Selection of mitigation actions based on financial and operational impact assessments

Gustavo Daniel Gonzalez Granadillo, Alexander Motzek, Joaquin Garcia-Alfaro, Hervé Debar

Selection of Mitigation Actions Based on Financial and Operational Impact Assessments

Gustavo Gonzalez Granadillo
Joaquin Garcia-Alfaro
Hervé Debar
Institut Mines Telecom, Telecom SudParis
CNRS UMR 5157 SAMOVAR, Evry, France
{first_name.last_name}@telecom-sudparis.eu

Alexander Motzek
Universität zu Lübeck
Institute of Information Systems,
Ratzeburger Allee 160, 23562 Lübeck, Germany
motzek@ifis.uni-luebeck.de

***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 relies on a response system that continuously quantifies risks, and decides how to respond to cyber-threats that target a monitored ICT system. Our Dynamic Risk Management Response (DRMR) model is composed of two main modules: a Response Financial Impact Assessor (RFIA), which provides an assessment concerning the potential financial impact that responses may cause to an organization; and a Response Operational Impact Assessor (ROIA), which assesses potential impacts that efficient mitigation actions may cause on the organization in an operational perspective. As a result, the DRMR model proposes response plans to mitigate identified risks, enable choice of the most suitable response possibilities to reduce identified risks below an admissible level while minimizing potential negative side effects of deliberately taken actions.**

I. Introduction

The impact of an event is defined as the magnitude of harm that is expected to be perceived by an organization as a result of the consequences from unauthorized disclosure, modification, destruction, unavailability, or loss of information [1]. It may be expressed relative to the nature of its consequences. A way of expressing such nature is to bind the consequences on security dimensions. Commonly admitted natures for impacts on Information Systems are: Confidentiality, Integrity and Availability (CIA).

Current research focus on considering the impact of attacks [2]–[5], by evaluating their severity and consequences, leaving aside the impact of security actions in mitigating the effects of such attacks. However, the analysis of current cyber events should also consider the impact of potential mitigation actions as well as time, geographic space and affected elements [6].

We adopt a quantitative risk-aware approach that considers the likelihood of success of the detected attacks, their induced impact, and the cost and consequences of response plans onto a higher goal, e.g., a company or mission. The research presented in this paper represents a work in progress towards the development of a comprehensive and practical dynamic risk management response system.

Our model considers two approaches: A Financial Impact Assessment (FIA) and an Operational Impact Assessment (OIA). A response FIA (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, attacks, errors) in a financial perspective, and to deploy mitigation actions accordingly. A response OIA (ROIA) considers that mitigation actions, while highly effective, could lead to 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. This is beneficial for applications, where highly critical missions and resources must be protected, without sacrificing missions in favor of security.

The rest of the paper is structured as follows: Sections II and III discuss preliminaries and theory of a financial- and operational-impact assessment. Section IV describes a proposed dynamic risk management response model based on financial- and operational impact-assessments. A real world application of the proposed system is presented in Section V showing the applicability of the proposed model. Section VI discusses related work and we conclude in Section VII.

II. Financial Impact Assessment

Cost sensitive metrics have been proposed as a viable approach to find an optimal balance between intrusion damages and response costs, and to guarantee the choice of the most appropriate response without sacrificing the system functionalities. Measurements are either absolute or relative. Absolute measurements use precise values that scale with a given unit (e.g. hundreds, thousands, millions, etc); whereas relative measurements are methods for deriving ratio scales from paired comparisons represented by absolute numbers [7]. 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 [8].

Examples of these models include the Return On Investment (ROI) and all its variants [8]–[10].

For the scope of this article, we consider sets of individual actions performed as a response to an adversary. Sets of these actions are called response plans:

Definition 1 (Response Plan)**.** *A response plan RP is a vector of mitigation actions, representing individual actions to be performed as a response to an adversary or threat opposed to an organization.* ▲

The Return On Response Investment (RORI) is a quantitative model for cost sensitive response based on a financial comparison of the response plans [11], [12]. RORI is an adaptation of the Return On Security Investment (ROSI) index, that provides a qualitative comparison of response plans. The RORI index considers not only the intrusion impact but also the effect of response plans, as shown in Equation 1.

$$RORI = \frac{(ALE \cdot RM) - ARC}{ARC + AIV} \cdot 100 \tag{1}$$

All parameters are defined as follows:

Definition 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), i.e.,*

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

▲

Definition 3 (Annual Infrastructure Value, AIV)**.** *AIV represents the fixed costs that are expected to be perceived by an organization regardless of the implemented mitigation action. 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), i.e.,*

$$AIV = C_e + C_p + C_s + C_o + V_r \ . \qquad ▲$$

Definition 4 (Risk Mitigation, RM)**.** *RM refers to the risk mitigation associated with a given mitigation action. RM takes values between zero and one hundred percent (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 (MC, 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), i.e.,*

$$RM = MC \cdot EF \ . \qquad ▲$$

Definition 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}) and indirect costs (C_i) that may originate from the adoption of a particular mitigation action, i.e.,*

$$ARC = C_{impl} + C_{maint} + C_{od} + C_i \ . \qquad ▲$$

Considering a RORI index alone, the best candidate response set is represented by a maximal positive RORI index.

III. Operational Impact Assessment

An operational impact assessment 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. Understanding these impacts is a pertinacious problem and current work uses adhoc solutions based on handcrafted algorithms. While such approaches deliver early results, their assessments need to be verified and validated by large amounts of data—which is not always available.

Motzek et al. introduce an approach towards OIA based on a probabilistic graphical model in [26], which defines a well-understood problem on which an OIA can be reduced. 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. It is this local view, which allows for a validation of defined data. This means, assessments from experts can be used directly without global normalization factors and experts are not forced into expertise which they can not understand.

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

Remark 1 (Impact)**.** *An abstract term of "impact" is used in this work in the sense of "not operating as fully intended". The underlying meaning of "intended operation" lies in an use case of the model.* ▲

A. Mission Dependency Model (Business View)

Motzek et al. [26] extend a model by Jakobson [22] and model mission dependencies as shown in Figure 1 as a graph of *mission nodes* (MN). 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 Figure 1), say, "provide access to customer data", failing, given required business-resource A, e.g., "customer-data-frontend", fails is 90%. [26] 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.

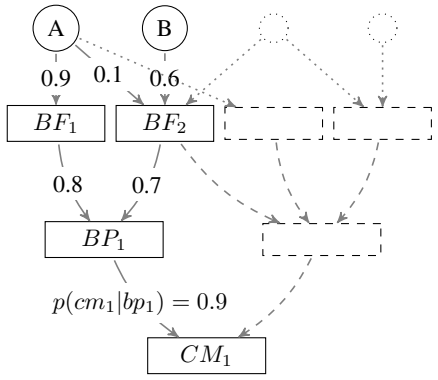

Fig. 1. Mission Dependency Model. Values along edges denote individual conditional probability fragments.

Definition 6 (Probabilistic Preliminaries)**.** *A node of a probabilistic dependency model is a random variable, denoted as capital X. Every random variable is assignable to one of its possible values $x \in \mathrm{dom}(X)$. Let $P(X = x)$ denote the probability of random variable X having x as a value. For the case $\mathrm{dom}(X) = \{true, false\}$ we write $+x$ for the event $X = true$ and $\neg x$ for $X = false$.* ▲

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

Definition 7 (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 $Y \in \vec{Y}$, [26] employs a non-leaky noisy-or combination function as described in [23], [24].* ▲

With Definition 7, 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.

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 can not be expected that an underlying distributed computing cluster is identified to actually provide this web-service. The following resource dependency model covers these dependencies.

B. Resource Dependency Model (Operation View)

Critical resources identified in a mission dependency model are *dependent* on further resources. Likewise, if a dependent resource is threatened, the identified critical resource might be threatened *transitively* as well. An operation expert, unlike a business expert, has an expertise to understand such dependencies, which we cover in an resource dependency model. The resource dependency model models dependencies between individual resources, which can be, e.g., individual ICT servers, ICS devices, software components or, in other use cases, manufacturing robots, suppliers, soldiers or vehicles. A "Bayesian" approach is followed as before, meaning that every dependency between two resources represents a local conditional probability of impact, if the dependence is impacted, as shown in Figure 2.

[26] argues that assessing resource dependencies is not manageable by hand. Complex operation structures render a manual dependency analysis infeasible and error prone. Further, dynamically adjusting infrastructures (e.g., as found in IT cloud use cases) make it even unknown to an expert to identify exact dependencies. However, [26] argues that an expert is able to validate a presented infrastructure dependency model for plausibility. Therefore, a solution based on heuristics from exchanged information amounts are proposed to obtain a resource dependency model, for which we present an example in Section IV-B.

