





The marriage between safety and cybersecurity: still practicing ★

Marielle Stoelinga^{1,2}, Christina Kolb¹, Stefano M. Nicoletti¹,
Carlos E. Budde¹, and Ernst Moritz Hahn¹

¹ Formal Methods and Tools, University of Twente, Enschede, the Netherlands
{m.i.a.stoelinga,c.kolb,s.m.nicoletti,c.e.budde,e.m.hahn}@utwente.nl

² Department of Software Science, Radboud University, Nijmegen, the Netherlands

Abstract Emerging technologies, like self-driving cars, drones, and the Internet-of-Things must not impose threats to people, neither due to accidental failures (safety), nor due to malicious attacks (security). As historically separated fields, safety and security are often analyzed in isolation. They are, however, heavily intertwined: measures that increase safety often decrease security and vice versa. Also, security vulnerabilities often cause safety hazards, e.g. in autonomous cars. Therefore, for effective decision-making, safety and security must be considered in combination.

This paper discusses three major challenges that a successful integration of safety and security faces: (1) The complex interaction between safety and security (2) The lack of efficient algorithms to compute system-level risk metrics (3) The lack of proper risk quantification methods. We will point out several research directions to tackle these challenges, exploiting novel combinations of mathematical game theory, stochastic model checking, as well as the Bayesian, fuzzy, and Dempster-Schafer frameworks for uncertainty reasoning. Finally, we report on early results in these directions.

Keywords: safety, security, model-based, interaction, Fault Trees, Attack Trees, Fault Tree-Attack Tree Integration

1 Introduction

New technology comes with new risks: drones may drop on to people, self-driving cars may get hacked, medical implants may leak in people's body. Such risks concern both accidental failures (safety) and malicious attacks (security). Here, *security* refers to the property that allows the system to perform its mission or critical functions despite risks posed by threats [74]. *Safety*, in contrast, is the absence of risk of harm due to malfunctioning behavior of technological systems [82].

Safety and security are heavily intertwined. Measures that increase safety may decrease security and vice versa: the Internet-of-Things offers ample opportunities to monitor the safety of a power plant, but their many access points are

★ This work was partially funded by ERC Consolidator Grant 864075 (*CAESAR*).

notorious for enabling hackers to enter the system. Passwords secure patients' medical data, but are a hindrance during emergencies. It is therefore widely acknowledged, also by international risk standards [64, 103], that safety and security must be analyzed in combination [9, 97]. The overarching challenge in safety and security risk management is decision making: which risks are most threatening, and which countermeasures are most (cost-)effective? Such decisions are notoriously hard to take: it is well-known (e.g. from Nobel prize winner Daniel Kahneman [71]) that people, have very poor intuitions for risks and probability.

Vision. To make effective decisions, *risk management should be accountable.*

1. *systematic*, so that no risks are overlooked;
2. *transparent*, so that experts can share and discuss their viewpoints;
3. *objective*, i.e. based on recorded facts and figures, rather than on (fallible) intuitions.

Hurdles. Tough hurdles that have hindered the effective integration of safety and security [3, 82] are their opposite perspectives on:

- H1. *User intention*: safety concerns unintended mishaps, while security is about malicious attacks.
- H2. *Dynamics*: Whereas safety analysis is often static, developing design-time solutions; security demands constant defence against new vulnerabilities.
- H3. *Risk quantification*: Whereas safety analysis can fruitfully exploit historic failure data, risk quantification for security is a major open problem. With hackers continuously changing their targets and strategies, historic data is of little value. Therefore, security decisions are often based on subjective estimates.

The demanding challenge in safety-security co-analysis is to overarch these diametrical viewpoints.

Challenges. To overcome these hurdles above and make decision making about safety and security less ad hoc, and more systematic, transparent, and quantitative, three challenges have to be solved.

- A systematic way to map safety and security risks, identifying how failures and vulnerabilities propagate through the system and lead to system level disruptions.
- Effective algorithms to compute system level risk safety and security metrics, together with diagnostic algorithms that explain how such metrics arise, and how one could improve these.
- Novel risk quantification methods. Reliable numbers are indispensable in decision making. Since objective data is often not available, we need algorithms that reason under uncertainty.

The ERC project CAESAR. The ERC-funded project CAESAR picks up the challenges above, exploring three novel directions:

1. *Game-theoretic methods* uniting the cooperative versus malicious user intention in safety versus security (H1). Our aim is to model the attacker versus defender as two players in a stochastic game. We will focus on a time dependent game (H2) that faithfully models the complex interaction between safety and security aspects.
2. *Stochastic model checking techniques* to compute safety-security risk metrics. Metrics are pivotal to prioritize risks and select effective countermeasures, as they clarify how failures and attacks affect system-level performance. Since risk is defined as a combination of likelihood and severity, metrics are stochastic by nature. Apart from computing numbers, we will also elucidate how these numbers arise.
3. *Risk quantification methods that handle data uncertainty.* Effective decision making requires insight in the most frequent failures and attacks. Since objective data is scarce, security decisions are often based on subjective estimates. We will combine objective and subjective probabilities, exploiting three prominent frameworks for data uncertainty: Bayesian reasoning, fuzzy logic and Dempster-Schafer theory. These explicate the underlying assumptions and (dis)agreements about risks.

Contributions. This paper outlines the first results and the approach taken in CAESAR: we present findings of a literature survey, where we compare existing formalisms for safety-security co-analysis and identify several gaps. An important outcome of our survey is that most of these formalisms are based on various combinations of the popular formalisms of attack trees for security analysis and fault trees for safety analysis. One important instance of a research gap is that in fault trees and attack trees OR-gates are interpreted in a different manner. This difference in interpretation is not considered in current analysis algorithms. To obtain a unified framework for safety and security building on these mechanisms, we thus have to unify the interpretation of such gates. Afterwards, we outline how CAESAR aims to solve such gaps. We discuss how recent results in attack tree analysis provides results for tree-shaped attack trees and fault trees as well as a formal semantics of DAG-shaped attack trees and fault trees. In these results, we exploit methods based on binary decision diagrams. Throughout the paper, we indicate current results as well a research gaps.

Organization of the paper. Section 1 has provided an introduction to this paper. Section 2 provides some background on the interaction between safety and security and on the two formalisms attack trees and fault trees. Section 3 presents an overview of formalisms of safety-security co-analysis. Section 4 provides a comparison between attack trees and fault trees, highlighting similarities and differences and defining metrics. Section 5 discusses analysis algorithms for attack and fault trees. Section 6 concludes the paper.

2 Background

Attack trees and fault trees. Attack trees and fault trees are popular models in dependability analysis, representing respectively how low-level attacks and failures propagate through the system and lead to system-level disruptions. As shown in Fig. 1, these are tree-like structures that culminate in a top level event (TLE), which models a system-level failure or attack. The TLE is thus refined via intermediate events, equipped with gates: the AND-gate indicates that all children must fail (be attacked) in order for the gate to fail (be attacked). For the OR-gate to fail (be attacked), at least one of its children need to fail (be attacked). The leaves in a FT are called *basic events (BEs)* and model atomic failures; the leaves in ATs model atomic attack steps, called *basic attack steps (BASs)*. Despite their names, FTs and ATs are directed acyclic graphs, since intermediate events can share children.

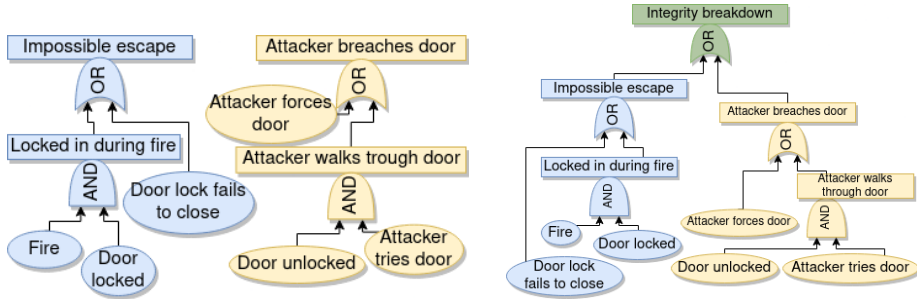


Figure 1: Fault tree (left), attack tree (center) and their combination (right). These represent respectively safety, security, and combined risks. In the FT, for the intermediate event "locked in during fire" to happen, both a fire and the door being locked have to occur, modelled through an AND-gate. In the AT, for an attacker to breach the front door he/she needs to either walk through an unlocked door or to force a locked door, modeled as an OR. On the right, a possible combination of ATs and FTs in the attack fault trees formalism [126].

