Motzek *et al. EURASIP Journal on Information Security* (2017) 2017:12
DOI 10.1186/s13635-017-0063-6

EURASIP Journal on
Information Security

RESEARCH Open Access

Selection of Pareto-efficient response plans based on financial and operational assessments

Alexander Motzek[2*], Gustavo Gonzalez-Granadillo[1], Hervé Debar[1], Joaquin Garcia-Alfaro[1] and Ralf Möller[2]

Abstract

Finding adequate responses to ongoing attacks on ICT systems is a pertinacious problem and requires assessments from different perpendicular viewpoints. However, current research focuses on reducing the impact of an attack irregardless of side effects caused by responses. In order to achieve a comprehensive yet accurate response to possible and ongoing attacks on a managed ICT system, we propose an approach that evaluates a response from two perpendicular perspectives: (1) A response financial impact assessment, considering the financial benefits of restoring and protecting potentially threatened operational capabilities while considering implementation and maintenance costs of responses. (2) A response operational impact assessment, which assesses potential impacts that efficient mitigation actions may inadvertently cause on the organization in an operational perspective, e.g., negative side effects of deploying mitigations. It is the key benefit of the presented approach to combine all obtained evaluations with a multi-dimensional optimization procedure such that a response plan is selected which reduces a state of risk below an admissible level while minimizing potential negative side effects of deliberately taken actions.

Keywords: Impact assessment, Dynamic response, Financial impact, Operational impact, Pareto-efficiency

1 Introduction

Finding adequate responses to ongoing attacks on information and communications technology (ICT) systems is a pertinacious problem and requires assessments from different perpendicular viewpoints. An adequate response to an ongoing attack or potential threat to a system must have the aim to minimize an associated degree of risk, minimize potential consequences of a successful attack, and to reduce the attack surface. However, most importantly, a chosen response plan shall not affect a mission inadvertently by itself, i.e., one shall not sacrifice a mission, e.g., a conglomeration of to-be-accomplished business processes, for a false sense of security.

Current research, however, focus on considering the impact of attacks [1–4], by evaluating their severity and consequences, but leave aside the potential impact of taken actions themselves onto higher goals. Moreover, the analysis of current cyber events should also consider the impact of potential mitigation actions as well as time, geographic space, and affected elements [5].

In this article, we present an approach to unify multiple perspectives on choosing adequate response plans from a set of response plans that have been proposed by some third-party entity, e.g., an operator or classical intrusion detection system. We consider a response plan as a collection of individual atomic actions called "mitigation actions," and we consider the likelihood of success of the detected attacks, their induced impact, their deployment and maintenance costs, as well as the consequences of these response plans onto a higher goal, e.g., success of a company or mission.

Our approach is based on two general impact assessments: (1) A financial impact assessment (FIA) and (2) an operational impact assessment (OIA), where both are applied to a cyber-defense domain for choosing adequate response plans. Effectively, one obtains a *response* FIA (RFIA) and a *response* OIA (ROIA).

RFIA relies on a return on response investment (RORI) index and a geometrical model (named attack volume) to estimate the impact of security incidents (e.g., intrusions,

*Correspondence: motzek@ifis.uni-luebeck.de
[2]Universität zu Lübeck, Institute of Information Systems, Ratzeburger Allee 160, 23562 Lübeck, Germany
Full list of author information is available at the end of the article

attacks, errors) in a financial perspective, and to rank responses accordingly. However, RFIA solely considers directly associated costs for deploying and maintaining a response, but does not consider their potential negative side effects onto a higher goal, e.g., loosing the ability to accomplish a business process. ROIA exactly considers that mitigation actions, while highly effective, could lead to such operational negative side effects inside the network and therefore onto a mission. ROIA evaluates proposed response plans based on validatable local impact- and dependency-assessments of dependencies inside an organization's business- and IT-infrastructure, but leaves aside the effectiveness of a response against an attack and neither considers associated implementation costs.

Both, RFIA and ROIA, establish impact assessments of proposed response plans, where a response plan represents a collection of multiple mitigation actions. Due to their nature, a financial impact assessments (using a RORI index) and an operational impact assessments are performed from perpendicular perspectives: On the one hand, the less invasive a response plan is, the less it can potentially cause collateral damage (probability of operational impact). On the other hand, a minimally invasive response plan will not significantly reduce a risk, i.e., will not yield a high return on response investment (RORI). It is the novel advantage of the presented approach of being able to combine both assessments using a multi-dimensional optimization procedure to find a "best-compromise" identified by a Pareto-efficient set.

The contributions on this paper are summarized as follows: We present a two-fold evaluation strategy for response plans such that a mission is protected from adversarial threats while not sacrificing the mission for a false sense of security. The strategy considers perpendicular perspectives: (1) A financial impact assessment of mitigation actions based on a cost-sensitive metric and geometrical models, and (2) an operational impact assessment of mitigation actions based on mission and resource dependency models. The combination of (1) and (2) yields vital benefits for choosing adequate responses in, e.g., critical infrastructures, and overcomes multiple discrepancies of related work, e.g., the neglect of negative side effect of responses. Further, we present multiple approaches to automatically obtain required models, e.g., by learning models from network traffic, and present practical implementation proposals for efficient evaluations and selections of responses. We validate and verify the presented approach with a real-world use case from the industry using real data.

This article extends previous work [6] by an in-depth description of the response financial and operational impact assessments, extended and improved theory of multi-dimensional optimization, and approaches towards automatically learning required models through machine learning and interviews with experts. Moreover, we discuss possible conflicts among mitigation actions and the relation between financial and operational assessments that makes it possible to propose a Pareto-efficient response plan. Furthermore, we discuss and compare various other optimization strategies.

Paper organization — Sections 2 and 3 discuss preliminaries and theory of a financial [7–9] and operational [10–12] impact assessment. Section 4 describes practical approaches to obtain required models and discusses an efficient implementation of evaluations. Section 5 presents the novel contribution of this article as a multidimensional optimization for selecting adequate response plans based on a combination of both impact assessment and the use of Pareto-efficiency. A real-world application of the presented approach is presented in Section 6 showing the applicability of the proposed models and approach. Section 7 discusses related work. In Section 8, we critically discuss our work, discuss potential conflicts among response plans, relate the financial and operational impact assessment to each other, and compare our approach to other selection strategies. Conclusions and perspectives for future work are presented in Section 9.

2 Financial impact assessment

Cost-sensitive metrics have been proposed as a viable approach to find an optimal balance between intrusion damages and the cost of implementing and maintaining a response over a period of time. Commonly, they guarantee the choice of the most appropriate response without sacrificing the system functionalities to an adversary, but do not consider the self-inflicted side effects of the responses themselves. Such cost-sensitive measurements are either absolute or relative: *absolute* measurements use precise values that scale with a given unit (e.g., hundreds, thousands, millions); whereas relative measurements are methods for deriving ratio scales from paired comparisons represented by absolute numbers [13]. *Relative* measurements are useful in obtaining an overall ratio scale ranking of the alternatives. If the ratio produces repeatable and consistent results, the model can be used to compare security solutions based on relative values [14]. Examples of these models include the return on investment (ROI) and all its variants [14–16].

2.1 Return on investment (ROI)

The simplest and most used approach for evaluating financial consequences of business investments, decisions, and/or actions is the return on investment (ROI). The ROI index is a *relative* measure that compares the benefits versus the costs obtained for a given investment [15, 17]. Informally, ROI basically shows how much a company earns from invested money. This metric supports decision makers to select the option(s) that have the highest return. ROI is calculated as the present value of accumulated net benefits over a certain time period minus the initial costs of investment, then divided by the initial costs of investment, as shown in Eq. 1.

$$ROI = \frac{B_t - C_t}{C_t} \cdot 100\,, \tag{1}$$

where B_t refers to all benefits during period t, and C_t refers to all costs during period t. The decision rule is that the higher the ROI value, the more interesting the investment.

2.2 Return on security investment

Return on security investment (ROSI) is a relative metric that compares the differences between the damages originated by attacks (with and without mitigation) against the cost of the mitigation action [14, 16, 18]. ROSI has been adapted from the ROI metric as presented in Eq. 2.

$$ROSI = \frac{(ALE_b - ALE_a) - Cost_{MA}}{Cost_{MA}} \cdot 100\,, \tag{2}$$

where ALE_b refers to the annual loss expectancy before mitigation, ALE_a refers to the annual loss expectancy after mitigation, and $Cost_{MA}$ is the cost of the mitigation action.

Similar to the ROI metric, the decision rule is that the higher the ROSI value, the more interesting the investment.

2.3 Return on response investment

Return on response investment (RORI) is a quantitative model for ranking response plans by a cost-sensitive financial comparison and has been introduced in [7–9, 19], and in the following. For the scope of this article, we consider sets of individual actions performed as a response to an adversary. Sets of these actions are hereinafter called response plans:

Definition 2.1 (*Response plan*) *A response plan RP is a set of mitigation actions, representing individual actions to be performed as a response to an adversary or threat opposed to an organization.*

RORI is an adaptation of the return on security investment (ROSI) index for a comparison of response plans. The RORI index considers the intrusion impact and direct financial costs as shown in Eq. 3.

$$RORI = \frac{(ALE \cdot RM) - ARC}{ARC + AIV} \cdot 100\,. \tag{3}$$

For every response plan, a RORI index may be calculated for a given attack scenario using Eq. 3. All parameters are defined as follows:

Definition 2.2 (*Annual loss expectancy, ALE*) *ALE corresponds to the attack impact loss that an organization is exposed to in the absence of mitigation actions. ALE is expressed in monetary values (e.g., \$/year) and depends directly on the attack's severity and likelihood. ALE includes the loss of assets (L_a), the loss of data (L_d), the loss of reputation (L_r), the legal procedures (LP), the loss of revenues from clients or customers (L_{rc}), as well as other losses (L_o), contracted insurances (Ins), to be multiplied by the annual rate of occurrence of the attack (ARO), as shown in Eq.* 4*,*

$$ALE = (L_a + L_d + L_r + LP + L_{rc} + L_o - Ins) \cdot ARO\,. \tag{4}$$

Definition 2.3 (*Annual infrastructure value, AIV*) *AIV represents the fixed costs that are expected to be perceived by an organization regardless of the deployed response. AIV is strictly positive and is expressed in monetary values (e.g., \$/year). It includes the following costs: equipment costs (C_e), personnel costs (C_p), service costs (C_s), and other costs (C_o), as well as the resell value (V_r), as shown in Eq.* 5*,*

$$AIV = C_e + C_p + C_s + C_o + V_r\,. \tag{5}$$

Definition 2.4 *[Risk mitigation, RM] RM refers to the risk mitigation associated with a given mitigation action. RM takes values between* 0 *and* 100% *(i.e.,* 0% $\leq$ *RM* $\leq$ 100%*). In the absence of mitigation actions, RM equals* 0%*. RM is computed as the product of the mitigation coverage (COV, which is the percentage of the attack covered by the mitigation action) by the effectiveness factor (EF, which is the percentage of reduction of the total incident cost given the enforcement of the mitigation action), as shown in Eq.* 6*,*

$$RM = COV \cdot EF\,. \tag{6}$$

Definition 2.5 (*Annual response cost, ARC*) *ARC refers to the costs associated to a given mitigation action. ARC is always positive and expressed in monetary values (e.g., \$/year). It includes direct costs such as the cost of implementation (C_{impl}), the cost of maintenance (C_{maint}), as well as other direct costs (C_{od}) that may originate from the adoption of a particular mitigation action, as shown in Eq.* 7*,*

$$ARC = C_{impl} + C_{maint} + C_{od}\,. \tag{7}$$

Considering a RORI index alone, the best candidate response set is represented by a maximal positive RORI index. Therefore, the RORI index of a response plan yields a ranking of response plans for, e.g., an ongoing attack, while considering the cost of implementing and maintaining this response (ARC), the expected monetary loss in the case of a successful attack (ALE), the absolute values of the protected good (AIV), and, obviously, a degree of effectiveness of the response against the attack (RM).

However, what a RORI index does not consider is a degree of self-inflicted side effects of the responses themselves, e.g., the severe issue of shutting down an extremely important central control server in order to "protect" it from an adversary.

3 Operational impact assessment

Operational impact assessment (OIA) is used to address potential impacts onto a higher goal, from widespread events which impact local operational capabilities. For example, a local impact caused by an event on a distant node might lead to a causal chain of operational failures, leading to an impact on a company. In this work, we utilize OIA and its associated property of being able to consider "spreads of impacts" through an infrastructure to address the negative side effects of responses onto a company or mission. For example, the shutdown of a central control server will certainly lead to a high probability that a central business process is not accomplishable anymore, which will certainly lead to the fact that the company is *impacted as well*. Likewise, a shutdown, patch, exchange, or maintenance of a server that supports the central control server may induce the same (negative) causal chain of events.

Motzek et al. introduce an approach towards OIA based on a probabilistic graphical model in [11], which defines a well-understood problem onto which an OIA can be reduced. To obtain a consistent model, various experts must be consulted, and information origins from various different perspectives coming from different expertises. In effect, experts, i.e., models, will contain disagreements and contain semantic and technology gaps [11]. By resorting to a probabilistic model, the use of conditional probability distributions allows for local views on assessments, without a need to understand a specific use case nor any algorithmic properties. Therefore, OIA introduced by [11] is able to directly include all views by experts directly into one consistent model, bridging semantic and technology gaps. This is highly beneficial, as experts are not forced to come one bad compromise, but the model is able to understand emerging disagreements.

