

Probabilistic Mission Defense and Assurance

Alexander Motzek and Ralf Möller

Universität zu Lübeck

Institute of Information Systems

Ratzeburger Allee 160, 23562 Lübeck

GERMANY

email: motzek@ifis.uni-luebeck.de, moeller@uni-luebeck.de

Abstract *Automatically generating adequate responses to ongoing or potential cyber threats and attacks is a pertinacious challenge and must have the aim to assure mission success, without sacrificing missions for security. To do so it must be understood how a threat may affect a mission, how a countermeasure diminishes potential threats, but also how a countermeasure might inadvertently impact the mission as well. Various approaches exist for all three subproblems and some for a partially combined solution. However, most suffer from one or more problems: (1) Approaches are holistic, delivering one acclaimed “optimal,” but intransparent solution. (2) Require unacquirable information that does not account for missing information, unforeseeable circumstances, or uncertainty. (3) Focus on cost optimization to mitigate direct affections without considering transitive impacts onto missions. In this paper we propose a probabilistic approach for cyber defense and assurance, decoupling mission impact assessments of threats and responses from a generation of those and from an optimal selection of those. We reduce mission impact assessments to commonly known mathematical problems, obtain directly validated and qualitative results, and greatly encompass missing information under uncertainty.*

1 Introduction

Automatically generating adequate responses to ongoing or potential cyber threats and attacks is a pertinacious challenge and must have the aim to assure accomplishment of missions, without sacrificing a mission for security. To do so it must be understood how a threat affects a mission, how a countermeasure diminishes potential threats, but also how a countermeasure might inadvertently impact a mission as well. For example, any potential compromise or procured failure of some node inside a network may lead to a causal chain of unforeseeable events and circumstances allowing an attacker to compromise further nodes until mission critical devices are affected as well. In order to mitigate these threats, the isolation or deactivation of all mission critical devices will definitely assure that no mission critical device will be adversarially compromised. Still, it is obvious that the mission will not succeed anymore.

Various approaches try to address these issues, but suffer from various problems. For example, various cost-minimizing approaches exist, but the cost of the abovementioned response is extremely low, as only some plugs need to be pulled. Moreover, various approaches do not encompass for the unknown: by trying to model exactly how an attacker will operate, e.g., in the form of attack-countermeasure-trees or attack graphs, any missing attack-step leads to failure of these approaches. We generally characterize frequent problems of existing approaches into three categories: **(1)** Approaches are holistic, i.e., try to solve a generation, evaluation and selection in a closed black-box approach delivering one acclaimed “optimal,” but intransparent solution. Informally and exaggeratedly said, holistic approaches may only provide information such as “*Response XYZ is best with metric 4589.32,*” which does not bear any meaning, requires a holistic reference set of all

other responses and deep training to vaguely understand it. (2) Require unacquirable information that does not account for missing information or uncertainty, e.g., require large and complex attack-countermeasure-trees explicitly identifying each and every possible attack and adequate countermeasures. (3) Focus on cost optimization to mitigate direct affections without considering transitive impacts onto a mission, i.e., do not consider unforeseen events and interactions between highly dependent nodes leading to mission failure.

In this paper, we take a fundamentally different perspective: We do not model an attacker, but model a mission from multiple perspectives as directly described by experts and automatically learned models. This is a paradigm shift, which allows us to consider all potential transitive and indirect effects from widespread events, i.e., positive and negative effects of threats and corresponding responses, potentially leading to unforeseeable chain of events. Moreover, we decouple the processes of generating responses, evaluating their effectiveness, and selecting an optimal response. Based on a well-defined mathematical problem, one obtains directly understandable assessments of responses and threats that do not require reference values or training to “judge” their optimality. Furthermore, we show how these assessments are used for an independent selection of adequate responses by the use of a multi-dimensional minimization problem, and we show how a mathematical graph-problem in the probabilistic model is used to generate adequate responses. The decoupling is highly beneficial, as the selection of an optimal response does *not* depend on the “correct” generation of response plans, i.e., obtained qualitative assessment deal as an independent and transparent validation of each response to assure mission success, and is suited for reporting along a command-chain.