FTs and ATs enable numerous analysis methods [77]: *Cut set analysis* indicates which combinations of BEs or BASs lead to the TLE. The set {Fire, Locked} is a cut set in Figure 1. Quantitative analyses compute *dependability metrics*, such as the system reliability, attack probabilities and costs. For example, by equipping the BEs and BASs with probabilities, one can compute the likelihood of a system level failure or attack to occur.

Both FTs and ATs are part of international standards [37] and have been used to analyse numerous case studies [135, 50, 30]. FTs and ATs also feature some differences: FTs often focus on probabilities, whereas ATs consider several other attributes, like cost, effort and required skills. Further, FTs have been extended

with repairs [113], and dynamic gates [66, 43]; ATs with defenses, and sequential AND (SAND) gates [76, 50].

Section 4 presents a more formal treatment of fault trees and attack trees, and in particular their different quantitative interpretation of the OR-gate. Their comparison is summarized in Table 3.

Safety-security dependencies. One of the key challenges in safety-security co-analysis is to model their interaction. The paper [82] has identified four safety-security dependencies:

- *Conditional dependency* means that fulfillment of safety requirements conditions security or vice-versa.
- *Mutual reinforcement* means that fulfillment of safety requirements or safety measures contributes to security, or vice-versa, thereby enabling resource optimization and cost reduction.
- *Antagonism* arises when safety and security requirements or measures, considered jointly, lead to conflicting situations.
- *Independency* means that there is no interaction between safety and security properties.

Figure 1 shows a classical example of antagonism: the door needs to be locked in order to prevent an attacker from entering the house (security requirement), but it has to be unlocked to allow the owner to escape during a fire (safety requirement). In this scenario, mutual reinforcement can be achieved by introducing a fire door: this contributes to safety by limiting the spread of an eventual fire and to security by increasing the robustness of the door, thus making it harder to breach. Conditional dependency is, in our view, always present: for the lock to ensure security, it must function properly. E.g., it must not break when locking door. Similarly, safety solutions must be secure and not be hacked: it must not be possible to easily force the door open.

3 Formalisms for safety-security co-analysis: an overview

3.1 The formalisms

As a first step in the CAESAR project, we carried out a literature survey [?], comparing the most prominent formalisms for safety-security co-analysis. Via a systematic literature search [25], which also considered earlier surveys on this topic [82, 100, 33], we have identified 10 important formalisms for model-based safety-security co-analysis. These are summarized in Table 1.

A first remarkable result of our survey is that the majority of safety-security formalisms combines attack trees (ATs) and fault trees (FTs). ATs and FTs are well established formalisms, extensively used in industry and academia. As previously mentioned, they are similar nature, and model respectively how failures and attacks propagate through a system. In that sense, combining attack trees and fault trees is a natural step. We further divided these approaches into two categories (plain combinations and extended combinations). A third category comprises architectural formalisms.

Formalism	Ref.	Year	#Citations
Fault Tree/Attack Tree Integration (FT/AT)	[49]	2009	170
Component Fault Trees (CFTs)	[126]	2013	52
Attack-Fault Trees (AFTs)	[83]	2017	56
State/Event Fault Trees (SEFTs)	[110]	2013	25
Boolean Driven Markov Processes (BDMPs)	[81]	2014	36
Attack Tree Bow-ties (ATBTs)	[15]	2017	9
STAMP	[52]	2017	120
SysML	[104]	2011	72
Architectural Analysis and Design Language (AADL)	[41]	2020	0
Bayesian Networks (BNs)	[80]	2015	46

Table 1: Overview of safety-security formalisms. Citations from Google Scholar, April 2021.

1. *Plain combinations of attack trees and fault trees.* These formalisms combine attack trees and fault trees without adding additional constructs: fault tree/attack trees (FT/AT) [49], which investigate how the a basic event of a FT can be triggered by an attacker, refining these with ATs with the event in question as goal, component fault trees (CFTs) [126] merge attack trees and fault trees without any restrictions, Attack-Fault Trees (AFTs) [83] merge dynamic attack trees and dynamic fault trees.

2. *Extensions of attack trees-fault tree combinations.* These merge attack trees, fault trees with additional constructs: State/Event Fault Trees (SEFTs) [110] exploit Petri nets to refine the basic attack steps in an attack tree and the basic component failures in a fault tree. The Petri nets can for instance model that the attack and failure behavior is different depending whether a door is open or closed. Boolean Driven Markov Processes (BDMPs) [81] extend attack trees and fault trees with both Petri nets and triggers. The latter model sequential behaviour, where one fault or attack triggers another one. Finally, Attack Tree Bow-ties (ATBTs) [15], extend bowties [99] with attack trees, where bowties themselves combine fault trees with event trees.

3. *Architectural formalisms and bayesian networks.* Apart from combinations of attack trees and fault trees, we identified a third category, containing formalisms that take as a starting point the system architecture: The Systems-Theoretic Accident Model and Processes (STAMP) [52], is an accident causality model, rooted in the observation that system risks do not come from component failures, but rather from inadequate control or enforcement of safety and security constraints. Systems-Theoretic Process Analysis then systematically identifies the consequences of incorrect control and feedback actions, e.g., when these happen too early, in the wrong order, or were maliciously inserted.

SysML-sec [111] extend the SysML modeling framework with safety and security requirements, which can be checked using model checkers. In particular, SysML-sec supports the modelling of communication channels between processes with the encryption methods and their complexity overhead.

The Architectural Analysis and Design Language (AADL) [41] enables safety analysis, via the AADL error model, and security analysis via the AADL LAMP extension. In this way the same AADL model can be separately analyzed to investigate safety and security properties.

Finally, albeit somewhat artificially, we also put Bayesian Networks (BNs) for safety-security analysis in this category [80]. This model allows to represent probabilistic dependencies between several variables via a directed acyclic graph. BNs are used to model safety and security dependencies. The two root nodes represent system safety and security. BNs can analyze which nodes influence other nodes and how (conditional independence analysis) calculate reliability metrics.

3.2 Findings

The analysis we performed highlighted some notable findings, summarized in Table 2. For each analyzed formalism we highlight the dependencies it captures, its modeling constructs, the analysis types it enables, case studies that were performed deploying this formalism and possible tools.

Formalism	Dependencies				Modelling	Analysis		Application	Tool
	A	CD	MR	I		QL	QT		
FT/AT	*	→		x	ATs refine FT leaves	x	x	Chemical plant	
CFTs	*	x	x	x	Merge ATs + FTs	x	x	Cruise control	
AFTs	x	x	x	x	Merge dynamic ATs + FTs	x	x	Pipeline, lock door	UPPAAL
SEFTs	x	→	x	x	FTs + Petri nets	x	x	Tyre pressure, lock door	ESSaRel
BDMPs	x	x	x	x	Triggers, Petri nets	x	x	Pipeline, lock door	KB3, Figaro
ATBTs	*	→		x	Bowties + FT/AT	x	x	Pipeline, Stuxnet	
STAMP				x	Process controller	x		Synchronous-islanding	
SysML					System components	x		Embedded systems	TTool
AADL				x	System components + ports	x		Lock door	Cheddar, Marzhin
BNs	x	x	x	x	Conditional prob.	x	x	Pipeline	MSBNx

Table 2: Comparison of safety-security formalisms. A= Antagonism, CD= Conditional Dependency, MR=Mutual reinforcement, I=Independence. * = capable when NOT-gate is supported. → = capable but only directional from security to safety.

Finding #1: The majority of approaches combine attack trees and fault trees. As stated, six out of the ten formalisms combine attack trees and fault trees. This is not surprising, since FTs and ATs are well established model-based formalisms, extensively used both in industry and academia.

Finding #2: No novel modeling constructs are introduced. Despite the shown combinations and extensions, existing safety-security co-analysis do not introduce novel modeling constructs to capture safety-security interactions. They do merge existing safety and security formalisms, however they do not add new operators. As such, they are suited to represent safety and security features in one model, but do not seem to appropriately capture the interaction between them.

