F}(S, S')$, we define $d_f(S, S') = \sqrt{\sum_{p \in S} d^2(p, f(p))}$, and the Euclidean distance $d(S, S')$ between two sets of predicates by

$$d(S, S') = \begin{cases} \infty, & \text{if } \mathcal{F}(S, S') = \emptyset \\ \min_{f \in \mathcal{F}(S, S')} d_f(S, S'), & \text{otherwise} \end{cases}$$

We restrict the distance measure above to bijective functions only. That is to say, we consider from now on $\mathcal{F}(S, S')$ to be the set containing all *bijective* functions $f: S \rightarrow S'$ such that $\forall p \in S: f(p) \Rightarrow p$. A consequence of such restriction is that predicate sets with different cardinality have distance ∞ , which simplifies the proof of [Theorem 5.3](#) below.

Next we provide sufficient conditions for a set of predicates to be part of a solution of the relaxed attack-tree decoration problem. It states that a set S , which minimizes its distance $\hat{d}(S, S(t))$ to $S(t)$, where $S(t)$ is the set of soft predicates for a given tree t , and that satisfies $\llbracket H(t) \cup S \rrbracket_t \neq \emptyset$, where $H(t)$ is the set of hard predicates, leads to a solution of the relaxed attack-tree decoration problem.

Theorem 5.3. *Let t be an attack tree and $H(t) \cup S(t)$ an attribute constraint-set for t , where $S(t) \subseteq \mathcal{P}_{ineq}$. Let $S \subseteq \mathcal{P}_{ineq}$ be a set of predicates such that $\llbracket H(t) \cup S \rrbracket_t \neq \emptyset$ and $\hat{d}(S, S(t))$ is minimum and defined, i.e. $\hat{d}(S, S(t)) \neq \infty$. Then there exists a valuation $\alpha \in \llbracket H(t) \cup S \rrbracket_t$ such that (S, α) is a solution of the relaxed attack-tree decoration problem with respect to \preceq_M .*

Proof. The interested reader can find the proof in [Appendix A](#). \square

This theorem makes it possible to reduce the relaxed attack-tree decoration problem for the maximal weakening relation \preceq_M on the given types of predicates to an optimisation problem that we solve via nonlinear programming. The formulation of the optimisation problem is given below.

Corollary 5.4. *Let t be an attack tree and $H(t) \cup S(t)$ an attribute constraint-set for t where $S(t) = \{p_1, \dots, p_n\}$ contains only predicates from \mathcal{P}_{ineq} . We create the set of predicates S by replacing every predicate $p_i \in S(t)$ by.*

$$\begin{aligned} p_\alpha &\equiv \ell \leq x_i & \text{if } p &\equiv \ell \leq a \\ p_\alpha &\equiv \ell \geq x_i & \text{if } p &\equiv \ell \geq a \\ p_\alpha &\equiv \ell \leq \ell' + x_i & \text{if } p &\equiv \ell \leq \ell' + a \end{aligned}$$

where x_1, \dots, x_i are variables. A solution α, x_1, \dots, x_n to the following optimisation problem leads to a solution (S, α) of the relaxed attack tree problem.

$$\begin{array}{lll} \text{minimize} & d(S, S(t)) & \text{subject to} \quad \alpha \in \llbracket H(t) \cup S \rrbracket_t. \\ & \alpha, x_1, \dots, x_n \end{array}$$

Implementation. To find a valuation function α and a set of weakening predicates S' so that the distance function $d(S, S(t))$ is minimised, we relied on the Sequential Quadratic Programming (SQP) problem interpretation of the relaxed attack-tree decoration problem. We further refer to this tool as the SQP-based tool. It is implemented using the Python `scipy` library that provides the Sequential Least Squares Programming (SLSQP) algorithm for solving optimisation problems of this type. Our implementation³ does not impose any burden on the analyst, as it allows a loose interpretation of all constraints together. In case the set of constraints is

satisfiable, our tool finds an optimal solution. In case of an unsatisfiable set of constraints, our implementation will find an optimal solution that minimises the distance function between sets of predicates.

5.2. Set inclusion weakening over propositional logic

As the analyst may require a predicate language richer than the one described above, we provide tool support for predicates written in the propositional logic. We do so via a transformation of a relaxed attack tree-decoration problem instance into a Satisfiability Modulo Theories (SMT) instance ([de Moura and Bjørner, 2008](#)). SMT is the problem of determining whether a formula in the first-order logic, where some operator symbols are provided with a theory, is satisfiable.

Given a set of predicates S , we use $\Gamma(S)$ to denote the first-order logic formula formed by all predicates in S in the conjunctive form. Then the SMT instance resulting from the relaxed attack-tree decoration problem instance with the attribute constraint-set $H(t) \cup S(t)$ is defined by $\Gamma(H(t) \cup S(t))$. If $\Gamma(H(t) \cup S(t))$ is satisfiable, it follows that the decoration problem does not need to be relaxed. Otherwise, we use [Algorithm 1](#) to find a subset of soft predicates $S \subseteq S(t)$ that solves the relaxed attack tree decoration problem with respect to the inclusion relation.

Algorithm 1 Solving the decoration problem w.r.t. the set inclusion weakening relation.

Require: The relaxed attack-tree decoration problem defined by attack tree t , attribute constraint-set $H(t) \cup S(t)$.

Ensure: S is a set of maximum cardinality such that $S \subseteq S(t)$ and $\llbracket H(t) \cup S \rrbracket_t \neq \emptyset$.

- 1: Let $S(t) = \{p_1, \dots, p_n\}$ be the soft constraint set.
- 2: $S \leftarrow \emptyset$
- 3: **for** $i = 1..n$ **do**
- 4: $S \leftarrow S \cup \{p_i\}$
- 5: **if** $\Gamma(H(t) \cup S)$ is not satisfiable **then**
- 6: $S \leftarrow S \setminus \{p_i\}$
- 7: **return** (S, α) where $\alpha \in \llbracket H(t) \cup S \rrbracket_t$.

Initially, the set S is empty. [Algorithm 1](#) works by iteratively adding predicates to the set S , until the formula $\Gamma(H(t) \cup S)$ is satisfiable, while the formula $\Gamma(H(t) \cup S')$ is unsatisfiable for any $S' \supset S$. It is easy to prove that such procedure provides a solution to the relaxed attack-tree decoration problem with respect to the subset inclusion weakening relation, based on the assumption that hard predicates are satisfiable.

Implementation. We implemented our transformation relying on the well-known theorem prover Z3 from Microsoft⁴. Z3 can be utilised as a constraint solver, i.e., it can find a solution satisfying a set of constraints expressed as equalities and inequalities. Our prototype⁵ can handle all attribute domains on real numbers that are defined in the ADTool ([Kordy et al., 2013a](#)), e.g. probability, minimal cost of attack, and minimal time; and it is trivially extensible to attribute domains defined on Boolean values, e.g. satisfiability of a scenario. The analyst can further specify additional constraints, if desired. The resulting set of constraints is passed to the Z3 prover, which reports whether the problem is solvable or not. In case the constraint satisfaction problem is solvable, the prover will report a complete consistent valuation for the given tree. This valuation will satisfy the tree structure and the constraints expressed by the analyst, and it will agree with the given initial valuation. If the

⁴ <https://github.com/Z3Prover/z3>.

³ Code available at https://github.com/vilena/at-decorator/tree/master/SQP_decorator.

⁵ Code available at https://github.com/vilena/at-decorator/tree/master/CSP_decorator.

constraint satisfaction problem is not solvable (i.e. the initial valuation is inconsistent), we use [Algorithm 1](#) to find a subset of soft predicates of maximum cardinality that is satisfiable, and report a solution satisfying the found maximal predicate set.

6. ATM case study

We now proceed to show how our approach to attack-tree decoration can be applied to a real-life security scenario. The goal is to show that decoration can be performed in a systematic way even when the analyst only has partial information about attribute values.

In this section, we introduce a case study related to capturing automated teller machine fraud scenarios as an attack tree, and we describe the historical data available for decorating this tree.