This paper can be summarized as follows: By reducing mission defense onto a mathematical problem in probabilistic graphical models, one obtains qualitative, directly understandable mission impact assessments raising situational awareness, neither requiring reference values nor training to understand those assessments. A probabilistic graphical model is based on directly acquirable information and from automatic analyzes. By the use of probabilistic inference, transitive and indirect implications onto a mission are considered from adversarial and self-inflicted perspectives, incorporating unforeseeable chains of events and missing information.

The remainder of this paper is structured as follows: In Section 2 we introduce a probabilistic mission impact assessment and show in Section 3 how it is directly applicable for cyber defense assessments. We demonstrate our approach in Section 4 on real data in a real world scenario. We dedicate Section 5 to a discussion how a semi-optimal response is selectable by the use of a multi-dimensional minimization and how commonly known graph theory problems aid to generate novel response plans. We critically discuss our approach and related work in Section 6 and conclude in Section 7.

2 Probabilistic Mission Impact Assessment

A mission impact assessment (MIA) is used to assess potential impacts of occurring, widespread events onto a higher goal, e.g., a mission or onto a company. For example, a local impact of a distant node, e.g., a potential harm caused by a vulnerability, may lead to a causal chain of failures, disclosures and violations inside a network and will eventually impact critical resources involved directly in a mission. We say that a mission is impacted transitively by these events. To do so, locally caused impacts are “spread” throughout a network, even over nodes about which no direct information is available, incorporating unforeseeable chains of events.

Motzek et al. introduces in [1] an approach to probabilistic mission impact assessments based on a probabilistic graphical model, in which each parameter is directly understandable and validatable locally. By reducing MIA to a known mathematical problem—probabilistic inference—obtained results are immediately validated once parameters are validated and, by the use of probabilities, obtained assessments are directly interpretable without requiring reference results. For example, an obtained assessment states “*The probability that our mission will be impacted by known vulnerabilities inside our network is 37%*”—without knowledge

of the precise probabilistic graphical model, inference procedure or reference results, this statement is directly understandable and not negligible. We call these understandable and validated results *qualitative assessments*.

In the following we briefly introduce this probabilistic MIA based on [1], [2], and [3] (Motzek et al., further referred to as Motzek), and utilize it to obtain qualitative assessments of attacks and threats under the consideration of in-place countermeasures and their negatively invoked side-effects.

Motzek considers mission impact assessment from three different perspectives, involving different experts and expertise. Every expert defines a different dependency model, where every modeled entity represents a random variable and a dependency between two entities is represented by a local conditional probability of impact.

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

2.1 Mission Dependency Model (Business View)

Motzek extends a model by Jakobson [4] and model mission dependencies as shown in Figure 1 as a graph of *mission nodes*. For the scope of this work a business perspective is used, where a set of business processes are highly critical for the success of a company. An adequate analogy is directly evident for missions and their individual objectives. 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. Every node inside a dependency model represents a random variable, defined as follows.

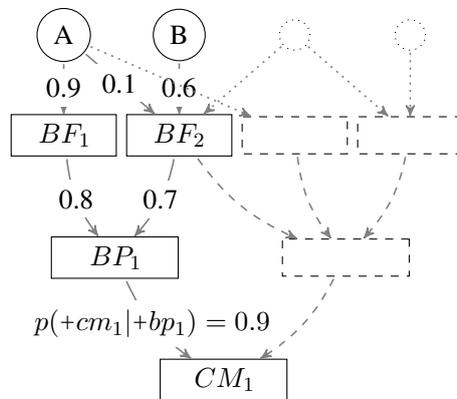


Figure 1: Mission Dependency Model. Values along edges denote individual conditional probability fragments.

Definition 1 (Random variables). *A random variable, denoted as capital X , is assignable to one of its possible values $x \in \text{dom}(X)$. Let $P(X = x)$ denote the probability of random variable X having x as a value. For our case we consider $\text{dom}(X) = \{true, false\}$ and we write $+x$ for the event $X = true$ and $-x$ for $X = false$.* ▲

The event $+x$ represents that node X is *impacted* and $-x$ that it is operating as intended, i.e., no impact is present. Dependencies are represented by local conditional probability distributions (CPDs) modeling probabilities of impact, given dependances are impacted. 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%. Motzek 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. For ease of parametrization of complete CPDs, every edge is associated with an individual local conditional probability of impact, e.g., for the above example $p(+bf_1|+a) = 0.9$. These probabilities are combined towards one distribution using so-called combination functions. Following [1], we employ a non-leaky noisy-or combination function in this work as described, e.g., by Cozman in [5]. Formally, Motzek defines in [2] a mission dependency model therefore as follows.