Finding #3: Safety-security interactions are still ill-understood. In spite of the definitions provided in [82], we are convinced that safety-security interactions can still be defined more thoroughly by adopting more rigour and by focusing on requirements and events. In particular, it is not so clear to which entities the safety-security interactions refer: do these concern safety-security requirements, measures, or something else? Clarifying their definitions is a prerequisite for properly modeling safety-security interactions in a mathematical formalism.

Finding #4: No novel metrics were proposed. Analyzed formalisms adopt classic metrics, such as mean time to failure and attacker success probabilities. However, none of them introduce new metrics to quantify safety-security dependencies or to analyze trade-offs, e.g., through Pareto analysis. New metrics and trade offs are paramount to understand the interaction between safety and security aspects.

Finding #5: No large case studies were carried out. To the best of our knowledge, no large case studies were carried out. The majority of analyzed papers present small examples used to showcase the formalism in question. Some notable exceptions are [81] and [83]: here, the medium-sized example of a pipeline is presented. However for safety and security, when considered separately, large case studies do exist [22].

Finding #6: Different formalisms model different safety-security interactions. As shown in Table 2, only AFTs, BDMPs and BNs can unconditionally model all four dependencies between safety and security. CFTs and SEFTs can model them provided with extensions/with some limitations.

Research gaps.

- Realistic *large-sized case studies* concerning safety-security interactions are still missing. Performing large-sized case study analysis would contribute to further address additional gaps;
- It would clarify the nature of *safety-security interactions*, that are currently still ill-understood. Furthermore, it would help improve standard definitions for safety-security interdependencies;
- From this understanding, *novel metrics* and *novel modeling constructs* for safety-security co-analysis - that are still missing - could be developed.

4 Attack trees versus fault trees

We saw that most safety-security formalisms combine attack trees and fault trees. This is a natural step, since attack trees and fault trees bear many similarities. What is less known, is that they also feature a number of remarkable differences, elaborated in [28]. In particular, the interpretation of the OR-gate is crucially different in attack trees than in fault trees, and therefore their analysis should not be mindlessly combined. Below, we present the most remarkable similarities and differences, summarized in Table 3.

4.1 Attack trees versus fault trees

It is no surprise that ATs and FTs are similar to each other, since ATs were inspired by FTs. FTs were introduced in 1961 at Bell Labs to study a ballistic missile [130, 117, 36]. Weiss introduced threat logic trees—the origin of ATs—in 1991, and its “similarity... to fault trees suggests that graph-based security modelling has its roots in safety modelling” [77].

Attack trees and fault trees come in various variants and extensions. Following [27], we categorize these along two axes. First, we distinguish between *static* and *dynamic* trees. Static attack and fault trees are equipped with Boolean gates. Dynamic trees come with additional gates to model time-dependent behavior. Second, we distinguish between *tree-shaped* and *DAG-shaped* trees. Trees are relatively easy to analyse via a bottom up algorithm. This algorithm works for all quantitative attributes (cost, time, probability) as long as they constitute an attribute domain.

	Attack trees	Fault trees
Syntax		
Leaves	Basic attack steps (BAS)	Basic events (BEs)
Non-leaves	Subgoals	Intermediate events (IEs)
Static gates	AND, OR	AND, OR, VOT
Dynamic gates	SAND	SPARE, FDEP, PAND
Other extensions	Defenses	Repairs, maintenance
Analysis		
Qualitative	(Minimal) attack vectors/scenarios	(Minimal) cut sets
Attributes	probability, cost, time, skill, impact	probability
Metrics	Min cost, time, skill Max impact, probability	Reliability, availability, MTTF, MTBF
Semantics		
Qualitative	Structure function	Structure function
Stochastic		
AND(a,b)	$p_a \cdot p_b$	$p_a \cdot p_b$
OR(a,b)	$\min(p_a, p_b)$	$p_a + p_b - p_a \cdot p_b$

Table 3: Differences between attack and fault trees

4.2 The static case

Syntax. The basic variants, called *static* or *standard* fault and attack trees, have the exact same syntax: trees or DAGs equipped with AND and OR gates. Fault trees often contain a (k, m) voting gate, which fails if k out of the n inputs fail; these can however be expressed in terms of AND and OR gates. We use the word *disruption tree* (DT) for either an attack tree or a fault tree.

Formally, DTs are rooted DAGs with typed nodes, for which we consider types $\mathbb{T} = \{\text{LEAF}, \text{OR}, \text{AND}\}$. For Booleans we use $\mathbb{B} = \{1, 0\}$. The edges of a DT are given by a function ch that assigns to each node its (possibly empty)

sequence of children. We use set notation for sequences, e.g. $e \in (e_1, \dots, e_m)$ means $\exists i. e_i = e$, and we denote the empty sequence by ε .

Definition 1. A disruption tree is a tuple $T = (N, t, ch)$ where:

- N is a finite set of nodes;
- $t: N \rightarrow \mathbb{T}$ gives the type of each node;
- $ch: N \rightarrow N^*$ gives the sequence of children of a node.

Moreover, T satisfies the following constraints:

- (N, E) is a connected DAG, where $E = \{(v, u) \in N^2 \mid u \in ch(v)\}$;
- T has a unique root, denoted $R_T: \exists! R_T \in N. \forall v \in N. R_T \notin ch(v)$;
- LEAF $_T$ nodes are the leaves of $T: \forall v \in N. t(v) = \text{LEAF} \Leftrightarrow ch(v) = \varepsilon$.

4.3 Semantics

Semantics pin down the mathematical meaning for attack and fault trees. The semantics of both fault trees and attack trees is given in terms of their *structure function*, indicating which sets of leaves cause the top level events to happen.

Thus, the *structure function* of a disruption tree T is a function $f_T: \mathbb{B}^n \rightarrow \mathbb{B}$. Technically, a status vector $\mathbf{v} = \langle v_1, \dots, v_n \rangle$ indicates for each leaf i whether it was disrupted, i.e., $v_i = 1$ if leaf i has failed or was attacked. Then $f_T(\mathbf{v}) \in \{0, 1\}$ indicates whether the whole system was disrupted. This function can be defined recursively in the nodes of T : $f_T(v, A)$ tells whether $A \subseteq \text{LEAF}$ suffices to disrupt node $v \in N$ of T , where A encodes \mathbf{v} as usual.

Definition 2. The structure function $f_T: N \times 2^{\text{LEAF}} \rightarrow \mathbb{B}$ of a disruption tree T is given by:

$$f_T(v, A) = \begin{cases} 1 & \text{if } t(v) = \text{OR} \quad \text{and } \exists u \in ch(v). f_T(u, A) = 1, \\ 1 & \text{if } t(v) = \text{AND} \quad \text{and } \forall u \in ch(v). f_T(u, A) = 1, \\ 1 & \text{if } t(v) = \text{LEAF} \quad \text{and } v \in A, \\ 0 & \text{otherwise.} \end{cases}$$

The structure function can be used to assess suites: a *disruption suite* $\mathcal{S} \subseteq 2^{\text{LEAF}}$ represents all ways in which the system can be compromised. From those, one is interested in disruptions $A \in \mathcal{S}$ that actually represent a threat. These correspond to (minimal) cut sets in fault trees and attack scenarios in attack trees. To find them, we let $f_T(A) \doteq f_T(R_T, A)$ and call disruption A *successful* if $f_T(A) = 1$, i.e. it makes the top-level of T succeed (resp. be attacked or failed). If, moreover, no proper subset of A is successful then A is a *minimal disruption*.

It is well known that attack trees and fault trees are *coherent* [12], meaning that adding attack steps/basic events preserves success: if A causes the TLE to happen, then so is $A \cup \{a\}$ for any $a \in \text{LEAF}$. Thus, the suite of successful disruptions of an DT is characterised by its minimal disruptions.

Definition 3. The semantics of a static DT T is its suite of minimal disruptions: $\llbracket T \rrbracket = \{A \subseteq 2^{\text{LEAF}} \mid f_T(A) \wedge A \text{ is minimal}\}$.

4.4 Metrics for attack and fault trees

Dependability metrics quantify key performance indicators (KPIs), that quantify several dependability characteristics of a system. Such metrics serve several purposes, e.g. allowing to compare different design alternatives w.r.t. the desired dependability features; computing the effectiveness of measures; verifying whether a solution meets its dependability requirements; etc.