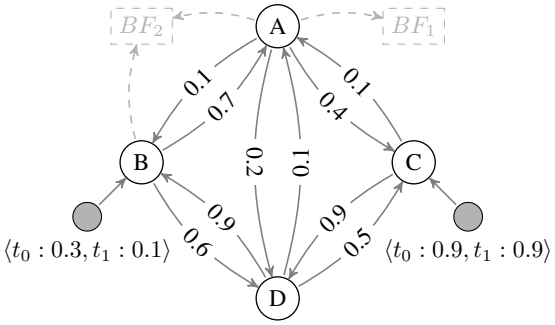

Fig. 2. A minimalistic resource dependency model. Conditional probability fragments are marked along the edges. Grey nodes represent external shock events leading to local impacts on resources. Connections to the mission dependency model are sketched in dashed gray.

C. Local Impacts (Security View)

Nodes of a resource dependency model might threatened directly by, so-called, external shock events. A security expert has the expertise to assess the local consequences on a node, given the presence of an shock event, e.g., the presence of a vulnerability or a direct shutdown of a node. An external shock event $SE \in \vec{SE}$ is a random variable and is present ($+se$) or not present ($\neg se$), for which a prior random distribution $P(SE)$ is defined. Every node X might be affected by one or

more external shock events $\vec{SE}$. Accordingly, the presence of an external shock event can be known or can be unclear and is assessed probabilistically through its prior random distribution $P(SE)$. The set of observed external shock events (known presence) is a set of instantiations $\vec{se}$ of observed random variables $\vec{SE_O} \subseteq \vec{SE}$. In the case that an external shock event is present (se), there exists a probability of it affecting a node X, expressed as a conditional probability fragment $p(+x|+se)$. If an external shock event exists and it is not inhibited, [26] speaks of a *local impact* on X. In the case that the external shock event is not present, i.e., $\neg se$, it does not affect random variable X. Every individual conditional probability fragment contributes to a random variable's CPD in the same way as a dependance on other nodes.

Definition 8 (Temporal Aspects)**.** *[26] defines a temporal aspect of an external shock event. In an abstract timeslices an effect of an external shock event changes. Every abstract timeslice represents a duplicate of the network- and mission dependencies with a different set of local conditional probabilities and prior probabilities of shock events. A time-varying probability is denoted as a sequence $\langle t_0 : p_0, \ldots, t_T : p_T \rangle$, with $T + 1$ abstract timeslices. In every abstract timeslice i, varying probabilities take their respective conditional or prior probability p_i defined for its timeslice t_i.* ▲

Note that a security expert does neither need to have any expertise in dependency analyses nor in business process analyses. An assessment of potential impacts is performed using a local, causal, view on resources and direct causes as external shock events.

D. 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.

Definition 9 (Mission Impact Assessment, MIA)**.** *The probability of a mission node MN being impacted, is defined as the conditional probability of MN being impacted $+mn$ given all observations of external shock events $se \in \vec{se}$, i.e. $P(+mn|\vec{se})$, where the effects of local impacts due to $\vec{se}$ are mapped globally based on mission-dependency and resource-dependency graphs. The mission impact assessment is therefore defined as the problem of obtaining $P(+mn|\vec{se})$, for all mission nodes MN defined defined in the mission dependency model.* ▲

Probabilistic inference is generally known to be NP-hard, and an exact solution for the MIA problem is 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, one can see the probabilistic model as 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. $P(+mn|\vec{se})$ is then the probability that at least one proof holds, or rather, the probability that the disjunction of conjunctions is satisfied, i.e.

$$P(+mn|\vec{se}) = \bigcup_i \mathbf{P}(w_i^{MN}) = P(\vec{w}^{MN}) = P(\{\bigvee_i w_i^{MN}\}) \ ,$$

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 [28]. 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 probabilistic MIA $P(+mn|\vec{se})$ directly originates from all defined dependency-models and represents an inference problem in a probabilistic graphical model. Therefore, [26] argues that if locally defined dependency-models are validated to be correct, an obtained impact assessment $P(+mn|\vec{se})$ is validated, too.

IV. Dynamic Risk Management Response System

We developed a system that proposes response possibilities to mitigate identified risks, enable choice of the most suitable response possibilities to reduce the identified risks below an admissible level, and then, compute the mitigation actions to be deployed on monitored or protected ICT systems. We adopt a quantitative risk-aware approach that provides a comprehensive view of the threats, by considering, (i) the likelihood of success of the considered attacks, (ii) their induced impact, and (iii) the cost and impact of the possible responses. Our dynamic response model is composed of two main modules: the Response Financial Impact Assessor (RFIA), and the Response Operational Impact Assessor (ROIA), as shown in Figure 3.