Definition 2 (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. 1) and are categorized according to their semantic as business-resources (\vec{BR}), -functions (\vec{BF}), -processes (\vec{BP}), and -company (\vec{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 $d \in \text{dom}(D), \forall D \in \vec{D}_V$. \blacktriangle*

With Definition 2, a mission dependency model represents a probabilistic graphical model, and, in particular, a Bayesian network, as, e.g., defined by Pearl and Russel in [6]. A key feature of Bayesian networks is the ability to locally interpret individual parameters, i.e., to locally interpret individual probabilities of CPDs. These properties are preserved in the presented probabilistic MIA as discussed in [2]. As all parameters are understandable locally, a mission dependency model is directly designable by an expert. Additionally, they can automatically be extracted from BPMN models. Further, mission dependency models are seen as persistent for a company, i.e., must only be designed once initially.

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, identifying a web-service as a business critical resource is reasonable, but it can not be expected that an underlying distributed computing cluster is identified in all extent providing the web-service. The following resource dependency model covers these dependencies.

2.2 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 probabilistic 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. [2] defines a resource dependency model formally as follows.

Definition 3 (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. 1) 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 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 of a resource dependency model is similar to the definition of a mission dependency model (Def. 2), and does represent a probabilistic graphical model as well, but does not introduce constraints of acyclicity, i.e., a resource dependency model can contain cyclic dependencies.

Motzek 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, [2] shows that an expert is able to validate a presented infrastructure dependency model for plausibility. Therefore, [2] presents an automatic learning approach for obtaining resource dependency models automatically from captured communication information, for which we present an example in Section 4. By incrementally relearning the model, the complete approach automatically adapts to changing environments.

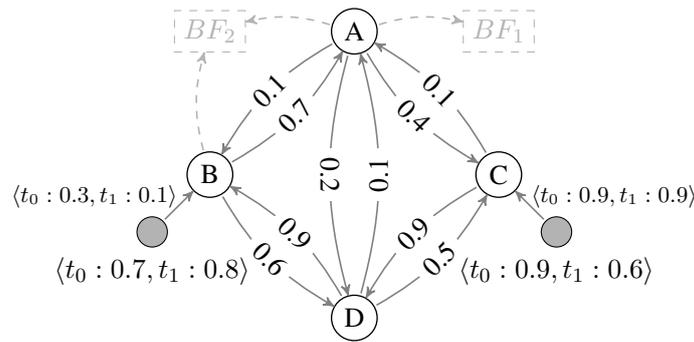


Figure 2: Resource Dependency Model. Dependencies between B, C would also be possible. Conditional probability fragments are marked along the edges. Grey nodes represent external shock events leading to local impacts. The time-varying conditional probability of local impact given an instantiated external shock event is given next to the edge and the time-varying shock event’s prior random probability is given below it. Connections to the mission dependency model are sketched in dashed gray.

2.3 Local Impacts (Security View)

Nodes of a resource dependency model might be 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 a shock event, e.g., the presence of a vulnerability or a direct shutdown of a node. Informally, an external shock event (SE) represents a source for an impact and is attached to a node in a resource dependency model, i.e., a SE threatens a node to be impacted. By representing SEs as random variables, one gains the ability to include uncertainty about the existence of SEs and uncertainty about whether a present threat leads to an impact on a node. Formally, Motzek defines external shock events in [2] as follows.

Definition 4 (External Shock Events). An external shock event SE is a random variable and 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 (−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 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. ▲

According to Definition 4, the presence of an external shock event can be known (observed) or can be unclear and is assessed probabilistically through its prior random distribution $P(SE)$. We denote the set of observed external shock events (known presence) as a set of instantiations \vec{se}_o of observed random variables $\vec{SE}_O \subseteq \vec{SE}$. This is highly beneficial for applications, where the actual presence of impact-sources is uncertain ($P(SE)$), and where evidence of existence and impacts is available, i.e., where SEs are observable ($+se \in \vec{se}_o$). To encompass varying effects over time, Motzek defines a temporal aspects of SEs as follows.