Metrics for attack trees focus on a wide variety of attributes, such as the cost of an attack, its time and, success probability. These can be conveniently summarized via an attribute domain [93]. Metrics for fault trees focus on probabilistic aspects, such as the system reliability (ie the probability that a system fails within its mission time T), the availability (i.e., the average percentage of time that the system is up), mean time to failure, etc.

Attribute metrics. We define dependability metrics for DTs in three steps: first an attribution α enrich the leaves with attributes, assigning a value to each $a \in \text{LEAF}$, then a dependability metric $\hat{\alpha}$ assigns a value to each disruption scenario A ; and finally the metric $\check{\alpha}$ assigns a value to each disruption suite.

Example 1. Consider the AT in Figure 1b. The metric we study is the time required to execute a successful attack. Thus, the attribution α equips each AT leaf with its attack time, setting e.g. $\alpha(\text{Attacker forces door}) = 5$, $\alpha(\text{Door unlocked}) = 0$ and $\alpha(\text{Attacker tries door}) = 2$. If all attack steps are executed sequentially, then the the time needed to execute an attack $A = \{a_1, \dots a_n\}$ is sum of the attack times of the BASs:

$$\hat{\alpha}(a_1, \dots a_n) = \sum_{i=1}^n \alpha(a_i)$$

Both for attackers and defenders of the system, it is relevant to consider the shortest attack in an attack suite \mathcal{S} :

$$\check{\alpha}(\{A_1, \dots A_n\}) = \min\{\hat{\alpha}(A_1), \dots \hat{\alpha}(A_n)\}$$

Other metrics give rise to other attribute definitions. For example, the success probability of an attack is the product of the success probabilities of the BAS. The probability of attack suite \mathcal{S} is is the probability of the most successful attack in \mathcal{S} . The cost of an attack is the sum of the cost of the BASs, and the cost of an attack suite is the minimum cost of its attacks. A formal definition is as follows.

Definition 4. *Given a DT and a set V of values:*

1. *an attribution $\alpha: \text{LEAF} \rightarrow V$ assigns an attribute value $\alpha(a)$, or shortly an attribute, to each leaf a ;*
2. *a dependability metric over disruptions is a function $\hat{\alpha}: \mathcal{A}_T \rightarrow V$ that assigns a value c to each disruption A ;*
3. *a dependability metric over disruption suites is a function and to a function $\check{\alpha}: \mathcal{S}_T \rightarrow V$ that assigns a value $\check{\alpha}(\mathcal{S})$ to each disruption suite \mathcal{S} .*

We write $\check{\alpha}(T)$ for $\check{\alpha}(\llbracket T \rrbracket)$, setting the metric of a DT to the metric of its minimal disruption suites.

Remark 1. We choose the notation $\hat{\alpha}$ for metrics over disruptions, since $\hat{\alpha}$ resembles Δ , and $\hat{\alpha}(A)$ corresponds to the interpretation of the AND gate. Similarly, $\check{\alpha}$ resembles ∇ , and corresponds to the OR gate, since $\check{\alpha}(S)$ often chooses the best disruption set among all $A \in S$.

Different metric interpretation of the OR-gate. It is important to realize that the quantitative interpretation of the OR-gate is different in attack trees than in fault trees. Fault trees assume that all components work in parallel. Thus, component i fails with probability p_i , the fault tree $OR(C_1, C_2)$ fails with probability $p_1 + p_2 - p_1 \cdot p_2$. In attack trees, the OR-gate works in parallel. The interpretation of the attack tree $OR(C_1, C_2)$ is that the attacker executes either C_1 or C_2 . Since the attacker maximizes their success probability, the probability on a successful attack in the tree $OR(C_1, C_2)$ equals $\max(p_1, p_2)$.

This distinction is completely ignored in the analysis methods for all six attack-fault combinations/extensions [49, 126, 83, 110, 81, 15]. In particular, the analysis methods for computing probabilities may not account for the different interpretation of the OR-gates related to safety or security events. This could further lead to incorrect computations of dependability metrics e.g., probability values.

5 Analysis algorithms for attack and fault trees

Numerous analysis methods for quantitative analysis of attack trees and fault trees exist [5, 18, 38, 69, 78, 108, 118]. In this section, we give an overview of two common algorithms for fault trees and attack trees.

5.1 Algorithms for tree-shaped DTs

We first provide the algorithms for tree-structured DTs, where every node in the graph has a single parent. These can be analyzed via a bottom-up algorithm, propagating the values from the bottom to the root of the tree. In order for this procedure to work for all metrics, we combine the inputs of the AND-gate using an operator Δ , and the inputs of the or gate via ∇ . Then this procedure works whenever the algebraic structure (V, Δ, ∇) constitutes a semiring [93].

Next, we treat the computationally more complex DAG-structured DTs. These can be analyzed by converting the DT to a binary decision diagram (BDD). Again, this works if (V, Δ, ∇) is a semiring [29].

<p>Input: Tree-structured DT T, node v from T, attribution $\alpha: \text{BAS}_T \rightarrow V$, semiring attribute domain $D = (V, \nabla, \Delta)$.</p> <p>Output: Metric value $\check{\alpha}(T) \in V$ from node v downwards.</p> <pre> 1 if $t(v) = \text{OR}$ then 2 return $\nabla_{u \in \text{ch}(v)} \text{BU}(T, u, \alpha, D)$ 3 else if $t(v) = \text{AND}$ then 4 return $\Delta_{u \in \text{ch}(v)} \text{BU}(T, u, \alpha, D)$ 5 else // $t(v) = \text{BAS}$ 6 return $\alpha(v)$ </pre> <p>Algorithm 1: BU for tree DTs</p>	<p>Input: BDD B_T from static DT T, node w from B_T, attribution $\alpha: \text{BAS}_T \rightarrow V$, semiring attribute domain $D_* = (V, \nabla, \Delta, 1_\nabla, 1_\Delta)$.</p> <p>Output: Metric value $\check{\alpha}(T) \in V$ from node w downwards.</p> <pre> 1 if $\text{Lab}(w) = 0$ then 2 return 1_∇ 3 else if $\text{Lab}(w) = 1$ then 4 return 1_Δ 5 else // non-terminal $\text{Lab}(w) = v \in \text{BAS}_T$ 6 return $\text{BDD}(B_T, \text{Low}(w), \alpha, D_*) \nabla$ ($\text{BDD}(B_T, \text{High}(w), \alpha, D_*) \Delta \alpha(v)$) </pre> <p>Algorithm 2: BDD for static DAG DTs</p>
--	---

Example 2. We illustrate the (straightforward) bottom up algorithm via the attack tree in Figure 2. We compute the time required to reach the top event, with the same attribute values as before: Abbreviating f = Attacker forces door, u = Door unlocked and t = Attacker tries door, we have $\alpha(f) = 5$, $\alpha(u) = 0$ and $\alpha(t) = 2$. The bottom up computation first computes the time required to achieve the subgoal "Attacker walks through door", abbreviated as w . Since the attack time metric interprets the AND-gate as the sum, we take $\Delta = +$ and obtain the value for w as the sum of the metric values of its children "Attacker forces door" and "Door unlocked", abbreviated u and f respectively.

$$\check{\alpha}(w) = \check{\alpha}(f) \Delta \check{\alpha}(u) = \alpha(f) \Delta \alpha(u) = 0 + 2 = 2.$$

Similarly, we compute the time required for the TLA "Attacker breaches door", abbreviated as b . Since the attack time metric interprets the OR-gate as the minimum, we take $\nabla = \min$ and obtain the value for b as the minimum of the metric values of its children f and w :

$$\check{\alpha}(b) = \check{\alpha}(f) \nabla \check{\alpha}(w) = \min(\alpha(f), \alpha(w)) = \min(5, 2) = 2.$$

The above procedure, formalized in [Algo. 1](#), works whenever the structure (V, ∇, Δ) constitute an attribute domain. Note that this algorithm is linear in the number of DT nodes.

Definition 5. Let V be a set:




Figure 2: [Algo. 1](#) for min. attack time.