6.1. ATM security: a case study

Automated teller machines (ATMs) are complex and expensive systems used daily by millions of bank customers worldwide. Because each carries a significant amount of cash, ATMs are the target of large-scale criminal actions. Only in 2015 more than 16,000 ATM incidents were reported in Europe, causing over 300 million Euros loss.⁶

In an attempt to provide structure to the risk assessment process and catalogue ATM threats, Fraile et al. created a comprehensive attack-defence tree capturing the most dangerous attack scenarios applicable to ATMs (Fraile et al., 2016). The tree modelled in Fraile et al. (2016) contains three main branches: brute-force attacks, fraud attacks, and logical attacks. The attacks of the logical type make use of malicious software, while a brute-force attack typically ends up destroying the ATM. Differently from these two attack scenarios, ATM fraud attacks involve conventional electronic devices (such as card skimmers) and require the participation of the victim.

In this empirical evaluation section we focus on ATM fraud, because we have more empirical data for these attacks than for the other types of attacks. [Fig. 3](#) presents an attack tree characterising such attacks that is loosely based on the attack-defence tree published by Fraile et al. (2016). In ATM fraud, criminals need covert access to the ATM, as this attack typically requires opening the machine's case either by force or with a generic key, and installing a special device (e.g. a skimmer). Then the attacker waits until a victim uses the ATM and, as a consequence, enables the installed device. Lastly, the attacker gets cash from the victim's account by means of various techniques, such as cash trapping, card cloning, etc.

6.2. Decorating the ATM fraud attack tree

The decoration process we propose in this paper consists of three independent steps that are executed for a given attack tree. First, an attribute is chosen. In this case study, we focus on *probability of success*, that is, the probability that a given ATM machine is used to successfully execute ATM fraud. The attack tree structure jointly with the attribute rules determine the hard constraint set, i.e. the standard bottom-up constraints derived from the attack tree structure. Second, statistical information (historical data) related to the chosen attribute is gathered. Such statistical values are used to provide tree nodes with probability values. For the ATM fraud scenario, the available statistical data is presented in [Section 6.2.1](#). Lastly, relations among nodes in the tree are established based on the analyst's insight and domain knowledge.

Table 1

Historical data values identified for some attack tree nodes from the ATM Crime report.

| Node | Prob. | Source |
|-----------------------|--------|------------------------------|
| <i>atm-fraud</i> | 0.0046 | ATM Crime Report 2015 (EAST) |
| <i>card-skimming</i> | 0.0172 | ATM Crime Report 2015 (EAST) |
| <i>card-trapping</i> | 0.0094 | ATM Crime Report 2015 (EAST) |
| <i>cash-trapping</i> | 0.0150 | ATM Crime Report 2015 (EAST) |
| <i>trans-reversal</i> | 0.0038 | ATM Crime Report 2015 (EAST) |

[Section 6.2.2](#) presents the corresponding analysis for the ATM fraud scenario. The full set of constraints for the ATM fraud scenario is summarised in [Section 6.3](#).

6.2.1. Statistical analysis

The statistical values we consider here have been derived from the ATM Crime Report 2015 (EAST). In our case, we analyse ATM fraud incidents in Lisbon, which hosts 300 ATMs. We remark, however, that these values have been derived for illustrative purposes only and may not be accurate.

Between 2010 and 2015, 83 ATM fraud attacks have been reported in Lisbon. This gives a 0.0461 probability of an ATM to be the target of fraud within a calendar year, if we assume the uniform distribution of these attacks. Because the report categorises ATM fraud into different attack types, we can provide probability values for some attack types by analysing the attack frequency as reported in the EAST report. The results can be found in [Table 1](#). For the results reported in this table, we assume that historical attacks were uniformly distributed, and we rely on frequencies of attacks over long time periods to estimate probabilities, like in the OCTAVE method (Caralli et al., 2007).

6.2.2. Domain knowledge constraints

In the previous subsection, we have shown how available statistical data can be used as a constraint in our decoration process. Another novelty of our approach is that we allow for domain knowledge constraints, that is, facts that must be additionally satisfied in the attack tree. The following list of predicates is based on historical observations from the previously mentioned ATM Crime Report 2015 and also on the European Central Bank report on card fraud (2015).

- *card-skimming* is more likely than *take-card-phys*. Moreover, *get-credentials* is more likely than *cash-trapping*, which is more likely than *trans-reversal*.
- *shoulder-surf* is more likely than *install-camera*.
- *install-camera*, *install-epp* and *install-skimmer* are all equally likely.
- *cash-trapping* and *card-trapping* are equally likely.

Generally, not all historical observations may be used as domain knowledge constraints. For example, even if some attacks were more rarely observed than others, they may become more frequent in the future if the miscreants adapt their strategies. Thus, experts must define domain knowledge constraints with caution. Yet, some observations still could be made based on the domain knowledge. For instance, one encryption system may be *known* to be easier to break than another one. Moreover, as our knowledge changes, the constraints could and should be revised.

6.3. Full set of predicates

We now list the full set of predicates that will be used by the tools to solve the decoration problem. To simplify the presentation,

⁶ <https://www.association-secure-transactions.eu/tag/atm-crime-report/>.

we use a short-hand notation. All predicates listed in this section can be straightforwardly transformed into our predicate notation.

Hard predicates. Considering the attribute domain of probability of success, the attack tree shown in Fig. 3 corresponds to the following set of hard predicates:

- $atm-fraud = \times (access-atm, execute-attack)$,
- $access-atm = + (break-in, social-engineer-staff)$,
- $execute-attack = + (trans-reversal, get-credentials, cash-trapping)$,
- $get-credentials = \times (get-pin, get-card)$,
- $get-pin = + (shoulder-surf, install-camera, install-epp)$,
- $get-card = + (card-skimming, take-card-phys, social-engineer-owner)$,
- $card-skimming = \times (install-skimmer, clone-card)$,
- $take-card-phys = + (card-trapping, steal-card)$.

Soft predicates. Historical data values from Table 1 are encoded in the form of soft predicates:

- $atm-fraud = 0.0046$,
- $card-skimming = 0.0172$,
- $card-trapping = 0.0094$,
- $cash-trapping = 0.0150$,
- $trans-reversal = 0.0038$.

We will subsequently refer to the soft predicates listed above as *historical data constraints*.

Domain knowledge from the ATM Crime report is encoded in the form of soft predicates as well:

- $take-card-phys \leq card-skimming$,
- $cash-trapping \leq get-credentials$,
- $trans-reversal \leq cash-trapping$,
- $install-camera \leq shoulder-surf$,
- $install-camera = install-epp$,
- $install-skimmer = install-epp$,
- $install-skimmer = install-camera$,
- $cash-trapping = card-trapping$,

Subsequently, we will refer to the set of soft predicates above as *domain knowledge predicates*.

6.4. Goals of the analysis

We consider that the analyst has designed an attack tree covering ATM fraud scenarios as presented in Fig. 3. This attack tree gives them the set of hard constraints given in Section 6.3. Furthermore, the analyst has elicited a set of soft constraints based on their knowledge of the problem space and the information available in the ATM Crime Report (also listed in Section 6.3). However, the analyst is not able to find enough data to estimate probabilities for all leaf nodes in the attack tree, what prevents them from straightforwardly applying the bottom-up evaluation procedure to compute the probabilities for all intermediate nodes and, ultimately, for the root node.

The analyst can, however, apply our methodology and decorate the attack tree. We consider the following possible analysis questions that can be investigated with our approach:

- Are the attack tree and the set of constraints elicited by the analyst compatible? I.e. does the corresponding decoration problem have a solution? In our notation, for an attack tree t , and attribute constraint-set $H(t) \cup S(t)$, is $\llbracket H(t) \cup S(t) \rrbracket_t \neq \emptyset$?
- What is a solution for the given decoration problem? In our notation, the analyst is interested in finding a solution $\alpha \in \llbracket H(t) \cup S(t) \rrbracket_t$.

- If the decoration problem has no solution, what is a solution that is the closest to satisfying all constraints? This question corresponds to solving the relaxed attack tree decoration problem formulated in Definition 4.2, for a chosen weakening relation.

We will demonstrate in the next section how our two implementations solve the relaxed attack tree decoration problem of the ATM case study, for the maximal and set inclusion weakening relations.

Note that one may argue that in our case study the analyst already has the probability of ATM fraud (the root node) from the historical data, and therefore, they can skip decorating the whole tree. However, the analyst may still want to perform a so-called *what-if* analysis, which consists in analyzing different but related scenarios. For example, the analyst could answer questions such as: What if the probability of this attack is in fact higher than I envisage? How will this affect my security posture? The *what-if* analysis requires a fully annotated tree, which can be provided using our decoration technique even over partially available data. Furthermore, in general, it cannot be assumed that the root node value will always be available from the historical data.