Moreover, local views allow one to validate data instead of results. This means that all parts of the model per se are understandable and validatable by using common sense or by validating small subparts against ground truth. No large, holistic ground truth datasets are required to validate algorithms, as the emerging problem and associated algorithms to solve these problems are inherently defined by the model itself.

The following sections introduce views on OIA from three different perspectives, each defining one dependency model as a probabilistic graphical model of random variables and respective dependencies.

Remark 3.1 (Impact) *The abstract term "impact" is used in this work in the sense of "not operating as fully intended." The underlying meaning of "intended operation" lies in a use case of the model.*

3.1 Mission dependency model (business view)

Motzek et al. [11] extend a model by Jakobson [20] and model mission dependencies as a graph of *mission nodes* (MN) as shown in Fig. 1. A *company* is dependent on its *business processes*. A business process is dependent on one or more *business functions*, which are provided by *business resources*. Figure 1 shows a dependency graph of business relevant objects for a small company consisting of two business processes, requiring a total of four functions provided by four resources.

Dependencies are represented by local conditional probability distributions (CPDs) modeling probabilities of failure, given dependances fail. For example, the probability of business function BF_1 (see Fig. 1), say, "provide access to customer data", failing, given required business resource A, e.g., "customer-data-frontend", fails is 90%. [10] argues that the meaning of local conditional probabilities are understandable using common sense (e.g., "in 9 out of 10 cases, customer data were not accessible for employees during frontend-server maintenance") and that the (numerical) assessment can be directly validated by either an expert or through ground-truth.

Definition 3.2 (*Probabilistic preliminaries*) *A node of a probabilistic dependency model is a random variable, denoted with capital letters, e.g., X. A random variable X is assignable to one of its possible values x* ∈ *dom*(*X*). *Let*

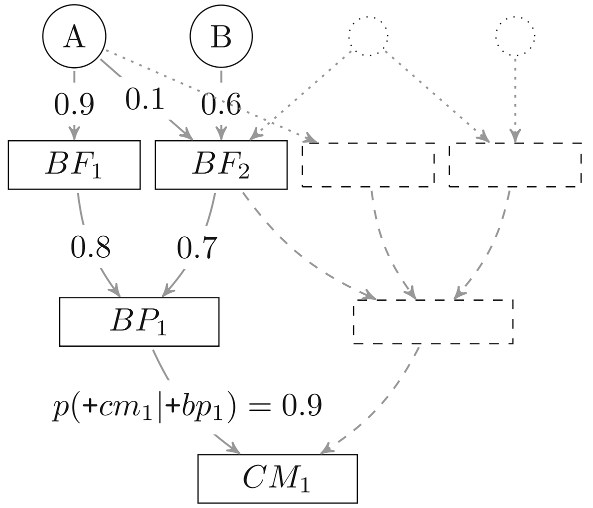

Fig. 1 Mission dependency model. *Values along edges* denote individual conditional probability fragments

$P(X = x)$ denote the probability of random variable X having x as a value. For the case $dom(X) = \{true, false\}$, we write ^+x for the event $X = true$ and ^-x for $X = false$.

The event ^+x represents the case that node X is operationally *impacted* and ^-x that is operating as fully intended, i.e., no impact is present.

Definition 3.3 (*From dependencies to distributions*) *Single dependencies of a random variable Y on X are modeled as individual conditional probability $p(x|y)$ and $p(x|\neg y)$. Such individual conditional probabilities are fragments of a complete CPD and are therefore denoted in lowercase. To acquire the local CPD $P(X|\vec{Y})$ of node X from all its fragments $p(X|Y)$ of all dependent nodes $Y \in \vec{Y}$,* [10] *employs a non-leaky noisy-or combination function as described in* [21, 22]. *Non-leakiness implies $p(^+x|^-y) = 0$ for every dependency and therefore $P(^+x|^-\vec{y}) = 0$. Non-leaky implies that an operational impact cannot origin "from nowhere" and there must be a cause for it inside the network.*

Using individual conditional probability fragments leads to a significantly easier design of mission dependency models, as no complete conditional probability distributions need to be designed. Per Definition 3.3, the complete conditional probability distributions are obtained via a deterministic combination function such as Noisy-OR. Other combination functions exist, which we briefly discuss in the following remark.

Remark 3.4 (Noisy-OR, Noisy-AND, Redundancy) *In this work, we solely consider a Noisy-OR model, i.e., every dependency can lead to a failure on its own. Notwithstanding, Noisy-AND is another possible modeling technique for certain resources, whose failure may only be provoked by a combined impact on all higher resources simultaneously, e.g., completely redundant systems. We define all models more generally in* [11], *and explicitly consider Noisy-AND cases. Still, we believe that a Noisy-OR assumption is often a safer approach: claiming that there exists absolutely no chance that a resource may be impacted given all, but one, dependencies are impacted, is a harsh assumption.*

Informally, a mission dependency model is a graph of nodes where every edge is associated with a probability value. A formal definition is given by [11] as follows, which provides valuable properties from a probabilistic perspective, which are discussed afterwards.

Definition 3.5 (*Mission dependency model*) *A mission dependency model M is a directed acyclic graph (DAG) as a pair $\langle \vec{V}, \vec{E} \rangle$ of vertices $\vec{V}$ and edges $\vec{E}$. Vertices $\vec{V}$ are random variables (Def. 3.2) and are categorized according to their semantic as business resources ($\vec{BR}$), functions ($\vec{BF}$), processes ($\vec{BP}$), and company (BC). For the scope of this work, we consider that a business dependency model is created for a single BC. The ordering $BR \prec BF \prec BP \prec BC$ represents the strict topological ordering of graph M. Every edge $E \in \vec{E}$ represents a dependency. Let $V \in \vec{V}$, then let $\vec{E}_V \subseteq \vec{E}$ be the set of edges directed to V, and let $\vec{D}_V$ be the set of vertices from which $\vec{E}_V$ origin, i.e., $\vec{D}_V$ is the set of dependencies of V. For every vertex $V \in \vec{V}$ a conditional probability distribution (CPD) $P(V|\vec{D}_V)$ is given, or, alternatively, a combination function is given for V and edges $E \in \vec{E}_V$ are associated with conditional probability fragments s.t. a $p(^+v|d)$ is given for all $d \in dom(D), \forall D \in \vec{D}_V$.*

With Definition 3.5, a mission dependency model is a Bayesian network, whose semantics is defined by the joint probability distribution over all mission nodes, i.e., random variables, as the product of all local defined CPDs. This simple definition of a global JPD for this mission dependency model is one of the most important aspects of this probabilistic approach to mission impact assessment: In the JPD, no global normalization factors are required. This means that every individual CPD and every local conditional probability fragment is interpretable individually and one does not require the complete big picture of other assessments to understand it. This property, i.e., the correspondence of a mission dependency model to a Bayesian network, is the mathematical foundation why the introduced mission impact assessment in [11] is context-free and bias-free, i.e., every parameter is immediately understandable without knowledge of other parts of the model, and obtained final mission impact assessments are directly understandable without requiring reference values, i.e., one cannot come to biased interpretation because of a dulling due to long-time seen results.

Business resources are part of an infrastructure perspective and—from an operational view—might be irrelevant, but are identified to be business critical by a business expert. Notwithstanding, such an assessment might be inaccurate, which is why transitive impacts must be considered. For example, a web-service might be identified as a business critical resource; it cannot be expected that an underlying distributed computing cluster is identified to actually provide this web-service. The following resource dependency model covers these dependencies.

Designing BPMN models is handled manually by an expert from a company or by an external business consultant having a precise expertise business analyses. Business analysis is performed on a pure business perspective and stops at a "resource" level. For example, a business analyst may identify a web-service, but will not describe the technical dependencies of the webservice on a database or a data center. This is a reasonable approach, as the latter perspective comes from a very different expertise and would require very broad-range experts. Further, an identification of a web-service as a business relevant object is causally more precise in the terms of an operation perspective: A failing database might cause a web-service to not operate as intended and therefore might lead to an unaccomplishment of a business-process. Still, the direct cause for the business-process being unaccomplishable is a rogue web-service and *not* the database. An "IT" expert might identify a web-service to be irrelevant, as the crucial point of failure or the point of interest lies in the availability of data from a database. Nevertheless, the latter dependencies must be covered and are discussed in the upcoming subsection. The main intention of using multiple models here is to reflect all views exactly and make the global model able to understand the discrepancies instead of enforcing a bad compromise.

3.2 Resource dependency model (operation view)

Critical resources identified in a mission dependency model are almost certainly *dependent* on further resources. For example, a web-server is likely dependent on a database-server. In effect, it is completely natural that multiple experts will disagree on the identification of critical devices. For example, a data analyst will solely be interested in data from the database, whereas the accountant will require data insights provided through the webserver. Essentially, no difference on the importance exists between two devices. No matter which one will be identified by a mission dependency model, the vast dependence between both must be covered, which is performed by the resource dependency model discussed here.

Formally, a resource dependency model is a probabilistic graphical model, where every resource represents a random variable, and a dependency is modeled as a conditional probability fragment as shown in Fig. 2. As before, the local and intuitive interpretation of those is preserved by following such a Bayesian approach.

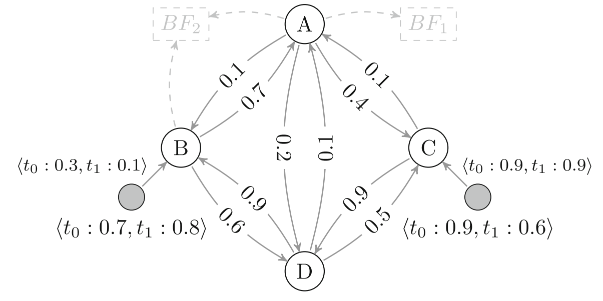
BF_2
A
BF_1
0.1
0.7
0.1
0.4
B
0.2
0.1
C
$\langle t_0 : 0.3, t_1 : 0.1\rangle$
$\langle t_0 : 0.9, t_1 : 0.9\rangle$
0.9
0.9
$\langle t_0 : 0.7, t_1 : 0.8\rangle$
0.6
0.5
$\langle t_0 : 0.9, t_1 : 0.6\rangle$
D

Fig. 2 A minimalistic resource dependency model. Conditional probability fragments are marked along the edges. *Gray nodes* represent external shock events leading to local impacts on resources. Connections to the mission dependency model are sketched in *dashed gray*

Definition 3.6 (*Resource dependency model*) *A resource dependency model R is a directed graph as a pair* $\langle \vec{V}, \vec{E} \rangle$ *of vertices* $\vec{V}$ *and edges* $\vec{E}$*. Every edge* $E \in \vec{E}$*, from vertex* $X \in \vec{V}$ *to* $Y \in \vec{V}$*, represents a dependency, and is associated with a conditional probability fragment* $p(^+y|^+x)$*. Vertices* $\vec{V}$ *are random variables (Def.* 3.2*) and represent resources in an infrastructure, where a subset of vertices semantically correspond to vertices of a corresponding mission dependency model M. Let* $V \in \vec{V}$*, then let* $\vec{E}_V \subseteq \vec{E}$ *be the set of edges directed to V, and let* $\vec{D}_V$ *be the set of vertices from which* $\vec{E}_V$ *origin, i.e.,* $\vec{D}_V$ *is the set of dependencies of V. For every vertex* $V \in \vec{V}$ *a conditional probability distribution (CPD)* $P(V|\vec{D}_V)$ *is defined by a non-leaky noisy-or combination of all conditional probability fragments of associated edges in* $\vec{E}_V$*. V is not contained in* $\vec{D}_V$*, i.e., a resource V is not dependent on itself.*

This definition is similar to the previous definition of a mission dependency model, but allows one to model cyclic dependencies, which will later be resolved during inference.

Considering impacts in a large infrastructure, where a resource is dependent on another, and a dependent resource is threatened, the identified critical resource might be threatened *transitively* as well. Effectively, an impact will "spread" throughout this model onto a mission dependency model which will later be covered during the impact assessment.

Resources in a resource dependency model may represent, e.g., individual ICT servers, ICS devices, software components or, in other use cases, manufacturing robots, suppliers, soldiers, or vehicles. The underlying meaning of resources in a resource dependency model depends on a use case similar to the meaning of an impact. These implications and mixtures between different types of resources are discussed in [11].

In large companies, resource dependency model may quickly grow and contain multiple hundreds of nodes with thousands of dependencies. Therefore, it is almost certain that a human expert will not be able to model such a resource dependency model manually by hand. Moreover, in novel, dynamically scaling environments, these infrastructure will quickly change over time and a model needs to be constantly adjusted. We introduce in [11] an approach to automatically learn resource dependency model constantly from network traffic analyses, which we discuss later in Section 4.2.1 and present results in Section 4.2.

Mission dependency models and resource dependency models already represent a probabilistic graphical model for a company, its business goals and its dependencies on its infrastructure. If a node inside this infrastructure becomes impacted, e.g., attacked, an impact might "spread" throughout the complete infrastructure and will eventually impact the company as well. To judge, i.e., to probabilistically correctly assess the probability of this eventuality, what is yet missing is a source for these potential impacts, which are addressed in the following subsection.