1. an attribute domain over V is a tuple $D = (V, \nabla, \Delta)$, whose disjunctive operator $\nabla: V^2 \rightarrow V$, and conjunctive operator $\Delta: V^2 \rightarrow V$, are associative and commutative;
2. the attribute domain is a semiring[†] if Δ distributes over ∇ , i.e. $\forall x, y, z \in V. x \Delta (y \nabla z) = (x \Delta y) \nabla (x \Delta z)$;
3. let T be a static DT and α an attribution on V . The metric for T associated to α and D is given by:

$$\check{\alpha}(T) = \underbrace{\bigtriangledown}_{\check{\alpha}}_{A \in \llbracket T \rrbracket} \underbrace{\bigtriangleup}_{\hat{\alpha}}_{a \in A} \alpha(a).$$

5.2 Algorithms for DAG-shaped DTs

DTs with shared subtrees cannot be analysed via a bottom-up procedure on their DAG structure. This is a classical result from fault tree analysis [?]. Intuitively, the problem is that a visit to node v in any bottom-up procedure that operates on the DT structure can only aggregate information on its descendants.

This is illustrated by the DAG-shaped AT in Figure 3: We assign attack time to the leaves: $\alpha(a) = 3$, $\alpha(b) = 2$ and $\alpha(c) = 4$. Then the bottom up algorithm yields the following results: for the OR-gates, we take the minimum value between the children, which both equal 2, and for the AND-gate we sum the values of the children, resulting in 4. However, this computation does not take the sharing of b into account. In fact, the shortest attack is to take the BAS b , which takes time 2.

As a matter of fact, computing metrics in a DAG-structured DT T is an NP-complete problem. Various methods to analyse DAG-structured DTs have been proposed: contributions over the last 15 years include [5, 18, 38, 69, 78]. We now detail our recent work on binary decision diagrams (BDDs) [29].

BDD algorithms. BDDs offer a very compact representation of Boolean functions, and can therefore represent the structure function of a DT T : Each BDD node v is labeled with a leaf a of T , and has two children: its left node v_L (reached via a dashed line) represents the structure function of T in case a has the value 0; its right child v_R (reached via a solid line) represents the structure function of T in case a evaluates to 1. The key insight in [29] is that the values of an




Figure 3: DAG-shaped AT (left) and its BDD (right).

[†]Since we require Δ to be commutative, D is in fact a commutative semiring. Further, rings often include a neutral element for disjunction and an absorbing element for conjunction, but these are not needed in Def. 5.

attribute domain can be computed recursively over this BDD, thereby avoiding duplication of values as in Figure 3. The idea is as follows. The value for the BDD terminal node labeled with 0 is set to the constant $1_{\nabla} \in V$; the BDD terminal node labeled with 1 is set to $1_{\Delta} \in V$. For an internal node v with children v_L and v_R , we proceed as follows: When choosing v_R , i.e. the basic event a occurs, we extend the value computed at v_R with the attribute value of a . We do so via the Δ operator, since taking the right child corresponds to executing both a and all leaves needed in $\alpha(v_R)$. If a is not executed, then we do not incur the value of $\alpha(a)$, and only take $\check{\alpha}(v_L)$. Now, the best disruption (i.e. attack or cut set) is obtained by choosing the best option, by deploying the ∇ operator: either one does not execute a , incurring $\check{\alpha}(v_L)$, or executes a and incurs $\alpha(a) \Delta \check{\alpha}(v_R)$. This yields the value $\check{\alpha}(v) = \check{\alpha}(v_L) \nabla (\alpha(a) \Delta \check{\alpha}(v_R))$. This is illustrated in (Fig. 3, in blue). As we can see, the TLE can fail either by the failure of b in 2 time units or by the failure of a and c but not of b , in 7 time units. Algo. 2 shows the pseudocode of this algorithm. The algorithm is linear in the size of the BDD, but that the BDD size can be exponential in the size of the DT. In particular, the BDD size heavily depends on the order for the variables. In practice good heuristics are available, making BDD-computations efficient in practice.

5.3 Research gaps

Research gaps.

- *An overarching formalism* is still missing. Since attack trees and fault trees interpret the OR-gate in a different manner, proper combinations must feature two variants of the OR-gate: one that coincides with the AT-interpretation and for the FT-variant.
- Another research gap concerns *proper analysis of OR-gates*. Analysis algorithms must handle both the aforementioned variants of OR-gates. Section 5 partially solves this problem for the case of tree-structured attack-fault trees. Efficient analysis of DAG-structured attack-fault tree combinations remains an open problem.

6 Conclusions

Conclusion. Safety and security interactions have been identified as important topics in complex systems, and multiple modeling methods have been developed in an attempt to account for their interactions. Our preliminary results show that most of these methods are based on extending and/or combining existing safety and security modeling methods. No specific metrics or novel modeling constructs are introduced. The majority of considered formalisms combine/extend attack trees and fault trees. As a consequential next step, we performed a thorough analysis of similarities and differences between ATs and FTs. Their static variants - SATs and SFTs - share the same syntax: we group them under the label of

disruption trees (DTs), for which we provide shared semantics. Furthermore, we show how to compute dependability metrics on DTs highlighting differences between ATs and FTs when needed, e.g., the different interpretation of the OR-gate. Finally, we propose analysis algorithms for ATs and FTs both for their tree-shaped and DAG-shaped variants.

Future work. While addressing some of the research gaps, this work also highlights future challenges. With the growing need for safety-security co-analysis, the urge of a better understanding of safety-security interactions arises:

Open problems.

- To foster this understanding, realistic *large-sized case study analysis* concerning safety-security interactions should be performed.
- This would clarify the nature of *safety-security interactions* - that are still ill-understood - and help *improve standard definitions* for safety-security interdependencies.
- As mentioned, *novel metrics* and *novel modeling constructs* for safety-security co-analysis - that are still missing - could then be developed, alongside an overarching formalism.
- This *overarching formalism* would need to account for two different OR-gates: one that coincides with the AT-variant and one for the FT-variant.
- Furthermore, *proper analysis of OR-gates* has to be performed, as analysis algorithms must handle both the aforementioned variants of OR-gates.

References

1. Ábrahám, E., Becker, B., Dehnert, C., Jansen, N., Katoen, J.P., Wimmer, R.: Counterexample generation for discrete-time markov models: An introductory survey. In: International School on Formal Methods for the Design of Computer, Communication and Software Systems. pp. 65–121. Springer (2014)
2. Al Amro, S., Chiclana, F., Elizondo, D.A.: Application of fuzzy logic in computer security and forensics. In: Computational Intelligence for Privacy and Security, pp. 35–49. Springer (2012)
3. Amorim, T., Schneider, D., Nguyen, V.Y., Schmittner, C., , Schoitsch, E.: Five major reasons why safety and security haven't married (yet). ERCIM News, vol. 102, pp. 16–17 (2015)
4. Apvrille, L., Roudier, Y.: Towards the model-driven engineering of secure yet safe embedded systems. arXiv preprint arXiv:1404.1985 (2014)
5. Arnold, F., Guck, D., Kumar, R., Stoelinga, M.: Sequential and Parallel Attack Tree Modelling. In: SAFECOMP. LNCS, vol. 9338, pp. 291–299. Springer International Publishing (2015)
6. Arroyo-Figueroa, G., Sucar, L.E.: A temporal bayesian network for diagnosis and prediction. arXiv preprint arXiv:1301.6675 (2013)
7. Ashok, P., Butkova, Y., Hermanns, H., Křetínský, J.: Continuous-time markov decisions based on partial exploration. In: International Symposium on Automated Technology for Verification and Analysis. pp. 317–334. Springer (2018)