Definition 5 (Temporal Aspects). *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 $\langle t_0 : p_0, \dots, t_T : p_T \rangle$, 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. An expert initially designs these local consequences or utilizes flat assumptions, based on which specific external shock events are automatically initialized from obtained information, as discussed in Section 3.

2.4 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 [2] call the mission impact assessment, defined as follows.

Definition 6 (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 $MN \in M$ 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. ▲*

To obtain a solution to the MIA problem, one can see the probabilistic model as a probabilistic logic program, as elaborated in [1, 2, 3], where the MIA problem can be reduced onto a probabilistic inference problem. As probabilistic inference is generally known to be NP-hard, a approximate inference techniques is used, and [1] and [2] show and verify a linearly-scaling approximation procedure for obtaining solutions to MIA problems even in very large scaled domains in the range of seconds.

A probabilistic MIA $P(+mn|\vec{s}\vec{e}_o)$ directly originates from all defined dependency-models and represents an inference problem in a probabilistic graphical model. Therefore, [2] shows that if locally defined dependency-models are validated to be correct, an obtained impact assessment $P(+mn|\vec{s}\vec{e}_o)$ is validated, too.

3 Multi Dimensional Probabilistic Mission Defense and Assurance

Probabilistic mission impact assessment delivers context and bias free results as demonstrated by [3] and [2]. This means that no reference values are required for understanding an assessment. Moreover, the use of a probabilistic graphical model directly allows one to integrate uncertainty into models, e.g., uncertainty over the existence of vulnerabilities or imprecision of raised alerts. Furthermore, external shock events allow one to model impacts caused by adversaries, impacts by individual countermeasures, and effects of countermeasures on threats individually from local perspectives. Therefore the introduced probabilistic approach can directly be employed for mission defense as discussed in this section.

We differentiate between an *adversarial impact* (AI), i.e., an impact uncontrollable by one and caused by, e.g., IDS alerts, vulnerabilities or known threats, and between an *operational impact* (OI), i.e., a self-inflicted impact by a set of countermeasure on the mission. Both are differentiated in three temporal dimensions of a short-term, mid-term and long-term impact. Every source of these impacts is an external shock event. In the following two examples we discuss how a potential action by an attacker and our response are represented as external shock events. The key advantage of this approach is that both dimensions are modelable individually, i.e., not each and every combination of response and attack must be considered on every resource.

Example 1 (Adversarial Impact Shock Events). [2] represents every *vulnerability* information, e.g., automatically obtained from network scans, as an external shock event *VULN* affecting one or more nodes \vec{X} . Respectively, a prior random distribution $P(VULN)$ and local conditional probability fragments $p(+x|vuln)$ is defined. $P(VULN)$ directly allows one to model uncertainty about the actual existence of this vulnerability, e.g., through inaccurate scanners, as well as an uncertainty about the exploitability, e.g., if an exploit is present in common frameworks. The latter probability is likely to vary over time for which temporal aspects can be used, representing an increasing exploitability-probability over time. Likewise, $p(+x|vuln)$ represents the probability that, given this vulnerability is present and exploitable, it does harm to a node. These parameters are directly extractable from CVSS databases as explained in [2].

Similarly, a **raised alert** by an IDS is modeled in the same way as a shock event *ALER*. $P(ALER)$ represents an accuracy of the IDS, and $p(+x|aler)$ a probability that, if the alarm is true, a harm is caused, e.g., a very high probability that an adversarial impact is created, given a true alarm about a gained root privilege is authored. Naturally, for every alert category a different type of local external shock event may be modeled, e.g., one for port-scanning, one for dos-attacks and one for gained accesses. For modeling, temporal aspects can be used to gain awareness about persisting impacts for raised alarms, e.g., a gained root access will lead to a constantly high impact, given the alarm is true, i.e., $p(+x|root) = \langle t_0 : 1.0, t_1 : 1.0, t_2 : 1.0 \rangle$. Given a present dos-attack, impacts are high in a short- and mid-term, but likely to be low in a longterm, i.e., $p(+x|ddos) = \langle t_0 : 0.9, t_1 : 0.85, t_2 : 0 \rangle$. ♦