The RFIA receives input data regarding the severity and likelihood of the potential threats, benefits and cost of mitigation actions, default security policies, and all elements that could be affected in the exploitation of the threat (e.g., users, channels, resources). The RFIA performs the evaluation of individual and combined mitigation actions and creates response plans based on the selected candidates. Such plans are sent to

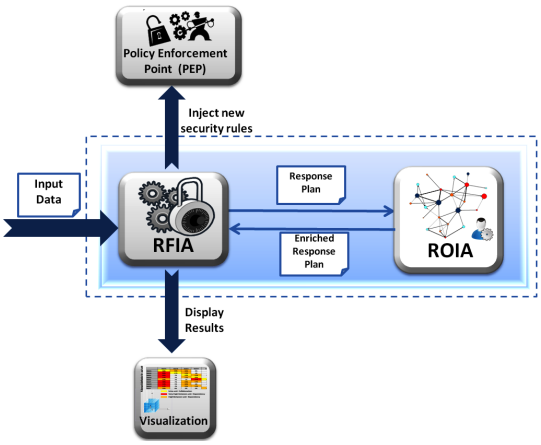

Fig. 3. Internal block diagram for the DRMRS components

the ROIA for evaluation. For every proposed response plan, the ROIA provides an operational impact assessment.

Assessed response plans are propagated to a visualization module and to an policy enforcement point, which transforms mitigation actions into security policies. A response plan selection strategy proposed in this paper used to assist a security operator in selecting the most suitable response plan.

A. Response Financial Impact Assessor (RFIA)

The Response Financial Impact Assessor quantifies the level of benefit perceived per response plan on a financial basis. It provides an assessment concerning the potential financial impact that a given response plan may cause to an organization. Response plans represent proposed mitigation of the assessed risks and are assumed to be composed of one or more mitigation actions. The RFIA is composed of two main components: the Return On Response Investment (RORI), and the Attack Volume (AV), as depicted in Figure 4.

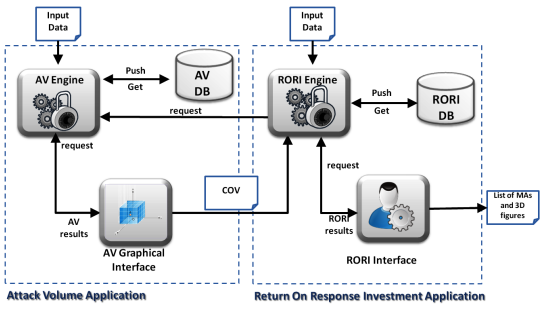

Fig. 4. Internal block diagram for the RFIA component

From Figure 4, the RORI interface requests the RORI Engine to perform the evaluation of mitigation actions (e.g., individual and combined) for a given threat scenario. The RORI Engine requests the input parameters (i.e., ALE, AIV, ARC, RM) to the RORI database through a get command. If the ALE or the RM are missing for that particular security incident, the RORI engine will request the AV Engine to perform the attack volume evaluation and to provide the corresponding values.

The AV Engine requests the input data to particular services (e.g., LDAP, databases, ACL, servers) and retrieves the associated RCU (Resource-Channel-User) information in order to calculate its volume and plot its graphical representation. The retrieved RCU data is stored in the AV database and displayed in the AV interface.

Upon reception of all the parameters, the RORI engine stores them in the RORI database through a push command, and performs the evaluation of the authorized mitigation actions.

1) Return On Response Investment (RORI): is a relative index that indicates the level of benefit perceived if a given mitigation action is implemented. The input parameters for the RORI calculation are of two kinds: fixed parameters include the Annual Infrastructure Value (AIV), which depends on the system, and the Annual Loss Expectancy (ALE), which characterizes the intrusion or attack; variable parameters include the Risk Mitigation (RM) and Annual Response Cost (ARC) which express the costs related to a mitigation action. RORI is calculated according to Equation 1.

The RORI index is used to evaluate optimal plans, by ranking them as a trade-off between their efficiency in stopping potential attacks, and their ability to preserve, at the same time, the best service to legitimate users. Details on the estimation of each parameter composing the RORI model can be found in [12].