7. Empirical evaluation results

We now show how our two implementations can be applied to analyse the ATM case study introduced previously.

7.1. The CSP-based implementation showcase

We first exemplify the results of the CSP-based implementation on the ATM case study tree presented in Fig. 3. The first solution in Table 2 presents a possible valuation for this attack tree found by the CSP-based tool with the hard predicates and the historical data predicates listed in Section 6.3. Fig. B.4 in Appendix B visualises this solution in the ATM fraud attack tree. We have used the open source ADTool software (Gadyatskaya et al., 2016c; Kordy et al., 2013a) for visualising attack trees.

If the analyst introduces the domain knowledge predicates as an additional set of constraints, the CSP tool indicates that the constraint satisfaction problem becomes unsatisfiable. Indeed, the constraints on *card-trapping* and *cash-trapping* are contradictory, because historical data does not indicate that the probabilities of these attacks are exactly equal.

The Z3 solver that we use as the underlying constraint satisfaction engine is capable of finding an unsatisfiable core of the problem. In the ATM fraud case, the solver reports that three predicates constitute the unsatisfiable core:

- $card-trapping = cash-trapping$,
- $card-trapping = 0.0094$,
- $cash-trapping = 0.015$.

The solver then proceeds to identify a maximal set of predicates that is still satisfiable. In our case, the solver is able to satisfy all predicates but $cash-trapping = 0.015$. The second solution in Table 2, illustrated in Fig. B.5 (Appendix B), shows a possible valuation found by the CSP-based tool after adding the set of domain knowledge constraints to the original assignment predicate set (i.e. excluding the constraint on equality of *cash-trapping* to 0.015).

Analysing this valuation, we can see that, for example, the probability of *take-card-phys* has dropped significantly, and the proba-

Table 2
Solutions of the ATM fraud attack tree decoration problem found by our tools.

| Node Label | Probability of attack success for tool and set of predicates | | |
|------------------------------|--|---|---|
| | CSP hard predicates + historical data | CSP hard predicates + historical data + domain knowledge | SQP hard predicates + historical data + domain knowledge |
| <i>atm-fraud</i> | 0.0046 | 0.0046 | 0.0046 |
| <i>access-atm</i> | 0.0068 | 0.0093 | 0.0184 |
| <i>break-in</i> | 0.0039 | 0.0078 | 0.0092 |
| <i>social-engineer-staff</i> | 0.0029 | 0.0015 | 0.0092 |
| <i>execute-attack</i> | 0.6683 | 0.4914 | 0.2493 |
| <i>trans-reversal</i> | 0.0038 | 0.0038 | 0.0038 |
| <i>get-credentials</i> | 0.6620 | 0.4847 | 0.2324 |
| <i>get-pin</i> | 0.875 | 0.9375 | 0.5780 |
| <i>shoulder-surf</i> | 0.5 | 0.75 | 0.3834 |
| <i>install-camera</i> | 0.5 | 0.5 | 0.0973 |
| <i>install-epp</i> | 0.5 | 0.5 | 0.0973 |
| <i>get-card</i> | 0.7566 | 0.5170 | 0.4021 |
| <i>card-skimming</i> | 0.0172 | 0.0172 | 0.0172 |
| <i>install-skimmer</i> | 0.5 | 0.5 | 0.0973 |
| <i>clone-card</i> | 0.0344 | 0.0344 | 0.1768 |
| <i>take-card-phys</i> | 0.5047 | 0.0171 | 0.0172 |
| <i>card-trapping</i> | 0.0094 | 0.0094 | 0.0113 |
| <i>steal-card</i> | 0.5 | 0.0078 | 0.0059 |
| <i>social-engineer-owner</i> | 0.5 | 0.5 | 0.3677 |
| <i>cash-trapping</i> | 0.015 | 0.0094 | 0.0131 |

bility of *steal-card* has decreased with the new constraint set. Indeed, in the first assignment set there was no constraint that the probability of *take-card-phys* is less than the probability of *card-skimming* (which is statistically a rare event itself). When we added this constraint found in the domain knowledge catalogue, the valuation for this subtree has changed. This result demonstrates that even if the analyst cannot obtain a valuation for some attack such as *steal-card*, they can get an estimation for it by using other known parameters and domain knowledge.

7.2. The SQP-based implementation showcase

When running the relaxed tree decoration method considering the historical data values in Table 1 and the domain knowledge constraints from Section 6.2.2, the optimisation task converged to the solution presented in the third solution of Table 2. Fig. B.6 in Appendix B presents the decorated attack tree.

Three of the soft constraints stemming from the historical data values (*atm-fraud*, *card-skimming*, *trans-reversal*) were satisfied precisely without weakening. *cash-trapping* and *card-trapping* constraints were satisfied by weakened predicates. The probability of *card-trapping* in historical data was 0.0094, and the optimised value is 0.0113. The probability of *cash-trapping* was 0.015 in historical data, and the optimised value is 0.0131.

Six out of seven domain knowledge constraints were satisfied precisely without weakening. The remaining constraint that was not satisfied precisely is the knowledge that *cash-trapping* and *card-trapping* are equally likely. In the optimised solution they have a small difference, the probability of *cash-trapping* is 0.0131 while the probability of *card-trapping* is 0.0113.

In the result of this process, four weakened predicates were found. The soft and weakened predicates for these four cases are shown in Table 3. It is the task of the analysts to decide on whether such weakened predicates fit their analysis. For example, except for the *take-card-phys* predicate, all other predicates in Table 3 satisfy that their distance to their weakening predicate is

lower than 10^{-2} , and one may argue that in the probability domain a 10^{-2} discrepancy is acceptable.

7.3. Discussion

We have showcased the application of our methodology on a practical scenario of ATM security. We have established in this empirical validation that the attack-tree decoration methodology is versatile, as it can be realised differently depending on the analyst's needs. In particular, the analyst can choose a suitable weakening relation and ask the solver to tackle the decoration problem depending on their confidence in the tree structure and the historical data available.

When the confidence in the predicate set is high, the analyst will work with the relaxed attack-tree decoration problem, e.g. using our CSP and SQP-based tools, to obtain a solution that is the closest one to satisfying all available data points. When the analyst has low confidence in the data, they can work with a constraint solver, for example, starting from our CSP-based implementation, to identify whether there is inconsistency in the data. If the constraint satisfaction problem is not satisfiable, the reported unsatisfiable core of the predicates can be the first candidate to be reviewed with other experts in order to revise the corresponding data points.

The limitation of our methodology is that it does not allow to draw precise conclusions about data accuracy. For example, an unsatisfiable core reported by the CSP-based implementation in case of inconsistency does not guarantee that this is indeed the set of wrong predicates. It might be that the real problem lies within another, co-dependent set of predicates, where considered data values are inaccurate. We plan to investigate the means to evaluate the decoration accuracy considering e.g. analyst's confidence level for each data point and the size of the solution space in future work.

Our methodology does not fully remove the intellectual burden from the analyst, but it equips them with an insight into conformity of the historical data with the given attack tree structure,