These external shock events create an adversarial impact (AI), varying over time, “spreading” throughout a network. For example, a gained root privilege might reveal passwords, useable to gain access on other nodes, or data is eavesdropped revealing information about dependent nodes. Our approach significantly different from existing approaches, which utilize, e.g., attack-graphs such as [7, 8]. Attack paths try to address the problem how exactly an attacker might compromise the network, i.e., they try to simulate an attacker. In contrary, we raise an amount of situational awareness that provokes an elimination of potential impact sources. In fact, Motzek shows in [2] that this approach is able to raise awareness for devices in a real world

scenario that were fully compromised by attackers through ways that were—in our opinion—completely unforeseeable by any classical “thinking like an attacker” or software-vulnerability focused analysis, e.g., were not foreseeable in classical attack graphs.

Naturally, attacks must be mitigated, for which we define a response in the form of a response plan formally as follows.

Definition 7 (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.* ▲

For example a response plan consists of multiple mitigation actions instructing the shutdown of some specific nodes. However, every mitigation action inside a response plans might cause an impact as well—an operational impact (OI), i.e., represent one or more external shock events aswell. On the other hand mitigation actions are able to mitigate or reduce adversarial impact probabilities. Probabilistic independencies of external shock events are used to model this interaction locally, as discussed in the following example.

Example 2 (Operational Impact Shock Events). *In the following we discuss three common cases of mitigation as given by [1]: Employing a **patch** on a node X may provoke collateral damage, i.e., represents a shock event $PATC$. During installation of a patch, there exists a (low) probability of immediate conflict, e.g., a flat assumption of 10% or a measure published by the software vendor. In a mean time, a patch might enforce a reboot of a network device. Finally, after one or more successful reboots and reconfigurations, the network device will fully resume its operational capability, and a vulnerability on a node (represented by shock event $vuln$) will be removed. One models a patching operation in three abstract timeslices and defines the local impact probabilities of this external shock event to be $p(+x|+pat) = \langle t_0 : 0.1, t_1 : 1.0, t_2 : 0.0 \rangle$. From a probabilistic perspective “removing a vulnerability” means that the node becomes independent of the external shock event: $t = 2 : P(X|+patc, vuln, \vec{Z}) = P(X|+patc, \vec{Z})$ or $t = 2 : X \perp vuln|+patc$. In the mid-term ($t = 1$), a vulnerability might or might not have been removed, which is represented by specifying $P(+x|+patc, +vuln, \vec{Z}) \leq P(+x|\neg patc, +vuln, \vec{Z})$ in the local CPD of the affected node X .*

*A restriction of a connection from node X to node Y , i.e., a new **firewall** rule, may invoke operational impact on Y , but prohibits “spreading” of adversarial impacts. From a technical perspective this operation forbids a transfer of data that might have been crucial for the operational capability of a node Y . As a connection between two devices resembles a dependency, one must further remove this dependency to prevent a double counting of impacts. To do so, [1] shows that one transforms a prohibited dependency to an observed external shock event $+se$ s.t. the local conditional failure probability $p(+y|+x)$ becomes a local impact probability $p(+y|+se)$. Temporal aspects can be used to model how long such a prohibition is intended to last. Multiple firewall rules can be used to completely **isolate** a node from the rest of a network, e.g., for inspection or repair.*

*Finally, a node can be **shutdown** as well, obviously creating operational impact by a shock event $SHUT$, but clearly avoiding all adversarial impacts $\vec{S}E_{AI}$ immediately, i.e., $p(+x|+shut) = \langle t_0 : 1.0, t_1 : 1.0, t_2 : 1.0 \rangle$ and $X \perp \vec{S}E_{AI}|+shut$. As a deactivated node is unable to communicate, a shutdown directly includes modeling an isolation.* ◆

This example shows how external shock events are used to model individual mitigation actions and their individual mitigation of adversarial impacts.

Note that neither interactions between all modeled AI and OI events nor all combinations of mitigation need be modeled. Only local operational impact effects of individual mitigation actions are modeled and some specific local effects. As an effect, these local impacts create time-profiles of a “fight” between impacts of different dimensions. A sketch of these profiles is visualized in Figure 4. Please note that these time-profiles are only examples of hypothetical effects of the modeled local impacts, i.e., “red” and “blue” are designed separately, automatically leading to the displayed effects.