2) Attack Volume (AV): is a graphical tool that evaluates 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 a parallelepiped in three dimensions. The resulting volume is computed as the product of the axes contribution to the execution of the incident, i.e.,

$$AV(A) = Co_{Acc}(A) \times Co_{Ip-Port}(A) \times Co_{Res}(A) \quad (2)$$

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[1] methodology [13].

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)

[1]A multi-criteria methodology that considers Criticality, Accessibility, Recuperability, Vulnerability, Effect, and Recognizability in the evaluation

times its associated weighting factor (that results from the implementation of the CARVER methodology), i.e.,

$$Co_{Axis}(S) = \sum_{i=0}^{n} Count(E \in Type_{Axis}(S)) \times WF(Type_{Axis}(S)) \ . \quad (3)$$

The attack volume interface provides a 3D view of the complete attack scenario, making it possible to calculate the impact of multiple attacks that originate simultaneously in the system. In addition, we are able to compute the coverage of such attacks in the system, and the level of coverage for one or more response plans against the detected attack(s). Details on the computation of the system, attack and countermeasure volumes can be found in [14].

B. Response Operational Impact Assessor (ROIA)

The ROIA is divided into three components, which acquire, define and evaluate all information needed for an response operational impact assessment based on the described probabilistic model.

1) Network Dependency Analyzer: In order to obtain a resource dependency model automatically, we propose a module which consecutively captures traffic inside an organization and analyses statistics in them based on a heuristic. 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 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 $p(x|y_i) = \frac{T_{Y_i,X}}{\sum_i T_{Y_i,X}}$. Depending on a network or company characteristics other heuristics might be appropriate, e.g. derivation from a mean received amount of data or a mapping onto a σ distribution. As long as no irrelevant information is consumed and distributed to other resources, this heuristic results in a plausible resource dependency model. As part of the ROIA, we implement this approach in an automatic module, periodically capturing traffic and analyzing obtained results, which are presented in Section V.

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. 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 extended [26]'s proposed external shock event transformation from response plans:

Definition 10 (Response Plan Side Effects)**.** *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* shut down *(+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,* restart*ing 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.* ▲

Further examples for shock events are given by Motzek et al. in [26].

3) Monte-Carlo Evaluator: 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-Poincar 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.

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 $P(+mn|\vec{se})$ 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 [26] and left out for brevity in this paper.

C. Selection of Response Plans

Both, RFIA and ROIA, perform impact assessments of a proposed response plan as a collection of individual 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 potential cause collateral damage. On the other hand, a minimally invasive response plan, will not significantly reduce a risk. It is the novel advantage of the proposed DRMRS of being able to combine both assessments.

However, finding an optimal response plan in all dimensions defined by an OIA and a FIA is not trivial. The proposed FIA results in a linear, relative metric, i.e., assessments depend on a use-case and context and are only interpretable, evaluable and comparable during one evaluation of a set of response plans. Still, among one evaluation there exists a well-defined ordering. However, relative reference points are required for obtaining an absolute scale from one evaluation.

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 evaluation. 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. However, an OIA consists an n-dimensional vector representing a temporal diversity.

Due to a missing absolute scale in FIA and an assumed incomparability of temporal dimensions, an optimization goal by a defined cost function is not available. We therefore propose a selection of response plans based on a best compromise, i.e., a semi-optimal solution among all impact assessment dimensions, related to a Pareto optimum.

Definition 11 (Semi-optimal 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 semi-optimal 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 a best compromise among an n-dimensional impact assessment is therefore defined as finding the smallest semi-optimal set.

Definition 12 (Smallest semi-optimum)**.** *Let $\vec{d}$ be the vector of all impact dimensions. Then, the smallest semi-optimal set of response plans $\mathring{RP}$ is the set*

$$\mathring{RP} = \min_{\varepsilon} \left(\left\{ \bigcap_{d \in \vec{d}} \dot{RP}^{d}_{\varepsilon} \right\} \neq \emptyset \right) \quad . \quad \blacktriangle \tag{4}$$

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 12 can efficiently be performed by using a binary search.

V. Use Case

This section studies an application of the proposed DRMRS in an infrastructure environment of an Energy Distribution Organization. The environment consists of a distributed network of remote terminal units (RTU) in energy stations of medium voltage (MV = 20,000 Volts) and high voltage (HV = 150,000 Volts). RTUs acquire data from electrical equipments (e.g., PLC, sensors, etc.), 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, 6,000 of which are controlled by the STU.

For testing purposes, we emulated the energy distribution organization (EDO) using the information shown in Table I.