8. Aslanyan, Z., Nielson, F., Parker, D.: Quantitative verification and synthesis of attack-defence scenarios. In: CSF. pp. 105–119. IEEE Computer Society (2016). <https://doi.org/10.1109/CSF.2016.15>
9. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.E.: Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Dependable Sec. Comput.*, vol. 1, pp. 11–33 (2004)
10. Baier, C., Katoen, J.P.: Principles of model checking. MIT press (2008)
11. Baier, C., Klein, J., Klüppelholz, S., Märcker, S.: Computing conditional probabilities in markovian models efficiently. In: International Conference on Tools and Algorithms for the Construction and Analysis of Systems. pp. 515–530. Springer (2014)
12. Barlow, R.E., Proschan, F.: Statistical theory of reliability and life testing: probability models. Intl. series in decision processes, Holt, Rinehart and Winston (1975)
13. Basagiannis, S., Katsaros, P., Pombortsis, A., Alexiou, N.: Probabilistic model checking for the quantification of DoS security threats. *Computers & Security* **28**(6), 450–465 (2009)
14. Baum, L.E., Petrie, T.: Statistical inference for probabilistic functions of finite state markov chains. *The annals of mathematical statistics* **37**(6), 1554–1563 (1966)
15. Bernsmed, K., Frøystad, C., Meland, P.H., Nesheim, D.A., Rødseth, Ø.J.: Visualizing cyber security risks with bow-tie diagrams. In: International Workshop on Graphical Models for Security. pp. 38–56. Springer (2017)
16. Bertrand, N., Schewe, S.: Playing optimally on timed automata with random delays. In: International Conference on Formal Modeling and Analysis of Timed Systems. pp. 43–58. Springer (2012)
17. Bianco, A., De Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: International Conference on Foundations of Software Technology and Theoretical Computer Science. pp. 499–513. Springer (1995)
18. Bobbio, A., Egidi, L., Terruggia, R.: A methodology for qualitative/quantitative analysis of weighted attack trees. *IFAC* **46**(22), 133–138 (2013)
19. Bossuat, A., Kordy, B.: Evil twins: Handling repetitions in attack–defense trees. In: *GramSec. LNCS*, vol. 10744, pp. 17–37. Springer International Publishing (2018)
20. Bouyer, P., Markey, N., Sankur, O.: Robustness in timed automata. In: *Lecture Notes in Computer Science*, pp. 1–18. Springer Berlin Heidelberg (2013)
21. Bozzano, M., Cimatti, A., Katoen, J.P., Nguyen, V.Y., Noll, T., Roveri, M.: Safety, dependability and performance analysis of extended aadl models. *The Computer Journal* **54**(5), 754–775 (2011)
22. Bozzano, M., Cimatti, A., Katoen, J., Nguyen, V.Y., Noll, T., Roveri, M., Wimmer, R.: A model checker for AADL. In: Touili, T., Cook, B., Jackson, P.B. (eds.) *CAV*. vol. 6174, pp. 562–565. Springer (2010)
23. Brázdil, T., Korenčíak, L., Krčál, J., Křetínský, J., Řehák, V.: On time-average limits in deterministic and stochastic petri nets. In: *Proceedings of the 4th ACM/SPEC International Conference on Performance Engineering*. pp. 421–422 (2013)
24. Brázdil, T., Krčál, J., Křetínský, J., Řehák, V.: Fixed-delay events in generalized semi-markov processes revisited. In: *International Conference on Concurrency Theory*. pp. 140–155. Springer (2011)

25. Brocke, J., Simons, A., Niehaves, B., Riemer, K., Plattfaut, R., Cleven, A.: Reconstructing the giant: On the importance of rigour in documenting the literature search process. In: ECIS (2009)
26. Bryant, R.E.: Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers* **C-35**(8), 677–691 (1986). <https://doi.org/10.1109/TC.1986.1676819>
27. Budde, C., Stoelinga, M.: Efficient algorithms for quantitative attack tree analysis. In: 42nd IEEE Computer Security Foundations Symposium CSF. IEEE (2021)
28. Budde, C.E., Kolb, C., Stoelinga, M.: Attack trees vs. fault trees: two sides of the same coin from different currencies. In: QEST (to appear)
29. Budde, C.E., Stoelinga, M.: Efficient algorithms for attack tree analysis. In: IEEE Computer Security Foundations Symposium CSF (2021)
30. Byres, E.J., Franz, M., Miller, D.: The use of attack trees in assessing vulnerabilities in scada systems. In: Proceedings of the international infrastructure survivability workshop. pp. 3–10. Citeseer (2004)
31. Calinescu, R., Češka, M., Gerasimou, S., Kwiatkowska, M., Paoletti, N.: Efficient synthesis of robust models for stochastic systems. *Journal of Systems and Software* **143**, 140–158 (2018)
32. das Chagas Moura, M., Zio, E., Lins, I.D., Droguett, E.: Failure and reliability prediction by support vector machines regression of time series data. *Reliability Engineering & System Safety* **96**(11), 1527–1534 (2011)
33. Chockalingam, S., Hadžiosmanović, D., Pieters, W., Teixeira, A., van Gelder, P.: Integrated safety and security risk assessment methods: A survey of key characteristics and applications. *LNCS* **10242**, 50–62 (2017)
34. Chockalingam, S., Pieters, W., Teixeira, A., van Gelder, P.: Bayesian network models in cyber security: a systematic review. In: Nordic Conference on Secure IT Systems. pp. 105–122. Springer (2017)
35. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: International Conference on Computer Aided Verification. pp. 154–169. Springer (2000)
36. Clifton, E., et al.: Fault tree analysis—a history. In: Proceedings of the 17th International Systems Safety Conference. pp. 1–9 (1999)
37. Commission, I.E., et al.: Iec 61025: Fault tree analysis. IEC Standards (2006)
38. Dalton, Mills, Colombi, Raines: Analyzing attack trees using generalized stochastic Petri nets. In: 2006 IEEE Information Assurance Workshop. pp. 116–123 (2006)
39. De Alfaro, L., Roy, P.: Magnifying-lens abstraction for markov decision processes. In: International Conference on Computer Aided Verification. pp. 325–338. Springer (2007)
40. Dempster, A.P.: A generalization of bayesian inference. In: Classic Works of the Dempster-Shafer Theory of Belief Functions, pp. 73–104. Springer Berlin Heidelberg (2008)
41. Dissaux, P., Singhoff, F., Lemarchand, L., Tran, H., Atchadam, I.: Combined real-time, safety and security model analysis. In: ERTSS (2020)
42. Dolev, D., Yao, A.: On the security of public key protocols. *IEEE Transactions on information theory* **29**(2), 198–208 (1983)
43. Dugan, J.B., Bavuso, S.J., Boyd, M.A.: Dynamic fault-tree models for fault-tolerant computer systems. *IEEE Transactions on Reliability* **41**(3), 363–377 (1992)

44. Dugan, J.B., Bavuso, S.J., Boyd, M.A.: Fault trees and markov models for reliability analysis of fault-tolerant digital systems. *Reliability Engineering & System Safety* **39**(3), 291–307 (1993)
45. Dullerud, G.E., Paganini, F.: *A course in robust control theory: a convex approach*, vol. 36. Springer Science & Business Media (2013)
46. D’Argenio, P.R., Jeannot, B., Jensen, H.E., Larsen, K.G.: Reduction and refinement strategies for probabilistic analysis. In: *Joint International Workshop von Process Algebra and Probabilistic Methods, Performance Modeling and Verification*. pp. 57–76. Springer (2002)
47. Finkelstein, M.: *Failure rate modelling for reliability and risk*. Springer Science & Business Media (2008)
48. Fishman, G.: *Monte Carlo: Concepts, Algorithms, and Applications*. Springer Series in Operations Research and Financial Engineering (1996)
49. Fovino, I.N., Masera, M., De Cian, A.: Integrating cyber attacks within fault trees. *Reliability Engineering & System Safety* **94**(9), 1394–1402 (2009)
50. Fraile, M., Ford, M., Gadyatskaya, O., Kumar, R., Stoelinga, M., Trujillo-Rasua, R.: Using attack-defense trees to analyze threats and countermeasures in an ATM: A case study. In: *PoEM. LNCS*, vol. 267, pp. 326–334. Springer (2016)
51. Fraile, M., Ford, M., Gadyatskaya, O., Kumar, R., Stoelinga, M., Trujillo-Rasua, R.: Using attack-defense trees to analyze threats and countermeasures in an atm: a case study. In: *IFIP Working Conference on The Practice of Enterprise Modeling*. pp. 326–334. Springer (2016)
52. Friedberg, I., McLaughlin, K., Smith, P., Laverty, D., Sezer, S.: Stpa-safesec: Safety and security analysis for cyber-physical systems. *Journal of information security and applications* **34**, 183–196 (2017)
53. Gadyatskaya, O., Hansen, R.R., Larsen, K.G., Legay, A., Olesen, M.C., Poulsen, D.B.: Modelling attack-defense trees using timed automata. In: *International Conference on Formal Modeling and Analysis of Timed Systems*. pp. 35–50. Springer (2016)
54. Gerhold, M., Hartmanns, A., Stoelinga, M.: Model-based testing of stochastically timed systems. *Innovations in Systems and Software Engineering* **15**(3), 207–233 (Sep 2019). <https://doi.org/10.1007/s11334-019-00349-z>, <https://doi.org/10.1007/s11334-019-00349-z>
55. German, R.: Markov regenerative stochastic petri nets with general execution policies: supplementary variable analysis and a prototype tool. *Performance Evaluation* **39**(1-4), 165–188 (2000)
56. Gupta, C., Jain, A., Joshi, N.: Fuzzy logic in natural language processing—a closer view. *Procedia computer science* **132**, 1375–1384 (2018)
57. Hartmanns, A., Hermanns, H.: In the quantitative automata zoo. *Science of Computer Programming* **112**, 3–23 (2015)
58. Hazewinkel, M.: *Bayesian approach to statistical problems* (1995)
59. Hermanns, H., Krämer, J., Krčál, J., Stoelinga, M.: The value of Attack-Defence Diagrams. In: *POST. LNCS*, vol. 9635, pp. 163–185. Springer Berlin Heidelberg (2016). https://doi.org/10.1007/978-3-662-49635-0_9
60. Hermanns, H., Krämer, J., Krčál, J., Stoelinga, M.: The value of attack-defence diagrams. In: *International Conference on Principles of Security and Trust*. pp. 163–185. Springer (2016)
61. Herzog, M.A., Marwala, T., Heyns, P.S.: Machine and component residual life estimation through the application of neural networks. *Reliability Engineering & System Safety* **94**(2), 479–489 (2009)