All transitive and global effects of these local events are assessed probabilistically correctly through inference in the obtained probabilistic graphical model. To be precise, one obtains two (AI/OI) three-dimensional (i =short-, mid-, long-term) assessments for the mission MI of the mission dependency model M as $P_i^{AI}(+mi|\vec{s}e_o)$ and $P_i^{OI}(+mi|\vec{s}e_o)$, where $\vec{s}e_o$ is the set of the observed external shock events, e.g., exact knowledge of presence of vulnerabilities or the known execution of a mitigation action.

Note that a mission dependency model is only designed once, a resource dependency model is learned automatically and adapts to changing environments by a periodic re-learning. Local impacts of shock events are designed directly without a need to understand the complete approach or other dependency models as demonstrated in Example 1 and 2. Moreover, both examples show that these external shock events are automatically initialized based on present and automatically acquirable information.

An obtained assessment, e.g., a long-term probability of 90% that an adversary may cause an impact is directly understandable and does not require a “comparison” with other options—it is clearly not acceptable and must be mitigated. On the other hand, an operational impact is an impact as well and may lead to the same consequences as an adversarial impact does, e.g., a long-term probability of 90% that an executed mitigation may cause an impact on the mission is not acceptable as well. These properties allow the global mission impact assessment to be a direct assessment for complete defense strategies, decoupling generation and selection from this evaluation. This means that no holistic approach is taken, but proposals for responses are integrable from any source and a selection remains transparent and directly understandable for an expert. In the following section we give a short demonstration of obtained assessments, and show in Section 5 how a semi-optimal minimization is used to select a best compromise and how such responses are generatable automatically from the defined models.

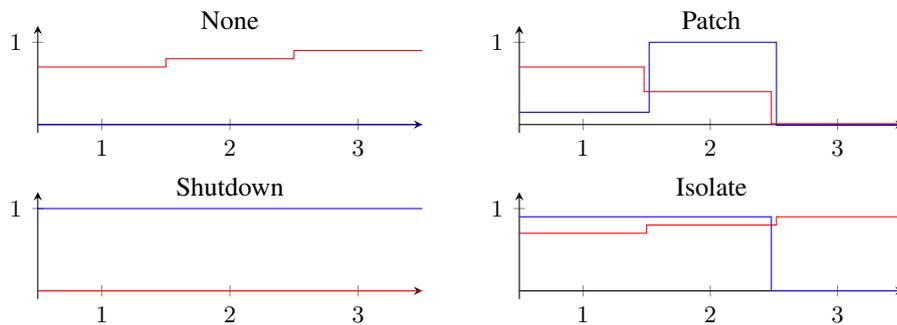


Figure 3: Sketch of local impact time profiles for adversarial impact (AI, red) and operational impact (OI, blue), while a vulnerability is present and some actions are executed (denoted as title). A vulnerability clearly poses a threat to a resources from an adversarial perspective (AI) raising over time (x-axis), transitively threatening other nodes. Patching said vulnerability might cause conflicts when installing, involves a period of uncertainty while rebooting provoking hardware failures and will eventually have removed the vulnerability (OI). Isolating for two timeslices provokes no immediate positive effect on the nodes, but prevents other dependent node from being adversarially impacted (compare Fig. 4). Respectively, the same holds for temporarily isolating and deactivating a node. There exists a tradeoff “game” between both impact dimensions over time.

4 Use Case Experiment

We evaluate and demonstrate the benefits of a probabilistically sound mission defense in a real world use case scenario involving real data. As part of the Panoptesec research project, we are able to apply our approach inside a backup-environment of ARETi, division of Acea SpA, Italy’s largest water services operator and one of the largest energy distribution companies in Italy [9]. ARETi is a division of Acea SpA in charge

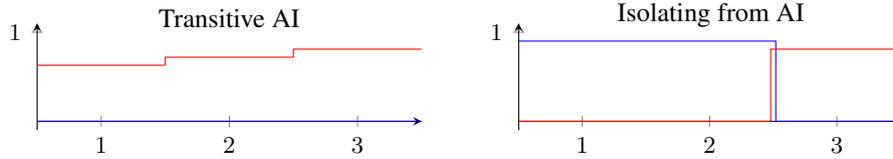


Figure 4: Sketch of transitive impact time profile of a node threatened by a vulnerability-impacted distant node (adversarial impact, red). Isolating this node from the impact source, e.g., for two timeslices, “removes” transitive AI during that period, but prohibits all potentially required dataflows as well leading to an operational impact (blue).

of distributing and controlling energy to the city and vicinity of Rome. We detailedly describe a process of obtaining required mission dependency models and resource dependency models in [2], which were in fact validated by IT-, security-, and business-experts to the company. A sketch of the obtained probabilistic graphical model is displayed in Figure 5.