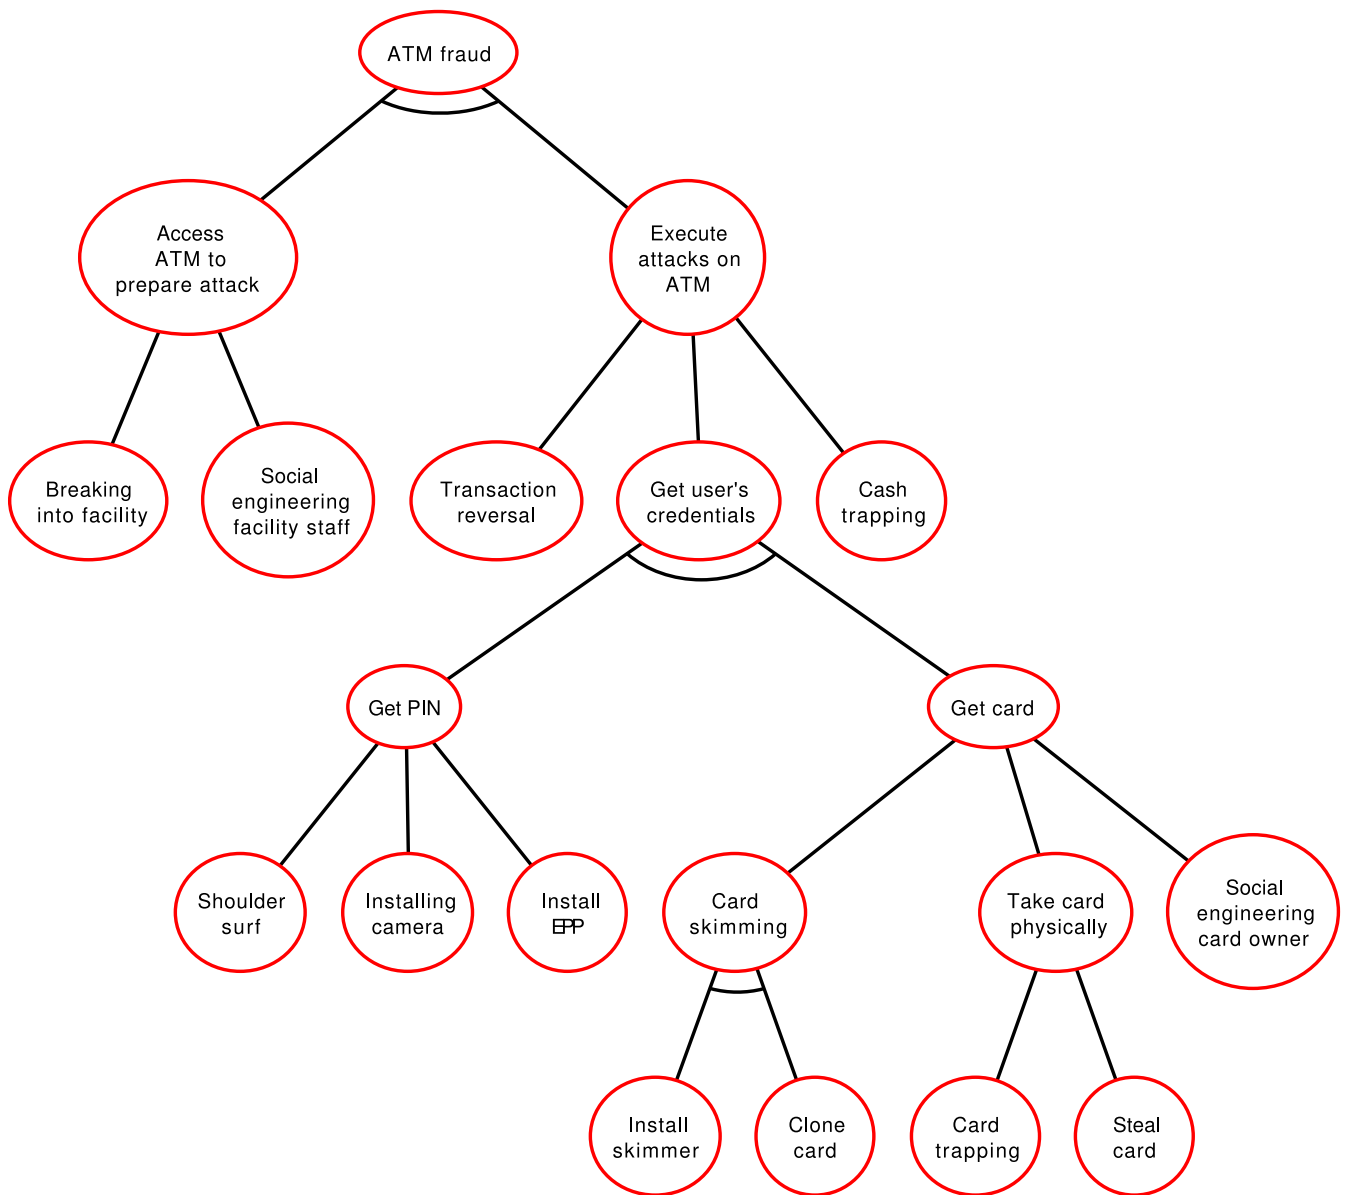


Fig. 3. An attack tree modelling ATM fraud. The tree is loosely based on the attack-defence tree published by Fraile et al. (2016).

Table 3

Weakened predicates found by the optimisation process.

| Soft predicate | Weakened predicate |
|---|--|
| $card-trapping \leq 0.0094$ | $card-trapping \leq 0.0113$ |
| $cash-trapping \leq 0.015$ | $cash-trapping \leq 0.0131$ |
| $take-card-phys \leq card-skimming + 0$ | $take-card-phys \leq card-skimming + 0.0172$ |
| $cash-trapping \leq card-trapping + 0$ | $cash-trapping \leq card-trapping + 0.0018$ |

and with an initial, complete attack-tree decoration. This decoration can be further improved by the analyst by incorporating more historical data or engaging more domain experts whenever needed.

Our two prototypes demonstrate that the methodology is practical and it works on real attack trees. In general, the non-linear constraint satisfaction problem used by the CSP-based implementation and the SQP problem are NP-hard (Manyem and Ugon, 2012; Tsang, 1993). Note that the complexity of the attack-tree

decoration task itself is unknown. Therefore, our implementations will not scale to very large attack trees. To the best of our knowledge, there are no statistics on the average or maximal sizes of attack trees in practice. However, manually designed attack trees, in our own experience, rarely have more than 100 nodes, because they quickly become incomprehensible for humans (Gadyatskaya and Trujillo-Rasua, 2017). Therefore, we expect that our implementations will work reasonably well with the

majority of attack trees created by practitioners, as this is our main target.

Still, manually designed attack trees can be very diverse, and it might turn out that no historical data is available at an organisation for the given tree. Then the analysts will fall back on the traditional attack tree application methodology, when experts assign values to the leaf nodes based on their experience and expectations. The benefit of our approach is to better integrate the attack trees and the *available* historical data, allowing to draw mutual insights from historical observations to envisaged threats, and vice-versa.

Finally, our methodology is showcased with the probability attribute, as we have a suitable collection of historical observations for the ATM fraud scenario. Probability or likelihood of a given attack is one of the most important attributes in the context of security risk assessment (Vose, 2008). However, our approach is applicable to other quantitative attributes defined for attack trees, e.g., cost, time, etc. (Bagnato et al., 2012; Kordy et al., 2013b).

8. Conclusions

In this article we proposed the first quantitative attribute approach for attack trees that is able to deal with incomplete information. On the one hand, we have well-known classical attributes on attack trees expressing relations between nodes in a tree. On the other hand, we recognise that reliable information can be obtained from historical data and domain knowledge. This type of information is typically not included in the semantics of attack trees, and so it has been largely ignored until now. We take these two views on the world and we verify that they are consistent. Since Schneier's definition of attack trees, bottom-up evaluation of attributes was the norm. We are the first to introduce a complete view on attribute values, including missing values and approximation.

The main benefit of our computational methodology is that it allows to obtain a consistent valuation for all attack tree nodes, even if some leaf nodes data is missing. This is not possible with the standard bottom-up decoration approaches. We have shown that the distinction between hard and soft constraints can be handy, as it allows the analyst to better understand the solution space and to interactively engage with the decoration problem. Lastly, we provide two implementations and we demonstrate the feasibility of the suggested approach on a case study. Our proof-of-concept implementations demonstrate the viability of the method, and they can be easily introduced into the established tools working with attack trees, such as the SecurITree tool and the ADTool.

Declaration of Competing Interest

None.

Acknowledgements

The research leading to these results has received funding from the [European Union Seventh Framework Programme](#)

under grant agreement number 318003 (TRESPASS), and from the [Fonds National de la Recherche Luxembourg](#) under grant C13/IS/5809105(ADT2P).

Appendix A. Proof of Theorem 5.3

Theorem Let t be an attack tree and $H(t) \cup S(t)$ an attribute constraint-set for t , where $S(t) \subseteq \mathcal{P}_{ineq}$. Let $S \subseteq \mathcal{P}_{ineq}$ be a set of predicates such that $\llbracket H(t) \cup S \rrbracket_t \neq \emptyset$ and $\hat{d}(S, S(t))$ is minimum and defined, i.e. $\hat{d}(S, S(t)) \neq \infty$. Then there exists a valuation $\alpha \in \llbracket H(t) \cup S \rrbracket_t$ such that (S, α) is a solution of the relaxed attack-tree decoration problem with respect to \preceq_M .