TABLE I
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

From Table I, we organize the information of the EDO according to their nature (dimension). We have for instance, 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.

The annual infrastructure value for the EDO is equivalent to 11,379,800.00 , 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.

A. Threat Scenario

Security experts from the use case partner identified a threat, called "AS02", which corresponds to a compromise of a specific target through vulnerability exploitation. More precisely, the threat will cause data corruption or leakage of a database in the ICT domain, e.g., from a file server.

There exists an attack vector via the ICT Network from, e.g., a vulnerable router "VR-08", targeting first a Web server "Web-SRV", second, a workstation "User-PC", and third, a file server "File-SRV".

B. Financial Impact Assessment

Threat AS02 has a severity defined as "serious", which corresponds to 1,000,000 €, and a likelihood defined as "high", which corresponds to a value of 12. The ALE is computed as 12,000,000 €/year. This threat has been associated to a set of mitigation actions. Combinations of associated mitigation actions form response plans that 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 deployment of actions (e.g., patch deployment). The detailed information of all authorized mitigation actions is shown in Table II.

TABLE II
MITIGATION ACTION RFIA EVALUATION

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	1,000	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	2,000	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

Table II summarizes the information about mitigation actions that are authorized as a response to the specified threat AS02. ARC and EF of each security mitigation action were estimated based on expert knowledge and historical data. The RM value is calculated as the product of the EF and coverage (COV). A coverage is obtained using geometrical operations from the attack volume model. The RORI index is calculated using Equation 1.

From the list of proposed mitigation actions, MA1 (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} can not be simultaneously implemented with MA_2, MA_4, and MA_7). All potential combinations, i.e., 797 response plans are 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 $\check{RP}^{RFI} = 0.21$.

The graphical representation of the best response plan vs. the evaluated threat is depicted in Figure 5, where hashed lines represent threat AS02 and colored lines represent the mitigation actions.

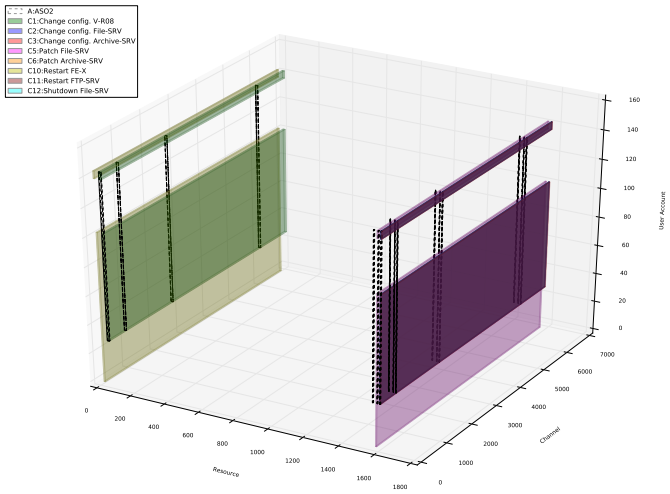

Fig. 5. Graphical representation of the threat and the best response plan.

C. Operational Impact Assessment

To perform a response operational impact assessment, a resource dependency model is needed. As described in Section III a manual assessment is said to be infeasible and a solution based on a heuristic of exchanged traffic information was proposed in Section IV-B. An evaluation of one week of traffic recordings inside the EDO resulted in a resource dependency model as shown in Figure 6 consisting of 35 nodes and 66 edges (dependencies); 470 other nodes resembled insignificant dependencies and were removed for visualization

and anonymization purposes. The average path length between two nodes resulted to be 2.5, s.t., a maximum search depth of 7 is likely to cover all paths.

Based on the resource dependency model and a mission dependency model defined by business experts from the use case partner, ROI assessments for all proposed response plans are evaluated. The mission dependency model for the EDO consists of four business processes and 26 identified critical resources, from which not all are yet simulated and not needed for the scope of this paper. For anonymization purposes and non-disclosure agreements, the mission dependency model can not be displayed here. A comparison between a RORI index and operational impact assessments in three temporal dimensions are given in Table III. Note how both lowest and highest probabilities of operational impact lead to extremely low RORI indices. In fact, Pearson's product-moment correlation coefficient between RORI and OI_0 and OI_1 for all evaluated response plan is ≈ 0.14 and between RORI and OI_2 even ≈ 0.01, showing that RORI and OI are almost uncorrelated.