3.3 Local impacts (security view)

The third view on OIA involves a security expert with an expertise to analyze *local* consequences of events. In effect, the security expert does neither require knowledge of the complete infrastructure nor of the company's business goals. In the style of classical reliability analyses using Bayesian approaches, as, e.g., early work investigated by [23], every event potentially affecting, i.e., impacting, a node represents a so-called external shock event. Informally, an external shock event (SE) represents a source for an impact and threatens one or more nodes in a resource dependency model to become impacted.

Every SE represents a random variable, by which one obtains two vital benefits in modeling: (1) one is able to model uncertainty about the existence of an SE and (2) one is able to model uncertainty about the effect of an SE, i.e., a probability belief in the impact-degree of an existing SE. Note that both probabilities are assessable individually and locally. We define in [11] external shock events as random variables formally as follows.

Definition 3.7 (*External shock events*) *An external shock event SE is a random variable. Let $\vec{SE}$ be the set of all known external shock events. An external shock event $SE \in \vec{SE}$ might be present (^+se) or not be present ($\neg se$), for which a prior random distribution $P(SE)$ is defined, i.e., SE is a prior random variable. Every vertex V of a resource dependency model $\underline{R}$ might be affected by one or more external shock events $\vec{SE}_V \subseteq \vec{SE}$. In the case that an external shock event is present ($SE =^+ se, SE \in \vec{SE}_V$), there exists a probability of it affecting node V, expressed as a local conditional probability fragment $p(^+v|^+se)$. If an external shock event exists and it is not inhibited, we speak of a local impact on V. In the case that the external shock event is not present, i.e., $\neg se$, it does not affect random variable V and we write $p(^+v|\neg se) = 0$. Every individual conditional probability fragment from an external shock event is treated in the same noisy-or manner as a dependency towards another node, and thus, multiple shock events can affect one node and one shock event can affect multiple nodes.*

Using Definition 3.7, the presence of an SE can be known (observed) or can remain unclear. If the presence of an external shock event remains unclear, its existence is assessed probabilistically through its prior random distribution $P(SE)$. The prior random distribution $P(SE)$ can, for example, represent the uncertainty about whether a vulnerability may or may not be present on a system due to imprecise system configuration knowledge. We denote the set of observed external shock events (known presence) as a set of instantiations $\vec{se}_o$ of observed random variables $\vec{SE}_O \subseteq \vec{SE}$. If an SE, say SE_1 is known to be present (observed $^+se_1$ in $\vec{se}_o$) or the existence-case is considered during probabilistic assessment, there exists a probability that SE_1 will affect an associated node X to become impacted. This probability is defined by $p(^+x|^+se_1)$ and can, for example, represent the uncertainty about whether an actually present ($^+se_1$) actually leads to an impact on the node, as no exploits are present or the node is additionally protected against, e.g., buffer-overflows.

Both probabilities, $p(^+x|^+se_1)$ and $p(^+se_1)$, are likely to vary over time, as, e.g., access complexities for a vulnerability lower and exploits will eventually emerge and become public. To capture this form of time-dependence, we define a form of temporal aspects.

Definition 3.8 (*Temporal aspects*) *We define a temporal aspect of an external shock event. We employ the idea of abstract timeslices in which the effect of an external shock event changes. Every abstract timeslice then represents a duplicate of the network- and mission dependencies with a different set of local conditional probabilities and prior probabilities of local impacts. We denote time-varying probabilities in a sequence notation as $\langle t_0 : p_0, \ldots, t_T : p_T \rangle$, with $T + 1$ abstract timeslices. In every abstract timeslice i, varying local impacts take their respective conditional or prior probability p_i defined for its timeslice t_i.*

With Definition 3.8, an independent model is created for each timeslice, i.e., impact assessments of time t_i are independent of assessments from time t_{i-1}. For example, for three timeslices, one obtains three complete probabilistic graphical models, in which one obtains an OIA for each business entity. Motzek et al. [11] call this the "independent-timeslice model," as no connection between nodes from one timeslice to another exist. This idea is extended by them in [11] towards a fully dynamic impact assessment, where entities of timeslice t_i *depend* on entities of timeslice t_{i-1}, i.e., a resource dependency model is a time-dependent model evolving over time with time-dependent, "conscious" nodes allowing for retrospective and predictive analyses of potential mission impacts.

As mentioned earlier, every local impact represents a potential threat and can be, for example, a consequence of a present vulnerability, a mitigation action, a failure, or an attack. It lies in the expertise of a security operator to assess a potential *local* impact of those threats, for which we present examples in Section 4.2.2. Note that he does not need to have neither any expertise in resource dependencies nor an understanding of missions to do so. Further, an assessment of local impact probability can be formally validated through experiments or be grounded on commonsense.

To summarize, external shock events affect nodes in a resource dependency model, where complex interdependencies lead to a "spread" of impacts until even business entities modeled by a mission dependency model become impacted. This assessment is based on a well-defined probabilistic inference problem, which is discussed in the following section.

3.4 Mathematical mission impact assessment

To summarize, one probabilistic graphical model is defined by a mission dependency network, a resource dependency network, and a set of external shock events with associated local impacts threatening nodes (or random variables) defined by the resource dependency network. As resource nodes are dependent on each other, a threatened node might again threaten another node, which leads to a global "spreading" of impacts induced by external shock events. In the end, there exists a probability that even a business process or the complete modeled company (mission) is threatened transitively by various external shock events, which is what we call the mission impact assessment.

We believe that a manually created mission dependency model will be constant over long periods of time, as a company will not change essentially. As a resource dependency model is automatically and constantly learned, the complete approach will adapt constantly to context drifts and will remain valid over long periods of time with very low manual workload during operation.

As mentioned before, a probabilistic approach is followed, in which every dependency is modeled as an individual local conditional probability. Every threat, i.e., local impact, is modeled as an instantiated or observed prior random variable (an external shock event). An (global) impact on a mission node X is equivalent to the conditional probability of the node being impacted (^{+}x) given all observed external shock events. Motzek et al. [11] formally define this probability as the mission impact assessment.

Definition 3.9 (*Mission impact assessment, MIA*) *Given a mission dependency model M, a resource dependency model R and a set of external shock events $\vec{SE}$, a mission impact assessment of a mission node MN is defined as the conditional probability of a mission node $MN \in M$ being impacted (^{+}mn), given all observed external shock events $\vec{se}_o$, i.e., $P(^{+}mn|\vec{se}_o)$, where the effects of local impacts due to all $\vec{SE}$ are mapped globally based on mission-dependency and resource-dependency graphs. Note that $\vec{se}_o$ includes present (^{+}se) and absent ($\neg se$) shock events and that some shock events are unobserved, i.e., are assessed probabilistically through their prior random distribution $P(SE)$. The task of obtaining $P(^{+}mn|\vec{se}_o)$ is defined as the MIA problem.*

Given Definition 3.9, it is the task of a mission impact assessment to solve the MIA problem, i.e., to obtain the probability $P(^{+}mn|\vec{se}_o)$. Probabilistic inference is generally known to be NP-hard, and exact solutions to MIA problems are only obtainable in small toy domains. However, approximate inference techniques are a valuable alternative for probabilistic inference. To obtain an algorithm determining an approximate solutions to the MIA problem, we show in [11] that the probabilistic model is a probabilistic logic program, where every "path" $w_i^{MN} \in \vec{w}^{MN}$ from an external shock event $SE \in \vec{SE}$ to the mission node MN is a conjunction of Boolean random variables and is a sufficient proof for satisfying $\{MN = true\} = ^{+}mn$. Due to the noisy-or assumptions, $\vec{w}^{MN}$ then represents a disjunction of conjunctions. Every proof w_i^{MN} exists with a probability $P(w_i^{MN})$, where $P(w_i^{MN})$ is the product of all probabilities in this proof. Let $\mathbf{P}(w_i^{MN})$ denote the probability viewed as a set. $P(^{+}mn|\vec{se}_o)$ is then the probability that at least one proof holds, or rather, the probability that the disjunction of conjunctions is satisfied, i.e.,

$$\begin{aligned} P\left(^{+}mn|\vec{se}_o\right) &= \bigcup_i \mathbf{P}\left(w_i^{MN}\right) = P\left(\vec{w}^{MN}\right) \\ &= P\left(\left\{\bigvee_i w_i^{MN}\right\}\right) , \end{aligned}$$

where not all $\mathbf{P}(w_i^{MN})$ are disjoint. Calculating $\bigcup_i \mathbf{P}(w_i^{MN})$ is also known as the probabilistic satisfaction problem and is also used in the Problog reasoning framework [24]. To reduce computational complexity, a search for all "paths" $w_i^{MN} \in \vec{w}^{MN}$ can be limited to a fixed depth, e.g., using a depth-limited depth-first search. It is reasonable to limit a depth to an average path length in a graph to at least visit every node once, i.e., to at least include every external shock event once.

A more detailed evaluation, derivation, and validation of this reduction and associated approximation algorithms are given by us in [11].

A probabilistic MIA $P(^{+}mn|\vec{se}_o)$ directly originates from all defined dependency-models and represents an inference problem in a probabilistic graphical model. Therefore, we argue in [11] that if locally defined dependency-models are validated to be correct, an obtained impact assessment $P(^{+}mn|\vec{se}_o)$ is validated, too.

4 Model acquisition, learning, and evaluation

The operational impact assessment, as well as the financial impact assessment, requires various models and parameters in order to yield a respective *response* OIA and *response* FIA. In this section, we discuss multiple approaches to obtain these models, i.e., to specify them from manual interviews with experts and to automatically learn them from data analyses using machine learning. Moreover, we discuss implementation details for efficient inference on impact assessments. The latter is highly important for the ROIA as it is based on a computationally expensive probabilistic inference problem.

4.1 Response financial impact assessments

Response financial impact assessments (RFIA) quantifies the level of benefit perceived per response plan on a financial basis. The main goal of RFIA is to calculate a RORI index considering various implementation and maintenance costs, and the effectiveness of a response against an attack. In order to obtain these parameters, expert interviews and operations on geometric models are utilized as discussed in the following subsections.

4.1.1 Return on response investment

Return on response investment (RORI) is a relative index that indicates the level of benefit perceived if a given mitigation action is implemented. Required parameters for obtaining RORI estimates by Eq. 3 include two kinds of parameters: (1) fixed parameter and (3) variable parameters.

The fixed parameter includes the annual loss expectancy (ALE), which characterizes the intrusion or attack and is directly acquirable from expert knowledge through interviews, as performed by us and presented in Section 6.

Variable parameters include (i) the annual infrastructure value (AIV), which depends on the system, (ii) the risk mitigation (RM), and (iii) the annual response cost (ARC) which expresses the costs related to a mitigation action. The RM parameter differs for every mitigation action type which may be present in a response plan, and represent the cost of installation, deployment and maintenance, but leave aside their transitive cost-impacts onto the company through negative side effects (which are assessed through ROIA). Therefore, RM parameters are obtainable from expert interviews and manufacturer information. To obtain an estimate for a risk mitigation (RM) of each mitigation type, a geometrical attack-volume model is used, as discussed in the following section.

4.1.2 Attack volume (AV)

The attack volume model is a geometrical model for evaluating the impact of one or multiple attacks and/or mitigation actions over a specific target. The representation of each attack is performed in a three-dimensional coordinate system, i.e., user account (Acc), channel (Ip-Port), and resource (Res). The same coordinates include also system assets and potential mitigation actions. The projection of the three axis in our coordinate system generates geometrical instances in three dimensions. The resulting volume is computed as the product of the axes contribution to the execution of the incident, i.e.,

$$AV(A) = Co_{Res}(A) \cdot Co_{Ip-Port}(A) \cdot Co_{Acc}(A). \tag{8}$$

The axis contribution is determined as the sum of the product of each set of axis category (e.g., user account type, port class, resource type, etc.) by its associated weighting factor. Each category within the axis contributes differently to the volume calculation. The weighting factor corresponds to the severity of a given category based on the CARVER methodology [25]. This latter assigns an appropriate weight to each entity composing the axes in our coordinate system based on multiple criteria (i.e., criticality, accessibility, recuperability, vulnerability, effect, and recognizability). CARVER assigns numerical values on a scale of 1 to 10 to each considered factor and places them in a decision matrix. The sum of the values indicates the priority of a given entity in an information system.

The volume calculation requires the computation of the contribution of each axis represented in the coordinate system. This contribution is determined as the sum of each set of axis entities (e.g., user account type, port class, resource type) times its associated weighting factor (that results from the implementation of the CARVER methodology), i.e.,

$$\begin{aligned} Co_{Axis}(A) = \sum_{i=0}^{n} & Count(E \in Type_{Axis}(A)) \\ & \times WF(Type_{Axis}(A)) \,. \end{aligned}$$

The attack volume yields a three-dimensional representation of a complete attack scenario, making it possible to calculate the impact of multiple security entities (e.g., system, attack, countermeasure) that originate simultaneously in the system. Furthermore, we use geometrical operations to compute the union and/or intersection of multiple volumes, making it possible to determine the impact of multiple attacks arriving simultaneously on the system and the effects of implementing multiple countermeasures as a reaction strategy. As such, we are able to compute the coverage of individual and combined attacks in the system, and the level of coverage for one or more response plans against the detected attack(s). For instance, considering that attack, A_1 affects resources R1:R3 (WF=5), channels Ch1:Ch3 (WF=3), and users U1:U3 (WF=2), the attack volume is equivalent to $(AV(A_1) = (3 \cdot 5) \cdot (3 \cdot 3) \cdot (3 \cdot 2) = 810\text{units}^3)$;

and response plan RP_1 protects resources R2:R5, channels Ch2:Ch5, and users U2:U5. Therefore, the resources elements that are covered by RP_1 respect to A_1 are the following: R2:R3, Ch2:Ch3, and U2:U3. The coverage volume of RP_1 with respect to A_1, i.e., the volume of $A_1^{RP_1}$, is therefore equivalent to

$$AV\left(A_1^{RP_1}\right) = [(2 \cdot 5) \cdot (2 \cdot 3) \cdot (2 \cdot 2)] = 240 \text{ units}^3 .$$

The coverage of RP_1 with respect to A_1 is calculated as:

$$COV = \frac{AV\left(A_1^{RP_1}\right)}{AV\left(A_1\right)} = \frac{240units^3}{810units^3} = 0.2962$$

As a result, only 29.62% of the total volume of A_1 is covered by RP_1. This value helps improving the accuracy in the evaluation and selection of response plans. The remaining 70.38% of the attack is considered as a residual risk. Details on the computation of the system, attack, and countermeasure volumes can be found in [26, 27].