Proof. First, if $\llbracket H(t) \cup S(t) \rrbracket_t \neq \emptyset$, then there exists $\alpha \in \llbracket H(t) \cup S(t) \rrbracket_t$ and $\hat{d}(S(t), S(t)) = 0$, which is minimum. Thus in the remainder of the proof we assume that $\llbracket H(t) \cup S(t) \rrbracket_t = \emptyset$.

Now, notice that $S(t) \not\preceq_M S$, otherwise $\llbracket H(t) \cup S \rrbracket_t = \emptyset$. Thus we obtain that $S \preceq_M S(t)$, implying that S satisfies the first condition of the relaxed attack tree decoration problem. Moreover, we also conclude that $|S| = |S(t)|$, given that $\hat{d}(S, S(t)) \neq \infty$. Next, we will show that S satisfies the third condition of the problem definition as well.

Suppose we have $S \preceq_M S' \preceq_M S(t)$, but $S' \not\preceq_M S$. Because $S' \not\preceq_M S$, there must exist $p' \in S'$ such that no $p \in S$ satisfies that $p \Rightarrow p'$. Let us analyse the three possible predicate types of p' .

1. Assume $p' \equiv \ell \leq a'$. Because $S' \preceq_M S(t)$, there must exist $p'' \equiv \ell \leq a''$ in $S(t)$ with $a'' \leq a'$. Now, let $f: S \rightarrow S(t)$ be a bijective function such that $d_f(S, S(t)) = \hat{d}(S, S(t))$. Such a function exists given that $\hat{d}(S, S(t)) \neq \infty$ and $S \preceq_M S(t)$. Let $p \equiv \ell \leq a$ be the predicate in S such that $f(p) = p''$. Overall we obtain that $a'' \leq a'$ and $a'' \leq a$. Now, given that p does not imply p' , according to Lemma 5.2 it must be the case that $a > a'$. Therefore we obtain the order $a > a' \geq a''$. Consider the set of predicates $S'' = S \setminus \{p\} \cup \{p'\}$. On the one hand, because $a' \geq a''$ it follows that $S'' \preceq_M S(t)$ and $f \in \mathcal{F}(S'', S(t))$. On other hand, because $a > a'$ and $a' \geq a''$, we obtain that $d_f(S, S(t)) > d_f(S'', S(t))$. Considering that $d_f(S, S(t)) = \hat{d}(S, S(t))$, then $\hat{d}(S, S(t)) > d_f(S'', S(t)) \geq \hat{d}(S'', S(t))$, which contradicts the assumption that $\hat{d}(S, S(t))$ is minimum.
2. The case $p' \equiv \ell = \ell' \leq a'$ is analogous to the previous one.
3. Finally assume $p' \equiv \ell \geq a'$. As in the first case we obtain that there exist predicates $p \equiv \ell \geq a$ and $p'' \equiv \ell \geq a''$ in S and $S(t)$, respectively, such that $p'' \Rightarrow p'$ and $f(p) = p''$. Given that p does not imply p' , then we obtain the following order, $a'' \geq a' > a$. Again, such an order implies that $d(p, p'') > d(p', p'')$. Therefore, the set of predicates $S'' = S \setminus \{p\} \cup \{p'\}$ satisfies that $\hat{d}(S, S(t)) > \hat{d}(S'', S(t))$, which contradicts the assumption that $\hat{d}(S, S(t))$ is minimum.

The proof concludes by remarking that S satisfies the second condition of the relaxed attack tree decoration problem, as stated in the body of the theorem. \square

Appendix B. Decorated attack trees for the ATM case study

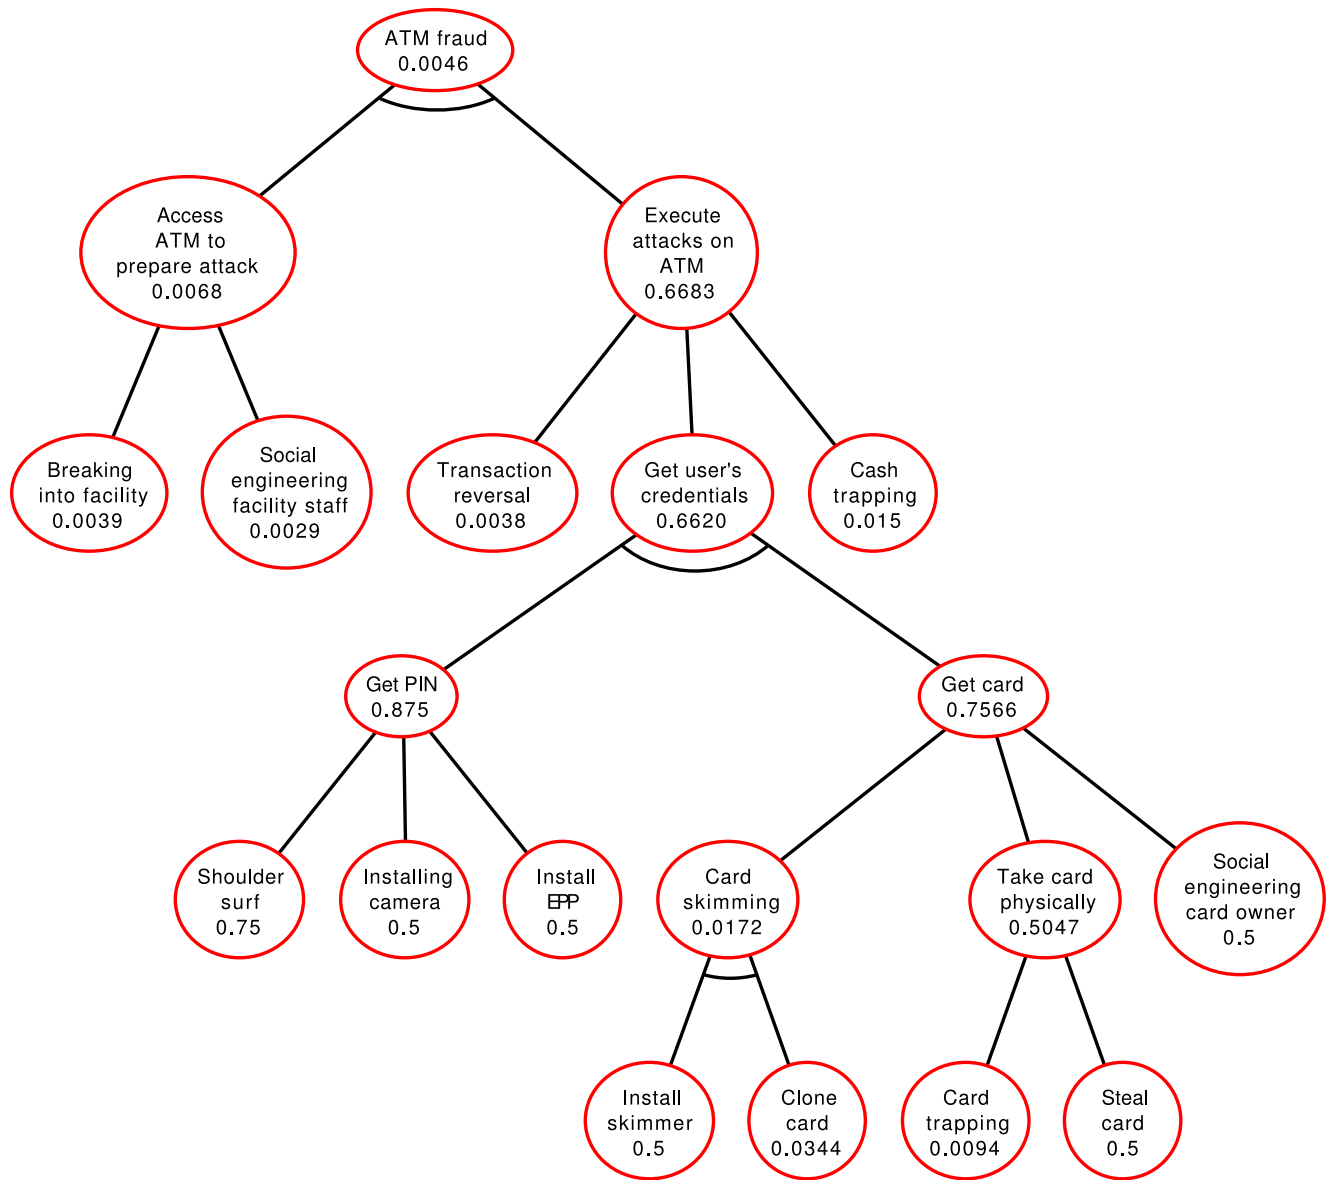


Fig. B.4. A valuation identified by the CSP-based tool for the attack tree in Fig. 3 with the hard constraints and the historical data predicates used as soft constraints (predicates are listed in Section 6.3).