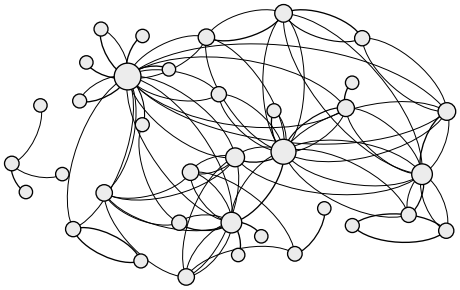

Fig. 6. Resource dependency model extracted from the use case partner. Thicker and darker edges represent higher dependency degrees. Visualized using Gephi [27].

TABLE III
MITIGATION ACTION ROIA EVALUATION. (*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%

D. Selection of Mitigation Actions

Judging from Table III a good compromise is to singly deploying mitigation action MA_1, resulting in both a low probability of operational impact and being financially attractive in terms of RORI. 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 IV-C we search for the best *compromise*: From Definition 12 one obtains the best semi-optimal response plan set $\mathring{RP} = \{\langle MA_1, MA_2, MA_4, MA_6, MA_7, MA_9 \rangle\}$ 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 semi-optimal response plan is found.

Notwithstanding, one could normalize all impact dimensions to a range between $[0, 1]$, where 1 represents the best (i.e., 0 for operational impact and the best evaluated RORI index for financial impact) and 0 the worst assessment and then define an equally weighted cost function f used for selection of an optimal response plan. Following such approach, one selects: $RP_f = \langle MA_1, MA_2, MA_4, MA_6, MA_9 \rangle$ with $\langle t_0 : 0.102\ t_1 : 0.09\ t_2 : 0.0\ rori = 82.7732 \rangle$, which is different from $\mathring{RP}$. The following example clarifies the difference: Say, there exists another response plan RP^* with an assessment $\langle t_0 : 0.191\ t_1 : 0.0\ t_2 : 0.0\ rori = 82.7732 \rangle$, i.e., an assessment similar to the one of RP_f, but where an impact of dimension t_1 is moved to t_0 with a small difference. RP^* would be assigned an even better cost than RP_f by f, but the t_0 assessments differs by $\varepsilon = 18.5\%$ of the theoretical optimum from the evaluated optimum in dimension t_0 instead of 14.75% as $\mathring{RP}$ does.

VI. RELATED WORK

Current researches focus 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 [2], for instance, propose a simulation framework that evaluates the impact of cyber-physical attacks, discusses the attack ranking process, and analyzes different mitigation actions. However, these latter are not considered in the calculation of the attacks' impact nor they are ranked according to their effectiveness in stopping or mitigating the attacks.

Kundur et al. [3], propose a paradigm for cyber attack impact analysis that employs a graph-theoretic structure and a dynamical systems framework to model the complex interactions amongst 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. [5] 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 researched as an adequate assessment of impacts or risks posed due to attacks or found vulnerabilities [15]–[17]. However, often imperfect knowledge is not considered [15] or dependency cycles pose a problem [17]. Other impact propagation approaches, able to handle such details, are not probabilistic based and degrade to a handcrafted propagation algorithm with arbitrary scores [18], [19].

Barreto et al. [20], [21] only consider direct impacts as approaches to mission modeling, leaving aside transitive impacts and/or defining a manual description of all dependencies between individual devices inside one organization, which is, in most of the cases an unfeasible 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.

VII. Conclusion

In this paper we have proposed an automatic response system, reacting to threats opposed on a company based on a multi-dimensional impact assessments. Two different impact assessment approaches have been incorporated, 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.

Based on a multi-dimensional minimization proposal, we propose the choice of semi-optimal 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 missions. This is beneficial for applications, where highly critical missions and resources must be protected, without sacrificing missions in favor of security.

Acknowledgment

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).

References

[1] R. Kissel,*Glossary of key information security terms*, National Institute of Standards and Technology, U.S. Department of Commerce, 2011.

[2] G. Dini, M. Tiloca,*On simulative analysis of attack impact in Wireless Sensor Networks*, 18th Conference on Emerging Technologies & Factory Automation (ETFA), 2013.

[3] D. Kundur, X. Feng, S. Liu, T. Zourntos, K.L. Butler-Purry, *Towards a Framework for Cyber Attack Impact Analysis of the Electric Smart Grid*, International Conference on Smart Grid Communications (SmartGridComm), pp. 244–249, 2010.

[4] P. Su, X. Chen, H. Tang, *DoS Attack Impact Assessment based on 3GPP QoS Indexes*, 3rd International Conference on Innovative Computing Information and Control, 2008.