4.2 Response operational impact assessment

The response operational impact assessment (ROIA) is supported by three learners, which acquire, learn, define, and evaluate all information needed for a response operational impact assessment based on the described probabilistic model.

4.2.1 Network dependency analysis

As mentioned earlier, a resource dependency model must be kept up-to-date over time and must automatically adapt to concept drifts, i.e., to changes in the network. Therefore, a resource dependency model is continuously learned from network traffic analyses inside an infrastructure. In more detail, a resource dependency model is learned from maximum likelihood estimates based on statistical analyses in network traffic meta data in a similar fashion to learning classical probabilistic graphical models.

In our use case, a resource dependency model consists of a medium-sized ICT environment, in which some ICT devices also represent gateways to an industry SCADA system. Further, it can be assumed that every device drives one purpose. This allows for a simple heuristic on exchanged information amounts, initially proposed by us in [10] and [11] to obtain a plausible resource dependency network, explained at the following simple example: A workstation X consuming different query results from multiple databases distribute gained and processed information from such queries to other devices. The percentage of received traffic $T_{Y_i,X}$ from every database Y_i towards the total received traffic gives a good guideline for the conditional dependency between them as $p(^+x|^+y_i) = \frac{T_{Y_i,X}}{\sum_i T_{Y_i,X}}$. This measure is seen as a maximum likelihood estimate for the conditional probability and directly reflects classical learning approaches for learning probabilistic graphical models from data, e.g., learning parameters of Bayesian networks [28, pp. 806-808] from (in)complete data using "counts."

Periodically capturing and analyzing conversation statistics from traffic metadata is a feasible process and, e.g., directly provided by Wireshark [29]. Therefore, this heuristic is trivially to implement and has low and constant memory requirements with low computational load over time.

The heuristic that the *relative amount* of transferred information directly represents the local conditional probability associated with an edge is a venturesome assumption on first sight. Nevertheless, it delivers great and validated results as demonstrated by [11] and in Section 6. From a theoretical perspective, this heuristic must hold under the assumption that all communication directed towards a node, e.g., device, is equally encoded and has similar entropy. In such a case, every received byte, kByte, MByte, etc. must directly correspond to the importance is the transferred information, as long as no irrelevant data is transferred. Similar encoding and similar entropy communication can often be assumed: For example, in industrial control systems, a SCADA server will communicate with remote terminal units (RTUs) over some simplistic protocol and send control commands. These control commands will be similarly encoded and share similar entropy. In effect, each RTU will be highly dependent on each SCADA server as almost all received traffic will originate from those. Additionally, each RTU will acknowledge received commands, i.e., each SCADA server will be slightly dependent on each RTU. Moreover, SCADA servers frequently synchronize with each other, leading to larger information transfers, and will frequently communicate with human machine interfaces (HMIs). Therefore, SCADA servers will be highly dependent on each other, moderately dependent on HMIs, and slightly on RTUs, as one expects. Similar assumptions can be made for ICT architectures, where distributed databases synchronize, data is received from and transferred to webservers, and analyzed in large computational clusters.

An implementation of this learning approach delivered great results for multiple use cases as discussed in [11] and outlined in Section 6.

4.2.2 Local impact definition

The introduced probabilistic mission impact model is based on general external shock events. In order to obtain a *response* OIA, a response plan must be transformed to external shock events. Every mitigation action inside a response plan represents a potential cause for local harm, i.e., represents an external shock event. Therefore, every response plan RP is a collection of external shock events $\vec{SE}$ and a vector of observations on those $\vec{se}_o$. Effectively, one may calculate an impact assessment on a mission node *MN*, i.e., $P(^{+}mn|\vec{se}_o)$ for a response plan *RP*.

For example, a shutdown of a node X might cause other transitively dependent nodes to not work as intended, i.e., become impacted. Assessing the global effects of a local action is intuitively not possible and is the goal of an ROIA. However, local assessments are validatable and can even be grounded on common sense: Given one shuts down a node X, the probability that it will be impacted, i.e., not work as intended, is 100%: $p(^{+}x|^{+}shutdown_x) = 1$. We extend [10] and [11]'s proposed external shock event transformation from response plans, which have been validated and verified against the expectations of Panoptesec's product owner.

Definition 4.1 (*Response plan side effects*) *We employ mission impact assessment to achieve a qualitative assessment of potential negative side effects of a proposed response plan to an ongoing or potential attack. We see a response plan as a collection of individual actions affecting a network. For example, a shutdown of a server might easily reduce the surface of a potential attack. Still, if a critical resource is highly dependent on that server, it might impact a mission even heavier than a potential attack. We consider three mitigation-action types and transform them to external shock events, possibly leading to local impacts. The aim is to achieve a qualitative mission impact assessment of a response plan.*

We define external shock events by using three abstract temporal timeslices: t_0 *representing a short-term impact,* t_1 *representing a mid-term impact, and* t_2 *representing a long-term impact.*

If a node is shutdown (^{+}se*: the external shock event is present) it is easy to assess a probability of local impact to be 1. This means,* $p(^{+}x|^{+}se) = \langle t_0 : 1, t_1 : 1, t_2 : 1\rangle$*. Likewise, restarting a resource has the same effect as a shutdown in* t_0*, and might likely lead to hardware failure during reboot in a mid-term* t_1*, but will locally not cause conflicts in a long-term:* $p(^{+}x|^{+}se) = \langle t_0 : 1, t_1 : 0.6, t_2 : 0\rangle$*.*

Employing a patch on a node X might produce collateral damage as well. During installation of the patch, there exists a (low) probability of immediate conflict, e.g., a flat assumption of 10% or a measure published by a software vendor. In a mean time, a patch might enforce a reboot of a resource. This leads to a temporal shutdown and might lead to hardware failure. Finally, after a successful reboot, a replacement of hardware, and/or a restore of a previous backup, the network device will fully resume its operational capability. Therefore, $p(^{+}x|^{+}se) = \langle t_0 : 0.1, t_1 : 1.0, t_2 : 0.0\rangle$*. We argue that every installation, update or change of software can be modeled from an impact perspective as a patching operation.*

Like software is exchanged by a patch, hardware can be reconfigured as well. A reconfiguration is likely to enforce a reboot, if an exchanged component is not hot swappable. Therefore, we assume the same local impact as induced by a reboot.

Our third considered mitigation action is the restriction of a connection from node X to node Y, i.e., a new firewall rule. From a technical perspective this operation forbids a transfer of data that might have been crucial for the operational capability of a node Y. Therefore, a firewall rule leads to an operational impact on Y. We must assess this impact locally. This is a special case requiring Pearl's [30] *do-calculus. As a connection between two devices resembles a dependency, we must further actually remove this dependency. Otherwise, we would infer further impacts over a dependency that was prohibited and already assessed locally. To do so, we simply "bend" the forbidden dependency to an observed external shock event* ^{+}se *s.t. the local conditional failure probability* $p(^{+}y|^{+}x)$ *becomes a local impact probability* $p(^{+}y|^{+}se)$*. Another approach, decidable by a security operator, would be to accumulate dropped connections and add an unified local impact for them.*

All of these external shock events are deliberately placed inside our domain, we model their prior probability to exist as a tautology, i.e., $p(^{+}se) = 1$, and, obviously, fully observe the presence of mitigation actions, i.e., all modeled shock events $\vec{SE}$ represent the observed events $\vec{se}_o$. Further examples for shock events are given by Motzek et al. in [10, 11, 31].

4.2.3 Monte-Carlo evaluation

As mentioned before, an exact calculation of $\bigcup_i \mathbf{P}(w_i^{MN})$ is possible by the inclusion and exclusion principle and the Sylvester-Poincaré equality, but is exponentially hard due to the subtraction of all overlapping sets and is therefore not practical. We therefore approximate a solution to the MIA problem by the use of an approximate inference technique proposed in [11] for which we discuss technical details and implementations in this section.

For every mission node *MN*, there exists a Boolean formula $\vec{w}^{MN}$ as a disjunction of conjunction over Boolean random variables $\vec{B}$. However, Boolean random variables in $\vec{B}$ take their respective truth value according to a probability distribution. To approximate $\bigcup_i \mathbf{P}(w_i^{MN})$, i.e., to find an approximate solution to the MIA problem, a complete instantiation of all Boolean variables $\vec{B}$ is drawn by sampling every Boolean variable according to its distribution, and $\vec{w}^{MN}$ is checked for satisfaction. Repeating this process n times, where n^{+} times a satisfaction was found, approximates $P(^{+}mn|\vec{se}_o)$ by n^{+}/n. Our results show that an upper three-sigma bound of expected error $\bar{E}$ is obtained by $\bar{E} = 0.775 \cdot \sqrt{n}^{-1}$. A detailed description and evaluation is given in [11] and left out for brevity in this paper. In summary, evaluations by [11] show that the utilized approximation method scales linearly with a resource dependency model's complexity in terms of nodes and edges, scales linearly with the number of simulation rounds, and scales linearly with the number of found Boolean disjunctions. As a consequence, evaluations in all three dimensions are obtained in the range of seconds even for large, deeply meshed graphs with hundreds of nodes.

In order to approximate $P(^{+}mn|\vec{se}_o)$, a two-step approximation technique is employed, which calculates the conditional probability through a probabilistic path search. We first acquire all paths leading to external shock events, for every mission node for which we would like to perform mission impact assessment. Often, we would like to perform this for every node in the mission dependency model. Finding paths for a node in the mission dependency graph is trivial, given found paths from business resources to external shock events. We therefore, as step one, acquire (all) paths leading to evidence for all business resources, which is a classic graph search problem.

Under the assumption that the number of business resources and external shock events is comparably small to all nodes in the network graph, a depth-limited search is a reasonable approach for finding paths leading to external shock events. If this assumption does not hold, an alternative is discussed after the following definitions and remarks.

Definition 4.2 (*Probabilistic paths*) *For every business resources, $BR_i \in \vec{BR}$ let $\vec{w}^{BR_i}$ denote the set of all paths leading to an external shock event and let $w_j^{BR_i}$ denote the jth path. Let $\vec{w}$ denote the super-set of all found paths. Every path $w_j^{BR_i}$ is a set of individual conditional probability fragments $p(x|y)$, representing an edge, i.e., a dependency, from y to x. The product of all probability fragments $p(x|y) \in w_j^{BR_i}$ is the exist-probability of a path $P(w_j^{BR_i})$. Every path $w_k^{BR_i}$ for which holds $\exists j : w_j^{BR_i} \subseteq w_k^{BR_i}$ is irrelevant for calculation and $\vec{w}$ is a finite set. Informally, this means during path search along one path, an already visited node must not be visited again and we cannot get stuck in infinite loops.*

After acquiring all paths $\vec{w}$ leading to all business resources, subsequent paths leading to business functions, processes, and the company are trivially acquired by following the paths leading to all children.

Step two is a Monte-Carlo simulation to approximate $P(\bigvee \vec{w}^{BR_i})$ for every business resource $BR_i \in \vec{BR}$. We draw a sample from $\vec{w}$ and from all dependencies in the mission dependency model. We check for every BR_i the satisfaction of $\bigvee \vec{w}^{BR_i}$ and mark the satisfaction result on BR_i. Subsequently, we check for satisfaction of any children, i.e., dependencies, of every node in the mission dependency model. Every satisfaction for a mission node MN found in the mission dependency model is marked as a hit in hit_{MN}. After n_S iterations, the desired conditional probability of MN being impacted (mn), i.e., the *mission impact*, given all external shock events $se_o \in \vec{se}_o$ is approximated by $P(^{+}mn|\vec{se}_o) = \frac{hit_{MN}}{n_S}$.

Remark 4.4 (Path check) *Checking all paths during one Monte-Carlo round is highly optimizable. $\vec{w}^{BR_i}$ can be sorted descending by $P(^{+}w_j^{BR_i})$, s.t. most likely existing paths are checked first and subsequent checks can be skipped once a path is found. Further, a path $w_j^{BR_i}$ can be sorted ascending by its individual local conditional probability fragments s.t. most unlikely random variables are checked first and further checks inside one path can be skipped. Further, a path w with $P(w) < \frac{1}{n_S}$ will statistically never be drawn, i.e., all such paths can be skipped during simulation and check. Notwithstanding, the complete process is highly parallelizable.*

Following this procedure, at first, the complete graph is searched for all proofs, followed by a simulation of these proofs. For certain graphs, various found proofs will share common subsets, which are simulated redundantly. This redundant simulation may represent a performance bottleneck. In such a particular situation, it may be beneficial to not base the simulation on a strict disjunction of conjunctions, but to preserve the underlying graph structure for simulation. However, such an optimization heavily depends on the graph structure, and evaluations in [10] and [11] have shown that evaluations are obtained in the range of seconds without further optimization on a graph's structure. Further implementation details and remarks on the procedure are given in [11].