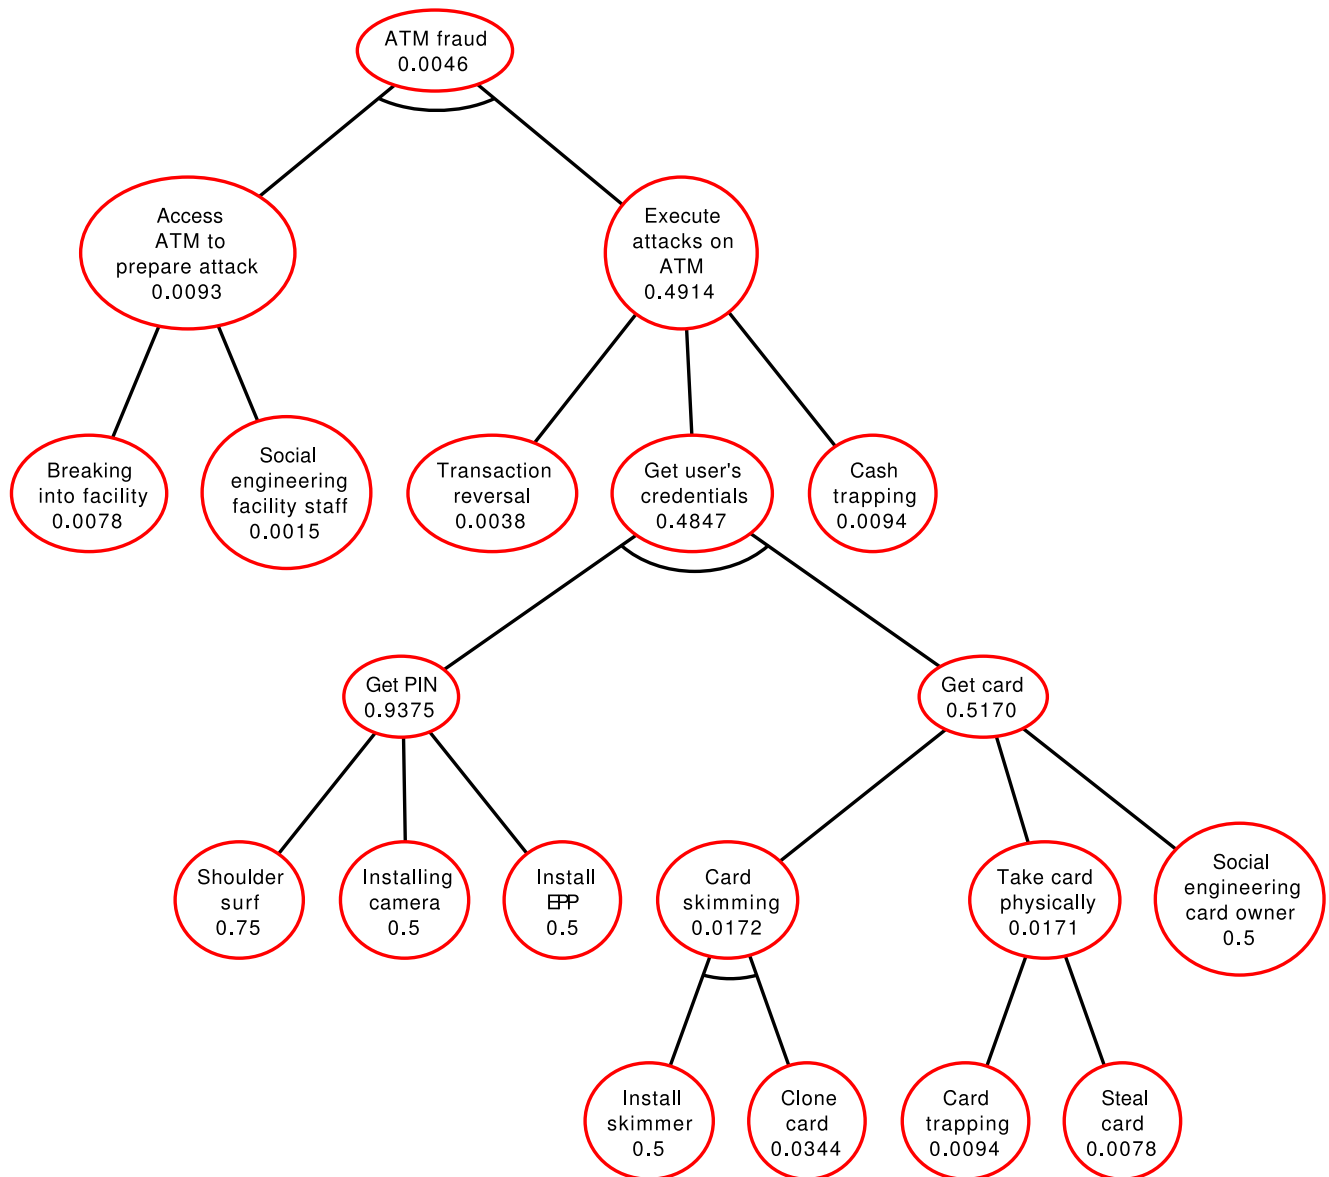


Fig. B.5. A valuation identified by the CSP-based tool for the attack tree in Fig. 3 with the predicates listed in Section 6.2.2.