[5] K. I. Sgouras, A. D. Birda, D. P. Labridis, *Cyber Attack Impact on Critical Smart Grid Infrastructures*, Innovative Smart Grid Technologies Conference (ISGT), 2014.

[6] B. Roberts, *The Macroeconomic Impacts of the 9/11 Attack: Evidence from Real-Time Forecasting*, Working Paper, Homeland Security, Office of Immigration Statistics, 2009.

[7] T. L. Saaty, *What is relative measurement? The ratio scale phantom*, Mathematical and Computer Modelling Journal, vol. 17, number 4-5, pp. 1–12, 1993.

[8] W. Sonnenreich, J. Albanese, B. Stout, *Return On Security Investment (ROSI) A Practical Quantitative Model*, Journal of Research and Practice in Information Technology, vol. 38, number 1, 2006.

[9] M. Jeffrey, *Return on Investment Analysis for e-Business Projects*, Internet Encyclopedia. Hossein Bidgoli Editor, vol. 3, pp. 211–236, 2004.

[10] Lockstep Consulting, *A Guide for Government Agencies Calculating Return on Security Investment*, Technical Paper, 2004

[11] N. Kheir, N. Cuppens-Boulahia, F. Cuppens, H. Debar, *A Service Dependency Model for Cost-Sensitive Intrusion Response*, European Symposium on Research in Computer Security (ESORICS), pp. 626–642, 2010.

[12] G. Gonzalez-Granadillo, M. Belhaouane, H. Debar, G. Jacob, *RORI-based countermeasure selection using the OrBAC formalism*, International Journal of Information Security, Vol. 13(1), pp. 63–79, 2014.

[13] T. L. Norman, *Risk Analysis and Security Countermeasure Selection*, CRC Press, Taylor & Francis Group, 2010.

[14] 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*, In Security and Privicay in Communication Networks, vol. 164, pp. 538–555, 2015.

[15] L. Wang, T.a Islam, T. Long, A. Singhal, S. Jajodia *An attack graph-based probabilistic security metric*, Data and Applications Security XXII. Springer Berlin Heidelberg, 283–296, 2008.

[16] L. Yu, H. Man *Network vulnerability assessment using Bayesian networks*. International Society for Optics and Photonics, 2005.

[17] P. Xie, J. Li, X. Ou, P. Liu, R. Levy, *Using Bayesian networks for cyber security analysis*, International Conference on Dependable Systems and Networks, pp. 211–220, 2010.

[18] N. Kheir, H. Debar, N. Cuppens-Boulahia, F. Cuppens, J. Viinikka *Cost evaluation for intrusion response using dependency graphs*, International Conference on Network and Service Security, 2009.

[19] J. Marko, C. Thul, P. Martini, *Graph based metrics for intrusion response measures in computer networks*, 32nd IEEE Conference on Local Computer Networks, 2007.

[20] A. Barreto, P. Costa, E. Yano, *A Semantic Approach to Evaluate the Impact of Cyber Actions to the Physical Domain*, 7th International Conference on Semantic Technologies for Intelligence, pp. 64–71, 2012.

[21] A. Barreto, P. Costa, E. Yano, *Using a Semantic Approach to Cyber Impact Assessment*, 8th International Conference on Semantic Technologies for Intelligence, pp. 101–108, 2013.

[22] G. Jakobson, *Mission Cyber Security Situation Assessment using Impact Dependency Graphs*, In Fourteenth International Conference on Information Fusion, IEEE, 2011, pp. 1–8.

[23] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, 2014.

[24] M. Henrion, *Practical Issues in Constructing a Bayes' Belief Network*, In Third Conference on Uncertainty in Artificial Intelligence, 1987.

[25] J. Pearl, *Causality: Models, Reasoning and Inference*, 2nd Edition, Cambridge University Press, New York, NY, USA, 2009.

[26] A. Motzek, R. Möller, M. Lange, S. Dubus, *Probabilistic Mission Impact Assessment based on Widespread Local Events*, NATO IST-128 Workshop: Assessing Mission Impact of Cyberattacks, June 2015.

[27] M. Bastian, S. Heymann, M. Jacomy, *Gephi: An open source software for exploring and manipulating networks*, In International AAAI Conference on Weblogs and Social Media, 2009.

[28] L. D. Raedt, A. Kimmig, H. Toivonen, *ProbLog: A Probabilistic Prolog and Its Application in Link Discovery*, In IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007, pp. 2462–2467.