5 Selection of Pareto-efficient response plans

RFIA and ROIA evaluate different forms of impacts implied by proposed response plans. As mentioned before, their nature is complementary to each other and it is neither trivial nor "intuitive" to rank all evaluated response and select "the best." In the following, we present an approach to unify all assessments without becoming biased towards one dimension.

Effectively, every response plan is associated with a four-dimensional impact assessment, i.e., four real valued numbers. There exists no ordering among four-dimensional values, and all impact assessments must be seen as equally important, e.g., one is not biased towards preferring RORI above some OI_i. This is why one cannot trivially reduce all four dimensions to a single one, and one cannot find a simple ordering. Notwithstanding, such naive dimension-reduction approaches exist and are frequently used, but such approaches suffer from significant problems as later discussed in Section 8.1.

As all impact assessments are seen as equally important, a response plan may dominate all other response plans in *all* impact dimension. In such a particular situation, "the best" response plan is well-defined. However, it is very likely that no such response plan exist, as, e.g., the do-nothing response plan will, by definition, lead to the most optimal response plan in terms of operational impact, but must be non-optimal in terms of financial impact. Effectively, no clear optimum will dominate, and a compromise among all dimensions must be found. Finding such a best compromise is known as finding a Pareto-efficient set.

The proposed FIA results in a linear, relative metric, i.e., assessments depend on a use case and context. Therefore, response plans are only interpretable, evaluable, and comparable for one common use case and context scenario. This means that there exists a well-defined ordering for all obtained FIAs from one scenario, but FIAs from one scenario cannot be compared to another scenario. In effect, relative reference points are required for obtaining an absolute scale for each scenario.

The proposed OIA is based on a probabilistic model resulting in a stable, absolute metric, e.g., an assessment of, say, 5% is understandable and interpretable independent of any context, use case, or scenario. For example, an OIA of 5% for a potential impact on a company, given a set of observed external shock event, is equivalent to a 5% of winning a lottery, given one plays the lottery, or a 5% probability of tossing a 1 on a twenty-sided cube. Each OIA consists of an n-dimensional vector representing a temporal diversity, e.g., short-, mid-, and long-term assessments. There exists a well-defined ordering for every temporal dimension by itself, but not for all in combination.

Based on these characteristics of OIA and FIA, we propose a selection of response plans based on a best compromise. Every dimension is considered equally and optimized individually until one finds a solution. We therefore define a selection of response plans based on a Pareto-efficient solution among all impact assessment dimensions as follows.

Definition 5.1 (*Pareto-efficient response plans*) *Let $\vec{RP}^d$ be a vector of proposed response plans, associated with a linearly scaled impact assessment of dimension d. Let $\dot{RP}^d \subseteq \vec{RP}^d$ denote the set of optimal-proposed response plans in terms of dimension d. Let $\hat{RP}^d$ denote the assessment of the theoretical optimal response plan, and let $\check{RP}^d$ denote the assessment of the theoretical worst response plan in terms of dimension d. Then, let $\dot{RP}^d_\varepsilon \subseteq \vec{RP}^d$ represent the set of Pareto-efficient response plans in terms of dimension d and easing factor $\varepsilon \in [0,1]$ representing the allowed deviation ε of the theoretical response plan range $|\hat{RP}^d - \check{RP}^d|$ from the evaluated optimal response plan $\dot{RP}^d$. Thus, $\dot{RP}^d_0 = \dot{RP}^d$ and $\dot{RP}^d_1 = \vec{RP}^d$.*

Finding the best compromise among an n-dimensional impact assessment is therefore defined as finding the smallest Pareto-efficient set.

Definition 5.2 (*Smallest Pareto-efficient set*) *Let $\vec{d}$ be the vector of all impact dimensions. Then, the smallest Pareto-efficient set of response plans $\mathring{RP}$ is the set*

$$\mathring{RP} = \min_{\varepsilon} \left(\left\{ \bigcap_{d \in \vec{d}} \dot{RP}^d_\varepsilon \right\} \neq \emptyset \right) \tag{9}$$

As the ROIA represents an absolute metric, $\check{RP}^{ROI} = 1$ and $\hat{RP}^{ROI} = 0$. For the relative RFIA metric, $\check{RP}^{RFI}$ and $\hat{RP}^{RFI}$ depend on $\vec{RP}^{RFI}$. If not all possibly allowed response plans are evaluated by the RFIA for performance criteria, $\check{RP}^{RFI}$ and $\hat{RP}^{RFI}$ are not uniquely identifiable and must be estimated by $\check{RP}^{RFI} = -1$ and $\hat{RP}^{RFI} = \dot{RP}^{RFI}$. This means, $\dot{RP}^{RFI}_\varepsilon$ might be too large. A selection of a response plan according to Definition 5.2 can efficiently be performed by using binary search.

By finding the smallest Pareto-efficient set, one selects a *set* of response plans which shows to be, to some degree, superior to all other response plans. By Definition 5.1, no further ordering exists inside the selected smallest Pareto-efficient set, and all selected response plans are considered equally superior. Nevertheless, it may be possible that some selected response plans are fully dominated by others in the case of draws in some dimensions. These cases are not considered further, as they are obviously seen, and it depends on a use case whether these may actually be superior. Notwithstanding, one may still order a smallest Pareto-efficient set by some criterion, e.g., by a short-time operational impact, for ease of visualization and operator assistance.

The choice of using a Pareto-efficient set as the optimization procedure bear vital benefits and is later discussed intensively in Section 8.1.

6 Use case demonstration

All presented models, approaches, and optimizations are implemented in the cyber-defense system PANOPTESEC, whose use case partner gives us the opportunity to evaluate and study the application of the proposed multi-dimensional evaluation and optimization of response plans in an infrastructure environment of an energy distribution organization (EDO). The environment consists of a distributed network of remote terminal units (RTU) in energy stations of medium voltage (MV = 20,000 V) and high voltage (HV = 150,000 V). RTUs acquire data from electrical equipments (e.g., PLC, sensors), and send data to a supervisor terminal unit (STU) of the headquarter. The RTU network utilizes Supervisory Control and Data

Acquisition (SCADA) protocols and is composed of over 13,000 energy stations, 6000 of which are controlled by the STU.

In the absence of security compromise, operators review the security status of the monitored system (SCADA and ICT environment). Security status indicators may note the presence of one or more system vulnerabilities due to known software security flaws as posted by publicly available vulnerability advisory services. Attack paths from hypothetical attack sources to known mission critical systems are analyzed, and the impact on critical business functions (e.g., energy distribution) is assessed resulting in a quantified risk assessment.

From Table 1, we organize the information of the EDO according to their nature (dimension). For example, we obtain servers, firewalls, IDs, etc. as resources, IP addresses and port numbers as channels, and operators as user accounts. Depending on the type of element and their importance to the mission of the organization, we assign a weighting factor. A basic operator is assigned a WF=1, whereas an advanced operator has a WF=4, and a supervisor has a WF=5. For those cases where the category regroups elements of different types (e.g., SCADA Servers, Web servers, NTP Server, etc are regroup as Servers), we assign a weighting factor for each type of element, going from one to five.

Table 1 Information of the EDO system

Dimension	Elements	Description	Q	WF
Resource	R1:R12	HV/MV Server	12	1-5
	R13:R16	HV/MV Front End	4	4
	R17:R22	HV/MV Gateway	4	4
	R23:R56	Routers	34	3-4
	R57:R63	Human-Machine Interface	6	2-3
	R64:R363	Remote Terminal Unit	300	5
	R364:R365	Firewall	2	2
	R366	PC	1	2
	R367:R368	IDS	2	2
Channel	Ch1:Ch2	Public IP address	2	3
	Ch3:Ch302	Private IP address	300	2
	Ch303:Ch698	UDP Port	396	1-5
	Ch699:Ch1712	TCP Port	1014	3-5
User Account	U1:U30	Basic Operator	30	1
	U31:U38	Advanced Operator	8	4
	U39:U52	High Voltage Operator	14	3
	U53:U70	Medium Voltage Operator	18	2
	U71	Supervisor	1	5

We have modeled 368 resources, 1712 channels, and 71 User accounts. The quantity of each type of element is shown in the column Q, and its weighting factor is shown in the column WF

The annual infrastructure value for the EDO is equivalent to 11,379,800 €, which represents the cost of operation, license, maintenance, and services incurred in a yearly basis for the regular operations of the organization. It considers the annual cost of all the policy enforcement points (PEPs) of the organization.

6.1 Threat scenarios

Use case providers have identified five threat scenarios that could lead to severe consequences on the target system. Table 2 summarizes the information associated to all possible threats, PEPs, and attack vectors of each identified threats.

Threat $AS02$ has been detailedly analyzed. $AS02$ corresponds to a compromise of a specific target through vulnerability exploitation, which will cause data corruption or leakage of a database in the ICT domain. For such a threat, there exists a specific attack vector, as shown in Table 2.

6.2 Financial impact assessment

Threat $AS02$ has a serious severity (1,000,000 €), and a high likelihood (f=12), which results into an $ALE = 12,000,000$ €$/year$. This threat has been associated to a set of mitigation actions. Combinations of associated mitigation actions form response plans shall improve the security status of the monitored system (e.g., patch deployment, shutdown, restart, or other system reconfiguration). They are selected and executed by operators resulting in automated deployment of mitigation actions where possible (e.g., firewall reconfigurations) or otherwise issuing instructions to senior operators for follow-up

Table 2 SCADA threat scenarios

Threat	Description	ALE (€)	Attack Vector
AS01HV	DoS to High Voltage nodes	20,000,000	EP=VGROUTER; T1=WEBSCADA; T2=FTPSRV; BD=FEXSCADA
AS01MV	DoS to Medium Voltage nodes	2,000,000	EP=RTUSCADA; T1=GWSCADA; T2=FEXSCADA; BD=SRVSCADA
AS02	Data corruption or leakage	12,000,000	EP=VGROUTER; T1=WEBSCADA; T2=USERPC; BD=FTPSRV
AS03	DoS against electrical devices	100,000	EP=RTUSCADA; T1=GWSCADA; T2=FEXSCADA; BD=SRVSCADA
AS04	DoS against business Services	2,000,000	EP=VGROUTER; T1=FTPSRV; T2=USERPC; BD=WEBSCADA

There are five threats that could affect the EDO system; from which Threat AS02, has been selected to be evaluated deployment of actions (e.g., patch deployment). Table 3 details information of the authorized mitigation actions for threat $AS02$.

From Table 3, the values of ARC and EF have been estimated based on expert knowledge and historical data. The RM value is calculated as the product of the EF and coverage (COV), where the coverage is obtained using geometrical operations from the attack volume model as discussed in Section 4.1.2. The RORI index is calculated using Eq. 3.

From the list of proposed mitigation actions, MA_1 (reconfiguration of V-R08) provides the highest RORI index. By taking this action, the risk is expected to be reduced to 60% (RM), resulting in a RORI index of 63.27. Response plans for this threat are formed by combinations of all possible mitigation actions, considering those actions that are mutually exclusive (e.g.,MA_{10} cannot be simultaneously implemented with MA_2, MA_4, and MA_7). All potential combinations, i.e., 797 response plans, were evaluated and the best response plan results in a RORI index of $\hat{RP}^{RFI} = 97.1435$ with a combination of mitigation actions as $\langle MA_1, MA_2, MA_3, MA_4, MA_5, MA_6, MA_7, MA_8, MA_9 \rangle$. The worst is represented by $\{\langle MA_7, MA_9 \rangle \, \langle MA_8, MA_9 \rangle\}$ with $\check{RP}^{RFI} = 0.21$. Please note that a response plan's RORI is not a linear combination of individual RORI scores associated with individual mitigation actions, but that an RFIA evaluation is performed for every response plan of all 797 response plans yielding an individual RORI score for every one of them.

An example of the graphical representation of the evaluated threat AS_02 on the EDO, associated channels, resources, and user accounts, and a visualization of a corresponding response plan is given in Fig. 3.