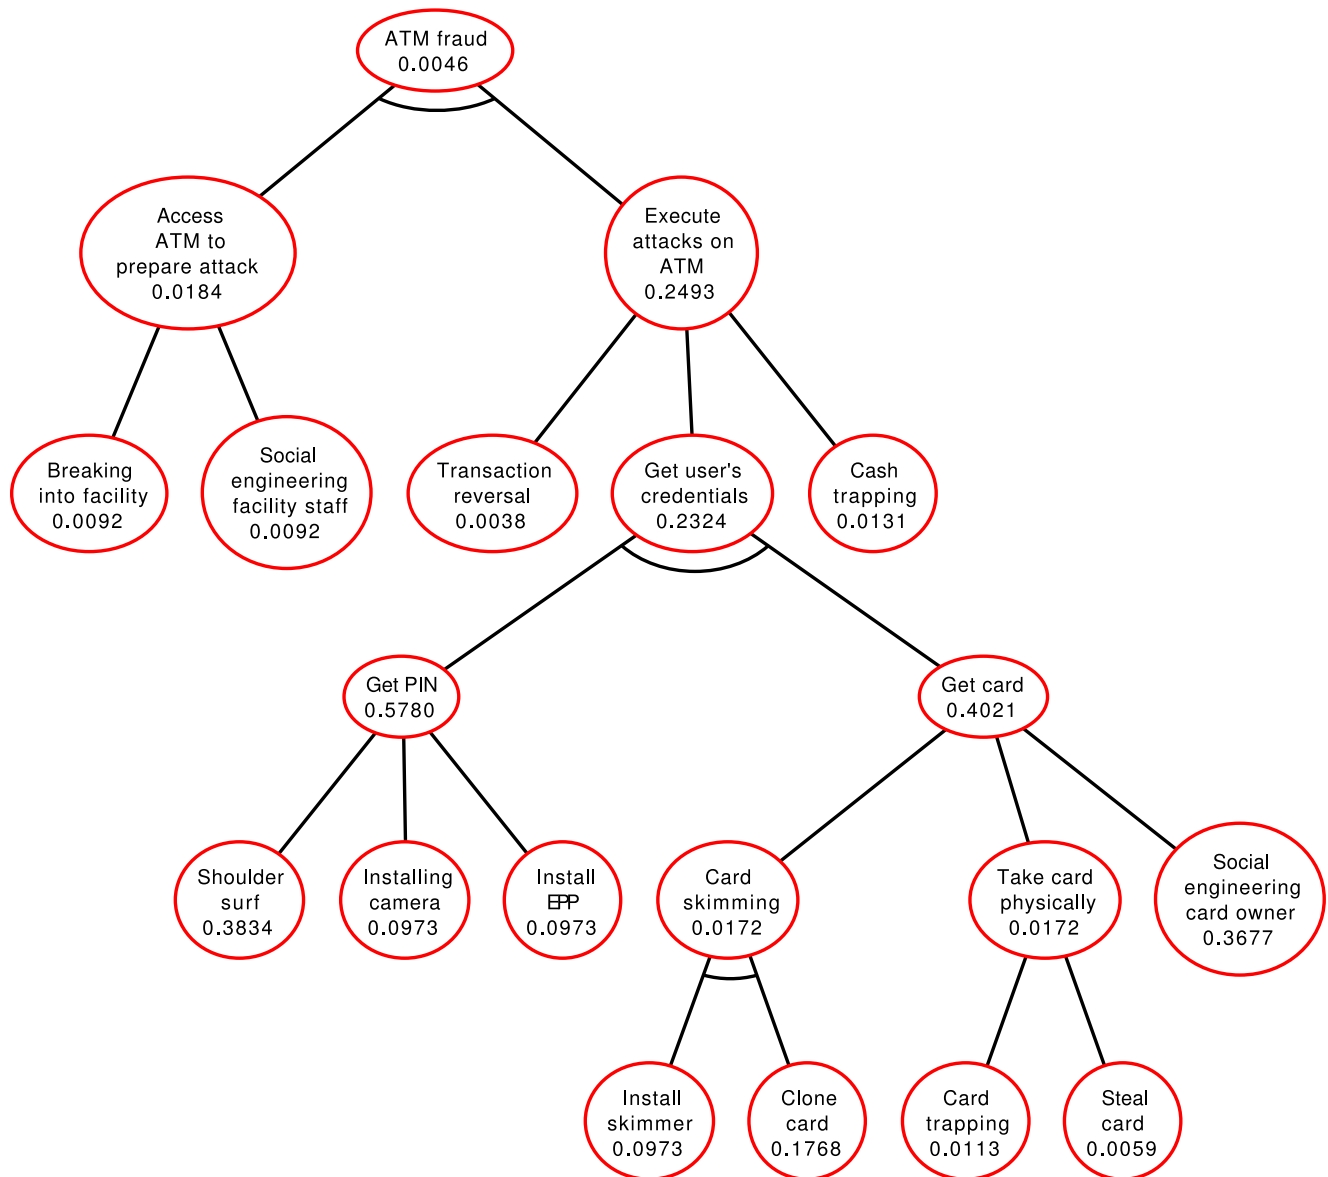


Fig. B.6. A valuation identified by the SQP-based tool for the attack tree in Fig. 3 with all hard and soft predicates listed in Section 6.2.2.