62. Hierons, R.M., Núñez, M.: Using schedulers to test probabilistic distributed systems. *Formal Aspects of Computing* **24**(4), 679–699 (2012)
63. Huth, M.: Possibilistic and probabilistic abstraction-based model checking. In: *Joint International Workshop von Process Algebra and Probabilistic Methods, Performance Modeling and Verification*. pp. 115–134. Springer (2002)
64. ISO/IEC 25010:2011, S., software engineering: Systems and software quality requirements and evaluation (square). *System and software quality models* (2011)
65. Jhawar, R., Kordy, B., Mauw, S., Radomirović, S., Trujillo-Rasua, R.: Attack Trees with Sequential Conjunction. In: *SEC. IFIPAICT*, vol. 455, pp. 339–353. Springer International Publishing (2015). https://doi.org/10.1007/978-3-319-18467-8_23
66. Junges, S., Guck, D., Katoen, J., Stoelinga, M.: Uncovering dynamic fault trees. In: *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. pp. 299–310 (2016)
67. Junges, S., Jansen, N., Dehnert, C., Topcu, U., Katoen, J.P.: Safety-constrained reinforcement learning for mdps. In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. pp. 130–146. Springer (2016)
68. Junges, S., Katoen, J.P., Stoelinga, M., Volk, M.: One net fits all. In: *International Conference on Applications and Theory of Petri Nets and Concurrency*. pp. 272–293. Springer (2018)
69. Jürgenson, A., Willemson, J.: Computing exact outcomes of multi-parameter attack trees. In: *OTM*. vol. 5332, pp. 1036–1051. Springer Berlin Heidelberg (2008)
70. Jürjens, J.: Umlsec: Extending uml for secure systems development. In: *International Conference on The Unified Modeling Language*. pp. 412–425. Springer (2002)
71. Kahneman, D.: A perspective on judgment and choice: mapping bounded rationality. *American psychologist* **58**(9), 697 (2003)
72. Katoen, J.P.: The probabilistic model checking landscape. In: *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science*. pp. 31–45 (2016)
73. Katz, J.: Bridging game theory and cryptography: Recent results and future directions. In: *Theory of Cryptography Conference*. pp. 251–272. Springer (2008)
74. Kimelman, D., Kimelman, M., Mandelin, D., Yellin, D.M.: Bayesian approaches to matching architectural diagrams. *Trans. Software Eng.* **36**(2), 248–274 (2010)
75. Klir, G.J., Yuan, B.: *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Prentice-Hall (1995)
76. Kordy, B., Mauw, S., Radomirović, S., Schweitzer, P.: Attack–defense trees. *Journal of Logic and Computation* **24**(1), 55–87 (06 2012)
77. Kordy, B., Piètre-Cambacédès, L., Schweitzer, P.: DAG-based attack and defense modeling: Don’t miss the forest for the attack trees. *Computer Science Review* **13–14**, 1–38 (2014)
78. Kordy, B., Wideł, W.: On quantitative analysis of attack–defense trees with repeated labels. In: *POST. LNCS*, vol. 10804, pp. 325–346. Springer (2018)
79. Korenčiak, L., Krčál, J., Řehák, V.: Dealing with zero density using piecewise phase-type approximation. In: *European Workshop on Performance Engineering*. pp. 119–134. Springer (2014)
80. Kornecki, A.J., Subramanian, N., Zalewski, J.: Studying interrelationships of safety and security for software assurance in cyber-physical systems: Approach based on bayesian belief networks. In: *2013 FedCSIS*. pp. 1393–1399. IEEE (2013)

81. Kriaa, S., Bouissou, M., Colin, F., Halgand, Y., Pietre-Cambacedes, L.: Safety and security interactions modeling using the bdmp formalism: case study of a pipeline. In: International Conference on Computer Safety, Reliability, and Security. pp. 326–341. Springer (2014)
82. Kriaa, S., Pietre-Cambacedes, L., Bouissou, M., Halgand, Y.: A survey of approaches combining safety and security for industrial control systems. *Reliability engineering & system safety* **139**, 156–178 (2015)
83. Kumar, R., Stoelinga, M.: Quantitative security and safety analysis with attack-fault trees. In: 2017 IEEE 18th International Symposium on High Assurance Systems Engineering (HASE). pp. 25–32 (2017)
84. Kumar, R., Ruijters, E., Stoelinga, M.: Quantitative Attack Tree Analysis via Priced Timed Automata. In: FORTE. LNCS, vol. 9268, pp. 156–171. Springer International Publishing (2015). https://doi.org/10.1007%2F978-3-319-22975-1_11
85. LeMay, E., Ford, M.D., Keefe, K., Sanders, W.H., Muehrcke, C.: Model-based security metrics using adversary view security evaluation (ADVISE). In: 2011 Eighth International Conference on Quantitative Evaluation of SysTems. pp. 191–200. IEEE (2011)
86. Letourneau, S., Famili, F., Matwin, S.: Data mining to predict aircraft component replacement. *IEEE Intelligent Systems and their Applications* **14**(6), 59–66 (1999)
87. Li, Y., Li, L.: Model checking of linear-time properties based on possibility measure. *IEEE Transactions on Fuzzy systems* **21**(5), 842–854 (2012)
88. Li, Y., Li, Y., Ma, Z.: Computation tree logic model checking based on possibility measures. *Fuzzy Sets and Systems* **262**, 44–59 (2015)
89. Lindemann, C., Thümmler, A.: Transient analysis of deterministic and stochastic petri nets with concurrent deterministic transitions. *Performance Evaluation* **36**, 35–54 (1999)
90. Liu, Y., Man, H.: Network vulnerability assessment using bayesian networks. In: Data mining, intrusion detection, information assurance, and data networks security 2005. vol. 5812, pp. 61–71. International Society for Optics and Photonics (2005)
91. Lye, K.w., Wing, J.M.: Game strategies in network security. *International Journal of Information Security* **4**(1), 71–86 (2005)
92. Manshaei, M.H., Zhu, Q., Alpcan, T., Başçar, T., Hubaux, J.P.: Game theory meets network security and privacy. *ACM Computing Surveys (CSUR)* **45**(3), 1–39 (2013)
93. Mauw, S., Oostdijk, M.: Foundations of Attack Trees. In: ICISC. LNCS, vol. 3935, pp. 186–198. Springer Berlin Heidelberg (2006)
94. Meedeniya, I., Moser, I., Aleti, A., Grunske, L.: Evaluating probabilistic models with uncertain model parameters. *Software & Systems Modeling* **13**(4), 1395–1415 (2014)
95. Modarres, M., Kaminskiy, M.P., Krivtsov, V.: Reliability engineering and risk analysis: a practical guide. CRC press (2016)
96. Murray, J.F., Hughes, G.F., Kreutz-Delgado, K., Schuurmans, D.: Machine learning methods for predicting failures in hard drives: A multiple-instance application. *Journal of Machine Learning Research* **6**(5) (2005)
97. Nicol, D.M., H.Sanders, W., Trivedi, K.S.: Model-based evaluation: From dependability to security. *IEEE Trans. Dep. Sec. Comput.*, vol. 1(1), pp. 48–65 (2004)
98. Nicol, D.M., Sanders, W.H., Trivedi, K.S.: Model-based evaluation: from dependability to security. *Trans. on dep. and sec. comp.* **1**(1), 48–65 (2004)
99. Nielsen, D.S.: The cause/consequence diagram method as a basis for quantitative accident analysis. Risø National Laboratory (1971)