Table 3 Examples of RFIA evaluations for proposed response plans RP_i consisting of a single mitigation action MA_i for threat scenario $AS02$

MA	Description	EF	COV	RM	ARC	Restriction	RORI
MA_1	Reconfig. V-R08	1.00	0.60	0.60	50	None	63.27
MA_2	Reconfig. Web-SRV	0.80	0.15	0.12	1000	MA_{10}	12.64
MA_3	Reconfig. File-SRV	0.80	0.15	0.12	500	MA_{11}	12.65
MA_4	Patch Web-SRV	1.00	0.15	0.15	2000	MA_{10}	15.8
MA_5	Patch File-SRV	1.00	0.15	0.15	500	MA_{11}	15.81
MA_6	Patch User-PC	1.00	0.10	0.10	500	MA_{12}	10.54
MA_7	Restart Web-SRV	0.01	0.15	0.00	50	MA_{10}	0.16
MA_8	Restart File-SRV	0.01	0.15	0.00	50	MA_{11}	0.16
MA_9	Restart User-PC	0.01	0.10	0.00	50	MA_{12}	0.11
MA_{10}	Shutdown Web-SRV	0.10	0.15	0.01	50	$MA_{2,4,7}$	1.58
MA_{11}	Shutdown File-SRV	0.10	0.15	0.01	50	$MA_{3,5,8}$	1.58
MA_{12}	Shutdown User-PC	0.10	0.10	0.01	50	$MA_{6,9}$	1.05

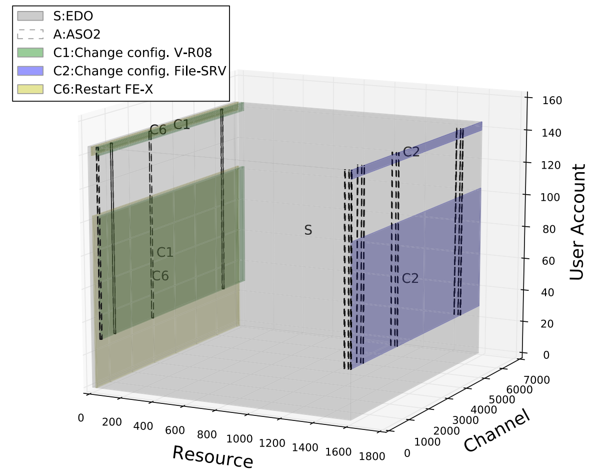

Fig. 3 Graphical representation of threat *AS*02 (*dashed*) and a corresponding response plan composed of three mitigation actions (C_1 *yellow*, C_6 *green*, and C_2 *blue*, from left to right). The *gray box* represents the full operational dimension of the EDO

6.3 Operational impact assessment

To perform a response operational impact assessment, a resource dependency model is needed. As described in Section 3, a manual assessment is said to be infeasible, and a solution based on a heuristic of exchanged traffic information was proposed in Section 4.2.1. Network traffic has been analyzed in a completely replicated backup environment of the use case partner involving all SCADA and ICT communications over multiple months. Metadata of all traffic, i.e., header information such as IPs and MAC-addresses, are recorded constantly over 50 min intervals and are analyzed postponed by the introduced NDA module. All obtained information is synchronized and cross-checked against a central network inventory s.t. one is able to distinguish communicating- and communication-establishing devices. The differentiation is based on an analysis of network inventory matches against IPs and MAC-addresses. The complete process for the use case partner is discussed in great detail in [11], where additional implementation details and proposals are made. The obtained network dependency model is shown in Fig. 4. Figure 4 shows the automatically learned and analyzed dependencies between ICT devices (shown in white) and directly business critical resources (shown in green). Further, all business processes (shown in orange) and their correspondingly required business functions (shown in blue) are highlighted in Fig. 4. Most interestingly, the automatic analysis revealed two "clouds" (see Fig. 4, lower right) of highly dependent nodes, which are clusters of remote terminal units communicating with the central control server. The obtained network dependency model has been validated by an external IT specialist consultant to the company to be reasonable, to contain the most

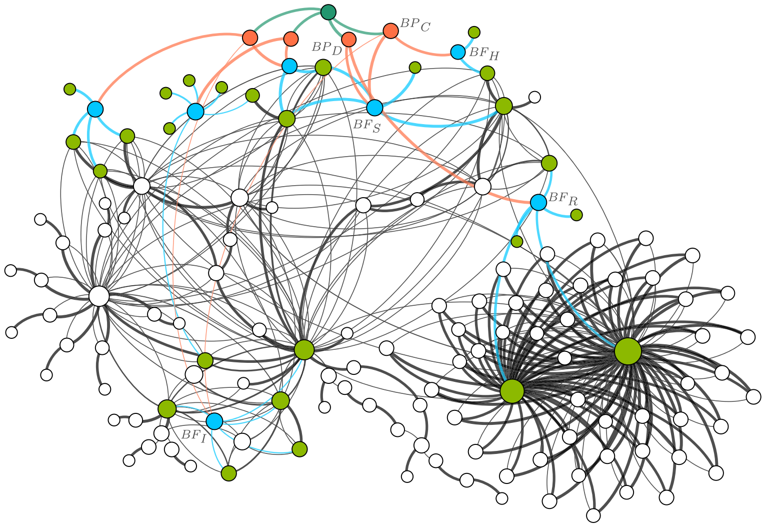

Fig. 4 Mission dependency model and resource dependency model obtained and learned from data of the use case partner. *BF*BF_R represents a business function (*blue*) subject to handling remote terminal units (RTUs), which are visible as the lower right clouds of nodes, where the central nodes are business critical devices (*green*). Business processes shown in *orange*, business company in *dark green*. Thicker and darker edges represent higher dependency degrees. Visualized using Gephi [50]

important devices, and to bear reasonable dependency degrees.

Figure 4 additionally shows an obtained mission dependency model in combination with the automatically learned resource dependency model. The mission dependency model was obtained by deep collaboration with various business and IT experts of the company in multiple iterations. As all dependencies are understandable by themselves, i.e., one does not require a complete global picture to grasp the semantic of one conditional probability, knowledge from multiple experts was collectible individually. This provided a great benefit, as every expert could be interviewed individually, and no large meeting had to be organized where all participants needed to come to one conclusion. Often, if such models are only generateable holistically, a bad compromise is chosen due to disagreements between multiple experts. In our approach, all view points of participants are includable, and the model is designed to accept, respect, and overcome these disagreements. We discuss these benefits further in [11] and show how individually collected information are mergeable into one model while containing all information content in [12].

Based on the resource dependency model and the mission dependency model from the use case partner, ROI assessments for all proposed response plans are evaluated, for which an excerpt is given in Table 4 corresponding to Table 3. As explained throughout Section 4.2, an OIA is an evaluated marginal probability of impact of a node X given observed shock events $\vec{se}_o$, i.e., $P(^+x|\vec{se}_o)$. Given time-varying probabilities, such an evaluation returns a multi-dimensional impact assessment, i.e., a probability for every time slice, e.g., $P(^+x|\vec{se}_o) = \langle t_0 : 0.1, t_1 : 0.9, t_2 : 0\rangle$. This probability value is obtained for every node of a mission dependency model. In the following, we solely discuss and consider the most-highest node in a mission dependency model, i.e., the mission or business company *CM*. Moreover, we abbreviate the probability nomenclature by simply writing $RP = \langle t_0 : 0.1, \ldots, t_2 : 0.3\rangle$ instead of $P(^+mn|\vec{se}_o) = \langle t_0 : 0.1, \ldots, t_2 : 0.3\rangle$, where $\vec{se}_o$ is a transformation of response plan *RP* to (observed) external

Table 4 Examples of ROIA evaluations for proposed response plans RP_i consisting of a single mitigation action MA_i for threat scenario *AS*02 corresponding to Table 3. (*RORI given for reference*)

MA	Description	RORI	OI_0	OI_1	OI_2
MA_1	Reconfig. V-R08	63.27	4.2%	2.4%	0
MA_2	Reconfig. Web-SRV	12.64	6.6%	3.6%	0
MA_3	Reconfig. File-SRV	12.65	36.6%	22.2%	0
MA_4	Patch Web-SRV	15.8	0.6%	6.6%	0
MA_5	Patch File-SRV	15.81	3.6%	37.2%	0
MA_6	Patch User-PC	10.54	0.6%	6.6%	0
MA_7	Restart Web-SRV	0.16	6.6%	4.2%	0
MA_8	Restart File-SRV	0.16	36.6%	22.2%	0
MA_9	Restart User-PC	0.11	6.6%	4.2%	0
MA_{10}	Shutdown Web-SRV	1.58	7.2%	7.2%	7.2%
MA_{11}	Shutdown File-SRV	1.58	40.8%	40.8%	40.8%
MA_{12}	Shutdown User-PC	1.05	7.2%	7.2%	7.2%

shock events. Further, we refer to individual timeslices simply by OI_i for the i-th timeslice.

The comparison between a RORI index and operational impact assessments in Table 4 shows how both lowest and highest probabilities of operational impact lead to extremely low RORI indices. In order to emphasize the latter, Fig. 5 shows a scatter plot of all evaluated response plans in AS03. From Fig. 5, it is evident that no simple correlation between OIA and FIA exists, i.e., both impact assessment evaluate a different kind of impact and are both required for a sound selection of optimal response plans. In fact, Pearson's product-moment correlation coefficient between RORI and OI_0 and OI_1 for all evaluated response plan is $\approx$ 0.14 and between RORI and OI_2 even $\approx$ 0.01, showing that OIA and FIA are almost uncorrelated. Furthermore, Fig. 5 shows the top four Pareto-efficient response plans (highlighted in red), which yield in high, i.e., good, RFIA, and low, i.e., good, probability of "collateral damage." However, please note that Pareto-efficiency cannot be seen or manually analyzed from Fig. 5, as the figure is a reduction of a four-dimensional hyperspace onto three planar plots. In order to find a geometrical visualization for an evaluated response plan and Pareto-efficiency, one needs to consider a four-dimensional hypercube, where OI_0, OI_1, OI_2, and RORI represent all four axis, such that every evaluated response plan is a point on its respective geometrical coordinate in the hypercube. The Pareto-efficient response plans then span a surface of Pareto-efficiency in this hypercube.

6.4 Pareto-efficient response plan selection

Judging from Table 4, a good compromise seems to be deploying mitigation action MA_1 alone, resulting in both a low probability of operational impact and being financially attractive in terms of RORI. Still, deploying MA_1 alone is not the best option. The most financially attractive response plan $RP_R = \langle MA_1, MA_2, MA_3, MA_4, MA_5, MA_6, MA_7, MA_8, MA_9 \rangle$ with a RORI index of 97.1435, however, is assessed to bear almost the highest probability of operational impact with $\langle t_0 : 0.408 \; t_1 : 0.402 \; t_2 : 0.0 \rangle$. Note that an OI assessment of a response plan is *not* a linear combination of individual mitigation actions, as a "double count" of probabilities is not allowed and would lead to spurious results. In terms of lowest short-term (t_0) OI probability, MA_4 and MA_6 alone show to be dominant, and in mid-term (t_1) MA_1 alone is dominant. In a long-term perspective (t_2), a large set of response plans is dominant with a 0 probability of impact. Thus RP_R, MA_4, and MA_6 represent a Pareto-optimal set. As proposed in Section 5, we search for the best *compromise*: From Definition 5.2, one obtains the best Pareto-efficient response plan $\mathring{RP} = \{\langle MA_1, MA_2, MA_4, MA_6, MA_7, MA_9 \rangle\}$ using $\varepsilon = 0.1475$, consisting of one response plan with an operational impact assessment of $\langle t_0 : 0.108 \; t_1 : 0.09 \; t_2 : 0.0 \rangle$ and a RORI index of 82.8514. This means, with a compromise of 14.75% of the theoretical optimum in every dimension from the evaluated optimum, a Pareto-efficient response plan is found.

Notwithstanding, multiple other optimization approaches exist, with which one may obtain similar results. In Section 8.1, we compare the proposed Pareto-efficient optimization with other, familiar approaches.

7 Related work

Current research focuses on considering the impact of attacks by evaluating their severity and consequences, leaving aside the impact of security actions in mitigating the effects of such attacks. Dini and Tiloca [1], for instance, propose a simulation framework that evaluates the impact of cyber-physical attacks, discuss the attack ranking process, and analyze different mitigation actions. However, the latter is not considered in the assessment of the attacks' impact nor they are ranked according to their effectiveness in stopping or mitigating the attacks.

Viduto et al. [32] present an approach to ponder risks posed by vulnerabilities against financial investments, while utilizing a multi-objective tabu search. Anyhow,

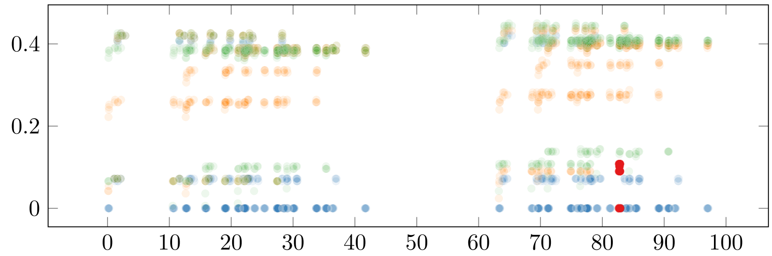

Fig. 5 Scatter plot of ROIA (ordinate, t_0 *blue*, t_1 *orange*, t_2 *green*) and RFIA (abscissa) evaluations for all response plans. Note that both impact evaluations are almost uncorrelated, i.e., both actually evaluate an impact from perpendicular perspectives, and all dimensions must be considered equally in the form of a Pareto-efficient set to obtain meaningful and valuable responses. The top four Pareto-efficient response plans are marked in *red*

they only consider direct impacts on devices, which leads to an assumption that, in the absence of vulnerability and a corresponding exploit, no harm can be caused at all. Moreover, the usage of relative metrics does not provide an intuitive interpretation of parameters. Therefore, their work represents a holistic approach without a considering transitive business impact relations and negative side effects of the proposed response plans.

Considering transitive and indirect effects is partially they Foo et al. [33] identify by the use of "spread" channels in a network, they notably identify that considering only local reactions to raised alarms is insufficient and they try to maintain subfunctions of business services. However, they utilize a novelly designed propagation algorithm for "spreading" impacts, which is not mathematically grounded and solely provides response analyses in an intransparent approach.