References

- Ahmed, A., Kayis, B., Amornsawadwatana, S., 2007. A review of techniques for risk management in projects. *Benchmarking* 14 (1), 22–36.
- Amenaza, 2017. SecurTree Software.
- Amoroso, E.G., 1994. Fundamentals of Computer Security Technology. Prentice-Hall, Inc.
- Arnold, F., Hermanns, H., Pulungan, R., Stoelinga, M., 2014. Time-dependent analysis of attacks. In: Proc. 3rd Int. Conf. on Principles of Security and Trust (POST'14). In: LNCS, vol. 8414. Springer, pp. 285–305.
- Aslanyan, Z., Nielson, F., 2015. Pareto efficient solutions of attack-defence trees. In: Proc. 4th Int. Conf. on Principles of Security and Trust (POST'15). In: LNCS, vol. 9036. Springer, pp. 95–114.
- Aslanyan, Z., Nielson, F., Parker, D., 2016. Quantitative verification and synthesis of attack-defence scenarios. In: Proc. 29th IEEE Computer Security Foundations Symposium (CSF'16). IEEE, pp. 105–119.
- Aven, T., 2007. A unified framework for risk and vulnerability analysis covering both safety and security. *Reliabil. Eng. Syst. Saf.* 92 (6), 745–754.
- Bagnato, A., Kordy, B., Meland, P.H., Schweitzer, P., 2012. Attribute decoration of attack-defence trees. *Int. J. Secure Soft. Eng.* 3 (2), 1–35.
- Baker, W.H., Rees, L.P., Tippet, P.S., 2007. Necessary measures: metric-driven information security risk assessment and decision making. *Commun. ACM* 50 (10), 101–106.
- Benini, M., Sicari, S., 2008. Risk assessment in practice: a real case study. *Comput. Commun.* 31 (15), 3691–3699.
- de Bijl, M.H., 2017. Using data analysis to enhance attack trees. In: Proc. Twente Student Conference.
- Bistarelli, S., Fioravanti, F., Peretti, P., 2006. Defense trees for economic evaluation of security investments. In: Proc. 1st Int. Conf. on Availability, Reliability and Security (ARES'06). IEEE.
- Böhme, R., 2010. Security metrics and security investment models. In: Proc. 5th Int. Workshop on Security (IWSEC'10). In: LNCS, vol. 6434. Springer, pp. 10–24.
- Buldas, A., Lenin, A., 2013. New efficient utility upper bounds for the fully adaptive model of attack trees. In: Proc. 4th Int. Conf. on Decision and Game Theory for Security (GameSec'13). In: LNCS, vol. 8252. Springer, pp. 192–205.
- Caralli, R., Stevens, J., Young, L., Wilson, W., 2007. Introducing OCTAVE Allegro: Improving the Information Security Risk Assessment Process. Technical Report CMU/SEI-2007-TR-012. Software Engineering Institute, Carnegie Mellon University.
- Dacier, M., Deswarte, Y., Kaënliche, M., 1996. Models and tools for quantitative assessment of operational security. In: Proc. IFIP Int. Conf. on ICT Systems Security and Privacy Protection (SEC'96). Springer, pp. 177–186.
- Frailie, M., Ford, M., Gadyatskaya, O., Kumar, R., Stoelinga, M., Trujillo-Rasua, R., 2016. Using attack-defence trees to analyze threats and countermeasures in an ATM: a case study. In: Proc. 9th IFIP Working Conference on the Practice of Enterprise Modeling (PoEM'16). In: Lecture Notes in Business Information Processing. Springer, pp. 326–334.
- Gadyatskaya, O., Hansen, R.R., Larsen, K.G., Legay, A., Olesen, M.C., Poulsen, D.B., 2016a. Modelling attack-defence trees using timed automata. In: Proc. Int. Conf. on Formal Modeling and Analysis of Timed Systems (FORMATS'16). In: LNCS, vol. 9884. Springer, pp. 35–50.
- Gadyatskaya, O., Harpes, C., Mauw, S., Muller, C., Muller, S., 2016b. Bridging two worlds: reconciling practical risk assessment methodologies with theory of attack trees. In: Proc. 3rd Int. Workshop on Graphical Models for Security (GraM-Sec'16). In: LNCS, vol. 9987. Springer, pp. 80–93.
- Gadyatskaya, O., Jhawar, R., Kordy, B., Lounis, K., Mauw, S., Trujillo-Rasua, R., 2016c. Attack trees for practical security assessment: ranking of attack scenarios with ADTool 2.0. In: Proc. 13th Int. Conf. on Quantitative Evaluation of Systems (QEST'16). In: LNCS, vol. 9826. Springer, pp. 159–162.
- Gadyatskaya, O., Trujillo-Rasua, R., 2017. New directions in attack tree research: catching up with industrial needs. In: Proc. of GraMSec. Springer, pp. 115–126.
- Holm, H., Shahzad, K., Buschle, M., Ekstedt, M., 2015. P2CySeMoL: predictive, probabilistic cyber security modeling language. *IEEE Trans. Depend. Secure Comput.* 12 (6), 626–639.
- Hong, J.B., Kim, D.S., Chung, C.J., Huang, D., 2017. A survey on the usability and practical applications of graphical security models. *Comput. Sci. Rev.* 26, 1–16.
- Horne, R., Mauw, S., Tiu, A., 2017. Semantics for specialising attack trees based on linear logic. *Fundam. Inform.* 153 (1–2), 57–86.
- Jaquith, A., 2007. Security Metrics. Pearson Education.
- Jhawar, R., Kordy, B., Mauw, S., Radomirović, S., Trujillo-Rasua, R., 2015. Attack trees with sequential conjunction. In: Proc. IFIP TC-11 Int. Information Security and Privacy Conference (IFIPSec'15). Springer, pp. 339–353, volume 455 of IFIPACT.
- Jhawar, R., Lounis, K., Mauw, S., 2016. A stochastic framework for quantitative analysis of attack-defence trees. In: Proc. 12th Workshop on Security and Trust Management (STM'16). In: LNCS, vol. 9871. Springer, pp. 138–153.
- Kordy, B., Kordy, P., Mauw, S., Schweitzer, P., 2013a. ADTool: security analysis with attack-defence trees. In: Proc. 10th International Conference on Quantitative Evaluation of Systems (QEST'13). In: LNCS, vol. 8054. Springer, pp. 173–176.
- Kordy, B., Mauw, S., Radomirović, S., Schweitzer, P., 2014a. Attack-defence trees. *J. Logic Comput.* 24 (1), 55–87.
- Kordy, B., Mauw, S., Schweitzer, P., 2013b. Quantitative questions on attack-defence trees. In: Proc. 15th Annual International Conference on Information Security and Cryptology (ICISC'12). In: LNCS, vol. 7839. Springer, pp. 49–64.
- Kordy, B., Piètre-Cambacédès, L., Schweitzer, P., 2014b. DAG-based attack and defence modeling: dont miss the forest for the attack trees. *Comput. Sci. Rev.* 13, 1–38.
- Kordy, B., Pouly, M., Schweitzer, P., 2012. Computational aspects of attack-defence trees. In: Int. Joint Conferences on Security and Intelligent Information Systems (SIIS'11). In: LNCS, vol. 7053. Springer, pp. 103–116.
- Kordy, B., Pouly, M., Schweitzer, P., 2014c. A probabilistic framework for security scenarios with dependent actions. In: Proc. Integrated Formal Methods (IFM'14). In: LNCS, vol. 8739, pp. 256–271.
- Kumar, R., Ruijters, E., Stoelinga, M., 2015. Quantitative attack tree analysis via priced timed automata. In: Proc. Int. Conf. on Formal Modeling and Analysis of Timed Systems (FORMATS'15). In: LNCS, vol. 9268. Springer, pp. 156–171.
- LeMay, E., Ford, M.D., Keefe, K., Sanders, W.H., Muehrcke, C., 2011. Model-based security metrics using adversary view security evaluation (ADVISE). Quantitative evaluation of systems (QEST). IEEE.
- Lenin, A., 2015. Reliable and Efficient Determination of the Likelihood of Rational Attacks. Tallinn University of Technology Ph.D. Thesis. TUT Press.
- Lenin, A., Buldas, A., 2014. Limiting adversarial budget in quantitative security assessment. In: Proc. 5th Int. Conf. on Decision and Game Theory for Security (GameSec'14). In: LNCS, vol. 8840. Springer, pp. 155–174.
- Lenin, A., Willemson, J., Charnamord, A., 2015. Genetic approximations for the failure-free security games. In: Proc. 6th Int. Conf. on Decision and Game Theory for Security (GameSec'15). In: LNCS, vol. 9406. Springer, pp. 311–321.
- Lenin, A., Willemson, J., Sari, D.P., 2014. Attacker profiling in quantitative security assessment based on attack trees. In: 19th Nordic Conference on Secure IT Systems (NordSec'14). In: LNCS, vol. 8788. Springer, pp. 199–212.
- Mahmood, Y.A., Ahmadi, A., Verma, A.K., Srividya, A., Kumar, U., 2013. Fuzzy fault tree analysis: a review of concept and application. *Int. J. Syst. Assur. Eng. Manag.* 4 (1), 19–32.
- Manyem, P., Ugon, J., 2012. Computational complexity, NP completeness and optimization duality: a survey. *Electron. Colloq. Comput. Complex.* 19 (9).
- Mauw, S., Oostdijk, M., 2006. Foundations of attack trees. In: Proc. 8th Int. Conf. on Information Security and Cryptology (ICISC'05). In: LNCS, vol. 3935. Springer, pp. 186–198.
- de Moura, L.M., Bjørner, N., 2008. Z3: an efficient SMT solver. In: TACAS, volume 4963. Springer, pp. 337–340. Lecture Notes in Computer Science.
- Oppliger, R., 2015. Quantitative risk analysis in information security management: a modern fairy tale. *IEEE Secur. Priv.* 13 (6), 18–21.
- Potteiger, B., Martins, G., Koutsoukos, X., 2016. Software and attack centric integrated threat modeling for quantitative risk assessment. In: Proc. Symposium and Bootcamp on the Science of Security (HotSOS'16). ACM, pp. 99–108.
- Roy, A., Kim, D.S., Trivedi, K., 2012. Scalable optimal countermeasure selection using implicit enumeration on attack countermeasure trees. In: Proc. IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN'12). IEEE.
- Ruijters, E., Stoelinga, M., 2015. Fault tree analysis: a survey of the state-of-the-art in modeling, analysis and tools. *Comput. Sci. Rev.* 15, 29–62.
- Salter, C., Saydjari, O.S., Schneier, B., Wallner, J., 1998. Toward a secure system engineering methodology. In: Proceedings of NSPW. ACM, pp. 2–10.
- Sarabi, A., Naghizadeh, P., Liu, Y., Liu, M., 2015. Prioritizing security spending: a quantitative analysis of risk distributions for different business profiles. In: Proc. 14th Annual Workshop on the Economics of Information Security (WEIS'15).
- Schneier, B., 1999. Attack Trees. *Dr Dobbs's J. Softw. Tools* 24 (12), 21–29.
- Schneier, B., 2000. *Secrets & Lies: Digital Security in a Networked World*. John Wiley & Sons, Inc., New York, NY, USA.
- Shostack, A., 2014. Threat Modeling: Designing for Security. John Wiley & Sons.
- Singhal, A., Ou, X., 2017. Security risk analysis of enterprise networks using probabilistic attack graphs. In: Network Security Metrics. Springer, pp. 53–73.
- Sommestad, T., Ekstedt, M., Johnson, P., 2009. Cyber security risks assessment with Bayesian defense graphs and architectural models. In: 2009 42nd Hawaii International Conference on System Sciences. IEEE, pp. 1–10.
- Tsang, E.P.K., 1993. Foundations of Constraint Satisfaction. Computation in Cognitive Science. Academic Press.
- Vose, D., 2008. Risk Analysis: A Quantitative Guide. John Wiley & Sons.

Ahto Buldas is a professor of cryptography at Tallinn University of Technology, and a senior research fellow at Cybernetica AS. Studied computer science at Tallinn University of Technology (1985–1991). M.Sc. on simulation techniques for Boolean circuits (1992). Ph.D. on Computational Algebraic Graph Theory (1999). His research interests are related to applied cryptography. Timestamping related research started from 1997, during which he has published papers in Crypto, Asiacrypt and PKC conferences. Participated in the development of the Estonian Digital Signature Act and ID-card (1996–2002). Current research interests also include risk analysis methods, including attack-tree semantics and game-theoretical approach to risk analysis.

Olga Gadyatskaya holds Ph.D. in mathematics from Novosibirsk State University, Russia. After a postdoc at the University of Trento, Italy, in 2014 she has joined the University of Luxembourg, Luxembourg as a research associate. Her research interests span from security risk management with attack trees to Android application security and reliability.

Aleksandr Lenin is a visiting lecturer in cryptography and information security risk management at Tallinn University of Technology, and a researcher at Cybernetica AS. Studied computer science at Tallinn University of Technology. M.Sc. in engineering (2012), Ph.D. in computer science (2015). His research interests include cryptography, quantitative aspects of system security, machine learning, and the foundations of mathematics. He publishes in venues for information security and cryptography. He has participated in the development of the SplitKey Authentication and Digital Signature Platform.

Sjouke Mauw is professor in computer security at the University of Luxembourg. He holds a master in mathematics and a Ph.D. in computer science from the University of Amsterdam. He is head of the SaToSS (Security and Trust of Software Systems) research group, which focuses on the application of formal methods to the design and analysis of secure systems. His research interests include security protocols, e-voting, security assessment, trust and risk management, privacy, and attack trees.

Rolando Trujillo-Rasua is a lecturer in Cyber Security at Deakin University. He completed a Masters and Ph.D. in computer engineering from Rovira i Virgili University, and shortly after joined the University of Luxembourg as a Postdoctoral researcher. His research interests span the areas of formal methods, computer security and privacy protection.