100. Nigam, V., Pretschner, A., Ruess, H.: Model-based safety and security engineering (2019)
101. Novák, V., Perfilieva, I., Mockor, J.: Mathematical principles of fuzzy logic, vol. 517. Springer Science & Business Media (2012)
102. Oortwijn, W., Dijk, T.v., Pol, J.v.d.: Distributed binary decision diagrams for symbolic reachability. In: SPIN. pp. 21–30. ACM (2017). <https://doi.org/10.1145/3092282.3092284>
103. Organization, I.S.: Iso/dis 26262: Road vehicles, functional safety. Technical report (2009)
104. Pedroza, G., Apvrille, L., Knorreck, D.: Avatar: A sysml environment for the formal verification of safety and security properties. In: 2011 NOTERE. pp. 1–10. IEEE (2011)
105. Piètre-Cambacédès, L., Bouissou, M.: Modeling safety and security interdependencies with bdmp (boolean logic driven markov processes). In: 2010 IEEE International Conference on Systems, Man and Cybernetics. pp. 2852–2861. IEEE (2010)
106. Pita, J., Tambe, M., Kiekintveld, C., Cullen, S., Steigerwald, E.: Guards—innovative application of game theory for national airport security. In: Twenty-Second International Joint Conference on Artificial Intelligence (2011)
107. Popov, P.: Preliminary interdependency analysis (pia): method and tool support. In: International Workshop on Software Engineering for Resilient Systems. pp. 1–8. Springer (2011)
108. Rauzy, A.: New algorithms for fault trees analysis. *Reliability Engineering & System Safety* **40**(3), 203–211 (1993)
109. Rontidis, G., Panaousis, E., Laszka, A., Dagiuklas, T., Malacaria, P., Alpcan, T.: A game-theoretic approach for minimizing security risks in the internet-of-things. In: 2015 IEEE International Conference on Communication Workshop (ICCW). pp. 2639–2644. IEEE (2015)
110. Roth, M., Liggesmeyer, P.: Modeling and Analysis of Safety-Critical Cyber Physical Systems using State/Event Fault Trees. In: SAFECOMP 2013 (2013)
111. Roudier, Y., Apvrille, L.: Sysml-sec: A model driven approach for designing safe and secure systems. In: MODELSWARD
112. Rugina, A., Kanoun, K., Kaâniche, M.: An architecture-based dependability modeling framework using aadl. *ArXiv abs/0704.0865* (2007)
113. Ruijters, E., Guck, D., Drolenga, P., Stoelinga, M.: Fault maintenance trees: Reliability centered maintenance via statistical model checking. In: 2016 Annual Reliability and Maintainability Symposium (RAMS). pp. 1–6 (2016)
114. Ruijters, E., Guck, D., Drolenga, P., Peters, M., Stoelinga, M.: Maintenance analysis and optimization via statistical model checking. In: International Conference on Quantitative Evaluation of Systems. pp. 331–347. Springer (2016)
115. Ruijters, E., Guck, D., Drolenga, P., Stoelinga, M.: Fault maintenance trees: reliability centered maintenance via statistical model checking. In: RAMS. pp. 1–6. IEEE (2016)
116. Ruijters, E., Guck, D., Van Noort, M., Stoelinga, M.: Reliability-centered maintenance of the electrically insulated railway joint via fault tree analysis: a practical experience report. In: 2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). pp. 662–669. IEEE (2016)
117. Ruijters, E., Stoelinga, M.: Fault Tree Analysis: A survey of the state-of-the-art in modeling, analysis and tools. *Computer Science Review* **15–16**, 29–62 (2015)

118. Ruijters, E., Stoelinga, M.: Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools. *Computer Science Review* **15-16**, 29 – 62 (2015)
119. Saltelli, A., Chan, K., Scott, E.: *Sensitivity Analysis*. Wiley (2009), <https://books.google.nl/books?id=gOcePwAACAAJ>
120. Schneier, B.: Attack trees. *Dr. Dobb's journal* **24**(12), 21–29 (1999)
121. Shafer, G.: *A mathematical theory of evidence*, vol. 42. Princeton university press (1976)
122. Sommestad, T., Ekstedt, M., Holm, H.: The cyber security modeling language: A tool for assessing the vulnerability of enterprise system architectures. *IEEE Systems Journal* **7**(3), 363–373 (2012)
123. Sommestad, T., Ekstedt, M., Johnson, P.: Combining defense graphs and enterprise architecture models for security analysis. In: 2008 12th International IEEE Enterprise Distributed Object Computing Conference. pp. 349–355. IEEE (2008)
124. Song, L., Zhang, L., Godskesen, J.C.: Bisimulations and logical characterizations on continuous-time markov decision processes. In: International Conference on Verification, Model Checking, and Abstract Interpretation. pp. 98–117. Springer (2014)
125. Stamatelatos, M., Vesely, W., Dugan, J.B., Fragola, J., Minarick, J., Railsback, J.: *Fault Tree Handbook with Aerospace Applications*. Office of safety and mission assurance, NASA headquarters (2002)
126. Steiner, M., Liggesmeyer, P.: *Combination of safety and security analysis - finding security problems that threaten the safety of a system* (2016)
127. Wakileh, B., Gill, K.: Use of fuzzy logic in robotics. *Computers in industry* **10**(1), 35–46 (1988)
128. Wang, J.X., Roush, M.L.: *What every engineer should know about risk engineering and management*. CRC Press (2000)
129. Wang, Y., Wang, Y., Liu, J., Huang, Z., Xie, P.: A survey of game theoretic methods for cyber security. In: 2016 IEEE First International Conference on Data Science in Cyberspace (DSC). pp. 631–636. IEEE (2016)
130. Watson, H.: *Launch control safety study*. Tech. Rep. Sec. VII, Vol. 1, Bell Labs (1961)
131. Weiss, J.: A system security engineering process. In: Proceedings of the 14th National Computer Security Conference. *Information System Security: Requirements & Practices*, vol. 249, pp. 572–581 (1991)
132. Wideł, W., Audinot, M., Fila, B., Pinchinat, S.: Beyond 2014: Formal methods for attack tree-based security modeling. *ACM Comput. Surv.* **52**(4) (2019). <https://doi.org/10.1145/3331524>
133. Wilczyński, A., Jakóbiak, A., Kołodziej, J.: Stackelberg security games: Models, applications and computational aspects. *Journal of Telecommunications and Information Technology* (2016)
134. Young, W., Leveson, N.: Systems thinking for safety and security. In: Proceedings of the 29th Annual Computer Security Applications Conference. pp. 1–8 (2013)
135. Zampino, E.J.: Application of fault-tree analysis to troubleshooting the NASA GRC icing research tunnel. In: Annual Reliability and Maintainability Symposium. 2001 Proceedings. pp. 16–22 (2001)
136. Zhao, X., Robu, V., Flynn, D., Dinmohammadi, F., Fisher, M., Webster, M.: Probabilistic model checking of robots deployed in extreme environments. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8066–8074 (2019)

137. Zonouz, S.A., Khurana, H., Sanders, W.H., Yardley, T.M.: Rre: A game-theoretic intrusion response and recovery engine. *IEEE Transactions on Parallel and Distributed Systems* **25**(2), 395–406 (2013)
138. Åström, K.: Optimal control of markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications* **10**(1), 174–205 (1965). [https://doi.org/https://doi.org/10.1016/0022-247X\(65\)90154-X](https://doi.org/https://doi.org/10.1016/0022-247X(65)90154-X), <https://www.sciencedirect.com/science/article/pii/0022247X6590154X>