Considering the negative side effects of responses is often only performed by a cost-perspective on the implementation, i.e., how much money must be spent to implement some response plans. [34] presents an interesting and well-formalized approach for situations where too few budget exists to fully implement all mitigation actions to known attack surfaces. [35], [36], and Fiedler [37] consider a defender-attacker interaction as strategic games and present well-formalized definitions and well-defined problems for winning these games based on a cost optimization. Other cost-focused approaches are proposed in [32, 36, 38, 39]. However, only considering the cost of *implementation* has a significant drawback: as mentioned in the beginning, the cost of shutting down a highly critical node will certainly eliminate an attack surface and involves almost no costs at all for implementation. Notably, Fiedler et al. consider indirect costs where a degree of performance disruption is covered in performing required business tasks. However, Fiedler et al.'s consideration does not consider the transitive impacts implied by one mitigation action leading to potential chains of failures in a network. Moreover, it is common for all abovementioned works that a detailed analysis of a defender-attacker interaction is required. However, acquiring a detailed analysis of defender-attacker interactions boils down to a detailed prediction of the future on how an attacker will infiltrate a network. Examples such as Stuxnet have shown that a prediction at the required detail level is impractical and nearly impossible. Due to the automatic analyses of transitive impacts in our ROIA and the volume-centric operation in the RFIA, our approach does not require detailed analyses of attacker behaviors. Notwithstanding, such analyses help one to estimate attacker volumes, as discussed in [26, 27], but our approach does not centrally built on a detailed analysis and a rough analysis is solely used for parameter estimation.

A further cost-focused and business economical focused approach is presented by [40] who resort to a meta-discussion of involved human factors in deriving various business economical figures related to security aspects. Unfortunately, what is yet missing in approach is a clear mathematical formalization of underlying mathematical problems and a formalized description to apply their derived cost metrics to an IT security scenario.

Kundur et al. [2] propose a paradigm for cyber attack impact analysis that employs a graph-theoretic structure and a dynamical systems framework to model the complex interactions among the various system components. The approach involves quantifying the effects of given classes of cyber attack, providing information on the degree of disruption that such class of attacks enable, and identifying sophisticated dependencies between the cyber and physical systems, but leaves aside the impact of mitigation actions in the attack's impact calculation.

Squoras et al. [4] present a qualitative assessment of the cyber attack impact on critical Smart Grid infrastructures. Authors evaluate the impact of DoS/DDoS attacks on data availability without considering mitigation actions in the assessment of the overall impact calculation.

In terms of operational impact assessment, probabilistic models have been investigated as an adequate assessment of impacts or risks posed due to attacks or found vulnerabilities [41–43]. However, often imperfect knowledge is not considered [41] or dependency cycles pose a problem [43]. Other impact propagation approaches, able to handle such details, are not probabilistic based and degrade to a hand-crafted propagation algorithm with arbitrary scores [44, 45]. Similarly, Barreto et al. [46, 47] only consider direct impacts as approaches to mission modeling, leaving aside transitive impacts and require a manual description of all dependencies between individual devices inside one organization, which is an infeasible process.

Our approach proposes the evaluation and selection of mitigation actions based on the financial and operational assessment of security events (e.g., attacks and mitigation actions). The ultimate goal of our approach is to select the set of mitigation actions that provides the maximal positive financial gain and the minimal operational negative side effect. The response financial assessment considers the RORI as an index that ranks mitigation actions based on multiple factors. The operational assessment evaluates threats according to their nature expressed as local time-varying impacts and considers transitive impacts based on a well-defined probabilistic model, for an organization's missions and resources.

8 Discussion

Two different impact assessment approaches have been proposed (i.e., financial and operational), which seem to be conflicting at first sight: Every action taken in order to reduce a potential attack vector bears a potential negative side effect that needs to be reduced. Both, RFIA and ROIA, perform impact assessments of mitigation actions (individuals and/or combined) that integrate a given response plan. Since the financial and operational impact assessments are of different nature, they perform the evaluation of response plans from different perspectives. On the one hand, the RFIA aims at assessing mitigation actions based on their financial benefits to the system, proposing the response plan with the highest RORI index, thus the maximal positive financial gain to the organization. On the other hand, the ROIA aims at assessing mitigation actions based on their potential collateral damage to the system, proposing the response plan with the lowest operational impact, thus the minimal operational negative side effect.

Combining both assessments is not a trivial task, since there exists no linear relationship between RFIA and the ROIA outputs. The RFIA provides relative measurements which are useful in obtaining an overall ratio scale ranking of the alternatives, whereas the ROIA provides absolute measurements that use precise values that scale with a given unit (e.g., hundreds, thousands, millions). Unlike absolute measurements, relative values derive ratio scales from paired comparisons represented by absolute numbers [13]. If the ratio produces repeatable and consistent results, the model can be used to compare security solutions based on relative values [14]. The absence of an absolute scale to compare RFIA and ROIA outputs makes it difficult to find an optimal response plan that satisfies all financial and operational criteria. We, therefore, propose a method that searches for an efficient solution related to a Pareto-optimum. The proposed Pareto-efficient optimization features significant advantages for this multi-dimensional optimization, which we discuss in more detail in the following subsection.

8.1 Comparison with other optimization and ranking approaches

A Pareto-efficient set of response plans might not be the best solution neither in financial nor in operational terms, but the Pareto-efficient response plans represent response plans that, on the one hand, bear the highest financial attractiveness on return on investment, and, on the other hand, bear the lowest probability of conflicting with a company's mission. This is beneficial for applications, where highly critical missions and resources must be protected, without sacrificing missions in favor of security.

From a mathematical perspective, the Pareto-efficient optimization is a multi-dimensional optimization with scaling-aware normalization and full information-content preservation. This means that, at first, every response plan RP of an evaluated set of response plan $\vec{RP}$ is normalized such that every dimension, i.e., *RORI*, OI_0, OI_1, OI_2, is based on a scale from 0 to 1, where 0 represents the best case of a dimension, and 1 represents the worst case of a dimension. It is extremely important to note that this normalization is scaling-aware, i.e., the normalization of each dimension d is based on $\hat{RP}^d$ and $\check{RP}^d$, and *not* on the current-best evaluated maxima and minima in each dimension. If one would naively normalize each dimension by the current-best evaluated maxima and minima, densely clustered evaluations in one dimension would be stretched artificially. As a counter example, say, one obtains evaluated response plans with extremely similar dimension-values $[0.5, 0.499, 0.501]$ on a theoretical scale from 0 to 1 (as for operational impact by default); naively normalizing these response plans by the maximum 0.501 and minimum 0.499 would completely distort these evaluations. In contrast, the Pareto-efficient optimization preserves the underlying semantic that all evaluations are extremely similar.

Based on these normalized response plans, the Pareto-efficient optimization searches for the best compromise in all dimensions by searching for a set of response plans where all evaluations are close to the current-best response plan in that dimension. In effect, some response plans may exist which are better in some dimension, say, OI_0, but which were not selected to be Pareto-efficient, because they deviated heavily in another dimension from the best solution.

Various other approaches exist for finding an "optimal" response plan from a set of multi-dimensional normalized values. An extremely common and naive approach is to simply reduce n-dimensional evaluations to a one-dimensional scalar by some combination function $f(\cdot)$. Note that this will inevitably reduce the information content during optimization, i.e., not all implications of the evaluated set of response plans can be considered. In contrast, the Pareto-efficient optimization preserves the complete information content, as all dimensions are preserved.

An example for such a naive combination function $f(\cdot)$ is a simple average or summation of all dimensions. Obtained one-dimensional scalar values can then be easily sorted, and an "optimum" is obtained trivially. Often, similar results are obtained to a Pareto-efficient optimization, but there exists significant and tenuous differences. Say, one is given three response plan evaluations, which have been normalized as described by our approach: $RP_1 = \langle t_0 : 0.5, t_1 : 0, t_2 : 0, rori : 0.901\rangle$, $RP_2 = \langle t_0 : 0.28, t_1 : 0.28, t_2 : 0, rori : 0.9\rangle$, and $RP_3 = \langle t_0 : 0.2, t_1 : 0.2, t_2 : 0.7, rori : 0.899\rangle$. Choosing a simple summation function $f(\cdot)$, one obtains $f(RP_1) = 1.401$, $f(RP_2) = 1.46$, and $f(RP_3) = 1.999$. Following, one would select RP_1 as it has the lowest score. However, RP_1 deviates heavily in OI_0 by 30% from the evaluated optimum in that dimension

($RP_3^{t_0} = 0.2$). RP_2 is Pareto-efficient, as it maximally deviates by only 8% from the evaluated optimum of each dimension.

Another frequently dimension-reduction function $f(\cdot)$ is a simple max operator which follows a similar intention as a Pareto-efficient optimization, as max tries to consider some worst-case instead of a naive average. However, continuing the above example, a ranking by max yields the ordering RP_3, RP_1, RP_2, as the max operator is unaware of the extremely small deviations in the RORI dimension and it cannot comprehend contextual information that a *RORI* evaluation deviating around 0.9 is (unfortunately) currently the only viable option. In contrast, the Pareto-efficient optimization is aware of this context, and, inherently, optimizes against all dimensions, i.e., it acknowledges the current situation and considers that a *RORI* around 0.9 is currently the only viable option.

Even though the proposed Pareto-efficient optimization delivers highly beneficial results, some technical conflicts of different mitigation actions in response plans may retain and are discussed in the following subsection.

8.2 Limitations of the proposed approach

An aspect to be discussed is related to conflicts among individual mitigation actions in response plans. Every mitigation action is associated to a generic type (e.g., patching, restart, shutdown), and each mitigation action type has an associated restriction (e.g., mutually exclusive, totally restrictive, partially restrictive). For instance, an action that suggests to shutdown equipment E_1 is totally restrictive with any other action associated to E_1 but it can be perfectly combined with actions to be implemented on another equipment as long as their implementation does not interfere with the normal operation of equipment E_1.

Conflicts of restrictive mitigation actions are assumed to be avoided at the first stage of the evaluation process, i.e., we assume that conflicting mitigation actions are never proposed as a response plan which is evaluated by RFIA or ROIA. However, our approach may not comprehend external restrictions on proposed response plans, i.e., RFIA and ROIA only perform syntactical verifications on response plans and do not implement any semantical validation of their implementability. For example, our approach cannot comprehend a situation where (i) a selected response plan requires to implement a mitigation action that is already activated in the system, i.e., an indirect increase of financial costs, or (ii) a selected response plan requires to deactivate an action that was previously active, i.e., a (potential) increase or decrease of operational impact and financial impact. Moreover, our approach does not consider semantic implications of individual mitigation actions. This means that, if a response plan is proposed which stands in conflict to the above-discussed restriction of shutdown of E_1, both RFIA and ROIA will still evaluate the response plan. Nevertheless, all of the beforementioned issues are directly resolved by an adequate approach of proposing response plans for evaluation.

In situations where not enough initial information was acquirable, RFIA may lead to inaccurate results. The reason for this is that any response plan in which a mitigation actions is missing information (e.g., cost, benefit, coverage) is directly discarded. To overcome this issue, future work is dedicated to find adequate heuristics and estimations for missing values in the RFIA evaluation. As ROIA is based on probabilistic inference, estimations for missing information is directly included and already considered in our models.

To acquire all parameters and models for the RFIA requires a great level of accuracy in estimations, i.e., requires detailed analyses of, e.g., monetary values. As the RFIA is based on direct calculations on acquired parameters, results can only be as accurate as the forecasts of loss event frequencies on which they rely. A significant amount of parameters, e.g., the annual loss expectancy (ALE), and the effectiveness (EF) used to compute the risk mitigation level, rely on expert knowledge, which will involve human errors. Effectively, an RFIA requires various estimated parameters whose kind is very similar to classical business economical operating figures. This means that, on the one hand side, a large amount of expert knowledge is required and many parameters must be manually assessed by experts. However, on the other hand side, experts from which this knowledge is acquirable, are trained business experts that are deeply familiar with this kind of estimation: estimating business operating figures. Even though the targeted security of an IT infrastructure is outside of the expert's subject area, the kind of required knowledge lies in their expertise.

A critical point to be discussed is that many calculations involved in the RFIA may seem simplistic, as they represent “simple” multiplications or additions of some business economical operating figures. On first sight, this is true, and, e.g., ROSI and RORI are classical adaptations of business economical key scores such as the return on investment (ROI). It remains true that, from a computer science perspective, these calculations are trivial and simple multiplication and additions do not provide an increase information content. However, these economical key scores such as ROI, ROA, ROIC, RONA, ROC, and ROCE [48] play a major role in business economics and are widely accepted as reference values. In addition with the previous paragraph, this provides an increased acceptance of our approach with the experts from which information is acquired.

It may not be computationally feasible to evaluate all proposed response plans, and a heuristic must be employed to prune the evaluation space. As discussed throughout this article, RFIA and ROIA targeted different impact dimensions whose performance criteria are contradictorily defined. We show throughout this article that their combination is highly beneficial and required. However, if both assessment perform a pre-pruning by their own standards, it is likely that disjoint subsets are evaluated. Therefore, a pre-pruning step must be based on a heuristic optimized against characteristics of RFIA and ROIA simultaneously, which is subject to future work.

Finally, the proposed ROIA is based on a model in which time is represented by fixed timeslices, where each timeslice is independent of each other. This implies that an impact of, e.g., the first timeslice is not considered anymore in the impact assessment of the second timeslice. To overcome this characteristic, the local impact probabilities must be adequately designed, as, e.g., proposed throughout this article. Another possibility is to utilize a time-dependent model, where the respective model of a timeslice t is dependent on nodes of the model representing timeslice $t-1$. Doing so creates a form of a dynamic Bayesian networks in which a near-realtime analysis of impacts is provided. Furthermore, such a network allows one to analyze chains of events that lead to an impact in retrospective. The design and mathematical implications of such a time-dependent dynamic impact assessment is considered by Motzek et al. in [11] utilizing so-called activator dynamic Bayesian networks [49].

9 Conclusions

We have proposed an approach for selecting adequate response plans as a reaction to threats opposed on a company based on a multi-dimensional impact assessments. On the one hand, we utilize a response financial impact assessment (RFIA) based on a cost-sensitive metric (i.e., return on response investment) and a geometrical tool (i.e., attack volume model). On the other hand, we utilize a response operational impact assessment (ROIA) based on mission and resource dependency models. The decision rule for the RFIA is that the higher the RORI value, the more interesting the response plan, whereas for the ROIA, the higher the impact values for the short-term (OI_0), medium term (OI_1), and long term (OI_2), the less interesting the response plan. We have shown that on most response plans only one dimension dominates, i.e., no clear optimal choice is present.

Based on a multi-dimensional minimization approach, we propose the choice of a Pareto-efficient response plan that bears the highest financial attractiveness on return on investment, and the lowest probability of conflicting with a company's missions. This is beneficial for applications, where highly critical missions and resources must be protected, without sacrificing missions in favor of security.

Future work is dedicated towards an evolution of the RORI metrics considering evaluations of multiple response plans simultaneously. Furthermore, future work is dedicated towards advancing operational impact assessments towards domains over near-continuous time dimensions for predictive and retrospective analyses.

Acknowledgements
The research in this paper has received funding from the PANOPTESEC project, as part of the Seventh Framework Programme (FP7) of the European Commission (GA 610416). We acknowledge financial support by Land Schleswig-Holstein within the funding programme Open Access Publikationsfonds.

Authors' contributions
All authors contributed equally to the manuscript. In addition, AM and GGG carried out the simulations and drafted the initial version of the manuscript. All authors read and approved the final manuscript.

Competing interests
The authors declare that they have no competing interests.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Author details
[1]SAMOVAR, Telecom SudParis, CNRS, Université Paris-Saclay, 9 Rue Charles Fourier, 91000 Evry, France. [2]Universität zu Lübeck, Institute of Information Systems, Ratzeburger Allee 160, 23562 Lübeck, Germany.

Received: 14 February 2017 Accepted: 18 June 2017
Published online: 10 July 2017

References

1. G Dini, M Tiloca, in *ETFA2013: 18th Conference on Emerging Technologies & Factory Automation*. On simulative analysis of attack impact in Wireless Sensor Networks (IEEE, Cagliari, 2013), pp. 1–8
2. D Kundur, X Feng, S Liu, T Zourntos, KL Butler-Purry, in *SmartGridComm: 1st International Conference on Smart Grid Communications*. Towards a framework for cyber attack impact analysis of the electric smart grid (IEEE, Gaithersburg, 2010), pp. 244–249
3. P Su, X Chen, H Tang, in *3rd International Conference on Innovative Computing Information and Control*. DoS attack impact assessment based on 3GPP QoS indexes (IEEE, Dalian, 2008), p. 103
4. KI Sgouras, AD Birda, DP Labridis, in *ISGT2014: Innovative Smart Grid Technologies Conference*. Cyber attack impact on critical smart grid infrastructures (IEEE, Washington, 2014), pp. 1–5
5. BW Roberts, The macroeconomic impacts of the 9/11 attack: evidence from real-time forecasting. Peace Economics, Peace Science and Public Policy 15.2 (2009)
6. G Gonzalez-Granadillo, A Motzek, J Garcia-Alfaro, H Debar, in *ARES2016: 11th International Conference on Availability, Reliability, and Security*. Selection of mitigation actions based on financial and operational impact assessments (IEEE, Salzburg, 2016), pp. 137–146
7. G Gonzalez-Granadillo, M Belhaouane, H Debar, G Jacob, RORI-based countermeasure selection using the OrBAC formalism. Int. J. Inf. Secur. **13**(1), 63–79 (2014)
8. G Gonzalez-Granadillo, H Debar, G Jacob, L Coppolino, in *INTECH2012: International Conference on the Innovative Computing Technology*. Combination approach to select optimal countermeasures based on the RORI index (IEEE, Casablanca, 2012), pp. 38–45
9. G Gonzalez-Granadillo, D Debar, G Jacob, C Gaber, M Achemlal, in *MMM-ACNS2012: International Conference Mathematical Methods, Models and Architectures for Computer Network Security,*. Individual countermeasure selection based on the return on response investment index (Springer, St. Petersburg, 2012), pp. 156–170
10. A Motzek, R Möller, M Lange, S Dubus, in *NATO IST-128 Workshop: Assessing Mission Impact of Cyberattacks*. Probabilistic mission impact assessment based on widespread local events (NATO IST, Istanbul, 2015), pp. 16–22

11. A Motzek, R Möller, Context- and bias-free probabilistic mission impact assessment. Comput.Secur. **65**, 166–186 (2017). ISSN 0167-4048. doi:10.1016/j.cose.2016.11.005
12. A Motzek, C Geick, R Möller, in *CBI2016: 18th IEEE Conference on Business Informatics*. Semantic normalization and merging of business dependency models, (Paris, 2016), pp. 7–15. doi:10.1109/CBI.2016.10
13. TL Saaty, What is relative measurement? The ratio scale phantom. Math.Comput. Model. J. **17**(4-5), 1–12 (1993)
14. W Sonnenreich, J Albanese, B Stout, Return on security investment (ROSI)-a practical quantitative model. J. Res. Pract. Inf. Technol. **38.1**, 45–56 (2006)
15. M Jeffrey, in *Return on investment analysis for e-business projects*, ed. by H Bidgoli. Internet Encyclopedia, vol. 3 (Wiley, 2004), pp. 211–236. doi:10.1002/047148296X.tie154
16. Lockstep Consulting, A guide for government agencies calculating return on security investment, Technical Paper (2004)
17. M Schmidt, *Return on investment (ROI): meaning and use*. (Encyclopedia of Business Terms and Methods, 2011). available at: https://www.business-case-analysis.com/return-on-investment.html. Accessed 26 June 2017
18. J Brocke, G Strauch, C Buddendick, in *ISTA: 6th International Conference of Information Systems Technology and its Applications*. Return on security investment—design principles of measurement system based on capital budgeting, vol. 107 (LNI, Kharkiv, 2007), pp. 21–32
19. N Kheir, N Cuppens-Boulahia, F Cuppens, H Debar, in *ESORICS2010: 15th European Symposium on Research in Computer Security, Athens, Greece*. A service dependency model for cost-sensitive intrusion response (Springer, Athens, 2010), pp. 626–642
20. G Jakobson, in *Fusion2011: 14th International Conference on Information Fusion*. Mission cyber security situation assessment using impact dependency graphs (IEEE, Chicago, 2011), pp. 1–8
21. J Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. (Morgan Kaufmann, 2014)
22. M Henrion, in *UAI1987: 3rd Conference on Uncertainty in Artificial Intelligence*. Practical issues in constructing a Bayes' belief network (AUAI, Seattle, 1987), pp. 132–139
23. JG Torres-Toledano, LE Sucar, in *IBERAMIA 98: 6th Ibero-American Conference on AI*. Bayesian networks for reliability analysis of complex systems (Springer, Lisbon, 1998), pp. 195–206
24. LD Raedt, A Kimmig, H Toivonen, in *IJCAI2007: 20th International Joint Conference on Artificial Intelligence*. ProbLog: a probabilistic prolog and its application in link discovery (AAAI, Hyderabad, 2007), pp. 2462–2467
25. TL Norman, *Risk analysis and security countermeasure selection*. (CRC Press, Taylor & Francis Group, 2010)
26. G Gonzalez-Granadillo, J Garcia-Alfaro, H Debar, Using a 3D geometrical model to improve accuracy in the evaluation and selection of countermeasures against complex cyber attacks. Secur. Priv. Commun. Netw. **164**, 538–555 (2015)
27. G Gonzalez-Granadillo, H Debar, G Jacob, in *CRiSIS2015: 10th International Conference on Risks and Security of Internet and Systems*. Attack volume model: geometrical approach and application (Springer, Mytilene, 2015), pp. 242–257
28. SJ Russell, P Norvig, *Artificial intelligence—a modern approach (3. internat. ed.)* (Pearson Education, 2010). ISBN 978-0-13-207148-2
29. G Combs, The Wireshark Foundation, *Wireshark*, (2017). Retrieved 13.02.2017 from http://www.wireshark.org/. Accessed 26 June 2017
30. J Pearl, *Causality: models, reasoning and inference*, 2nd edn. (Cambridge University Press, New York, 2009)
31. A Motzek, R Möller, in *NATO IST-148 Symposium on Cyber Defence Situation Awareness, Sofia, Bulgaria*. Probabilistic mission defense and assurance, (2016), pp. 4–1–4-18. doi:10.14339/STO-MP-IST-148
32. V Viduto, C Maple, W Huang, D López-Pérez, A novel risk assessment and optimisation model for a multi-objective network security countermeasure selection problem. Decis. Support. Syst. **53**(3), 599–610 (2012)
33. B Foo, Y Wu, Y Mao, S Bagchi, EH Spafford, in *DSN2005: International Conference on Dependable Systems and Networks, Yokohama, Japan, 28 June - 1 July, 2005*. ADEPTS: adaptive intrusion response using attack graphs in an e-commerce environment (IEEE, Yokohama, 2005), pp. 508–517
34. R Dewri, N Poolsappasit, I Ray, D Whitley, in *CCS2007: ACM Conference on Computer and Communications Security*. Optimal security hardening using multi-objective optimization on attack tree models of networks (ACM, Alexandria, 2007), pp. 204–213
35. SA Zonouz, H Khurana, WH Sanders, TM Yardley, RRE: a game-theoretic intrusion response and recovery engine. IEEE Trans. Parallel Distrib. Syst. **25**(2), 395–406 (2014)
36. S Bistarelli, M Dall'Aglio, P Peretti, in *FAST2006: 4th International Workshop on Formal Aspects in Security and Trust*. Strategic games on defense trees (Springer, Hamilton, 2006), pp. 1–15
37. A Fielder, E Panaousis, P Malacaria, C Hankin, F Smeraldi, Decision support approaches for cyber security investment. Decis. Support. Syst. **86**, 13–23 (2016)
38. A Roy, DS Kim, KS Trivedi, in *DSN2012: IEEE/IFIP International Conference on Dependable Systems and Networks*. Scalable optimal countermeasure selection using implicit enumeration on attack countermeasure trees (IEEE, Boston, 2012), pp. 1–12
39. N Stakhanova, C Strasburg, S Basu, JS Wong. J. Comput. Secur. **20**(2-3), 169–198 (2012)
40. R Alavi, S Islam, H Mouratidis, An information security risk-driven investment model for analysing human factors. Inf. Comput. Secur. **24**(2), 205–227 (2016)
41. L Wang, Ta Islam, T Long, A Singhal, S Jajodia, in *Data and Applications Security XXII, 22nd Annual IFIP WG 11.3 Working Conference on Data and Applications Security*. An attack graph-based probabilistic security metric (Springer, London, 2008), pp. 283–296
42. L Yu, H Man, in *Data Mining, Intrusion Detection, Information Assurance, and Data Networks Security*. Network vulnerability assessment using Bayesian networks (SPIE, Orlando, 2005), pp. 61–71
43. P Xie, J Li, X Ou, P Liu, R Levy, in *DSN2010: International Conference on Dependable Systems and Networks*. Using Bayesian networks for cyber security analysis (IEEE/IFIP, Chicago, 2010), pp. 211–220
44. N Kheir, H Debar, N Cuppens-Boulahia, F Cuppens, J Viinikka, in *International Conference on Network and Service Security*. Cost evaluation for intrusion response using dependency graphs (IEEE, Paris, France, 2009), pp. 1–6
45. J Marko, C Thul, P Martini, in *LCN2007: 32nd IEEE Conference on Local Computer Networks*. Graph based metrics for intrusion response measures in computer networks (IEEE, Dublin, 2007), pp. 1035–1042
46. A Barreto, P Costa, E Yano, in *STIDS2012: 7th International Conference on Semantic Technologies for Intelligence*. A semantic approach to evaluate the impact of cyber actions to the physical domain (CEUR, Fairfax, 2012), pp. 64–71
47. A Barreto, P Costa, E Yano, in *STIDS2013: 8th International Conference on Semantic Technologies for Intelligence*. Using a semantic approach to cyber impact assessment (CEUR, Fairfax, 2013), pp. 101–108
48. PW Farris, N Bendle, P Pfeifer, D Reibstein, Marketing metrics: the definitive guide to measuring marketing performance, Pearson Education (2010)
49. A Motzek, R Möller, in *IJCAI2015: 24th International Joint Conference on Artificial Intelligence*. Indirect causes in dynamic Bayesian networks revisited (AAAI, Buenos Aires, 2015), pp. 703–709
50. M Bastian, S Heymann, M Jacomy, in *International AAAI Conference on Weblogs and Social Media*. Gephi: an open source software for exploring and manipulating networks, (2009)