As in our approach the generation, selection and assessment of responses to cyber threats and attacks is completely decoupled, we demonstrate our approach on intuitive, “hand-crafted” response plans, whose global effects on the mission are somewhat foreseeable. This allows one to verify our approach in that sense by the results of the experiment carried out in this section. In the upcoming sections we describe how adequate response plans are generated automatically and how an expert is assisted in selecting an appropriate one.

For the experiment, we consider the presence of two hypothetical known software vulnerabilities on two distant nodes¹ (Fig. 5, black nodes). The vulnerabilities are designed to lower their access complexity over time, i.e., a potential impact rises over time. Note that the affected nodes are not mission critical and at least one “hop” away from mission critical devices. Still, other nodes are highly dependent on them (thick edges), leading to an immediate spread of impacts. Without considering transitive effects there exists a high chance that other approaches, solely focusing on direct costs and direct mitigation, sacrifice the mission in favor of security. Moreover, no “vulnerability-path” exists from these affected nodes to other devices, i.e., purely software-vulnerability focused analysis would miss the potential impacts of these vulnerabilities.

We propose the following six response plans for demonstration: **(1)** no response is taken, i.e., only adversarial impact is present, **(2)** shutting down all critical devices, i.e., guarantees that no adversarial impact is posed on the mission, but clearly will sacrifice it, **(3)** a direct shutdown of affected nodes, i.e., will eliminate an AI onto a mission directly, but significantly hampers the operation of the network (OI), **(4)** all vulnerability-affected nodes are patched, i.e., in a long-term perspectives threats will be eliminated, **(5)** patching all vulnerabilities while isolating them from the network until a mid-term time interval, i.e., focuses on eliminating the threat in the long term and preventing soon malicious activities, and **(6)** a random choice of shutting down arbitrary devices.

As evident from Table 1, globally self-inflicted and adversarial-inflicted impacts onto the mission correspond to intuitive assumptions, which are further discussed in Example 3. Note that these assessments are based on a well-defined probabilistic graphical model and probabilistic inference, where all parameters have been validated. Therefore, these assessments are seen as validated as well. Moreover, these assessments stand for their own (qualitative assessments): Probabilities of impacts are not negligible and do not require reference results to judge their likelihood, and, depending on what is at stake, directly raise situational awareness for the criticality of the situation and appropriate response. For example, without knowing the second column (response description) of Table 1, a response plan can be chosen, without knowing reference values, without knowing all other possible response plan, without a detailed description of all parameters and

¹We emphasize that these two vulnerabilities are of completely hypothetical nature and are not present in the environment of Acea SpA or ARETi.

without knowing the complete attack scenario. Therefore, these assessment are highly suitable for reporting along a command-chain involving different experts from different expertises: no indepth-knowledge about attack paths, vulnerabilities or cyber incidents must be known, as the survival of the mission is a clear objective and must raise an appropriate situational awareness.

Table 1: Cyber defense assessments in ARETi for a set of responses, showing the impact probability by an adversary (AI) and the self-inflicted probability of operational impact (OI).

#	Response	$P^{AI}(+mi \vec{e}_o)$	$P^{OI}(+mi \vec{e}_o)$
1	No response	$\langle 0.36, 0.64, 0.75 \rangle$	$\langle 0, 0, 0 \rangle$
2	Shutdown critical devices	$\langle 0, 0, 0 \rangle$	$\langle 1, 1, 1 \rangle$
3	Shutdown directly affected	$\langle 0, 0, 0 \rangle$	$\langle 0.8, 0.8, 0.8 \rangle$
4	Patch all	$\langle 0.36, 0.32, 0 \rangle$	$\langle 0.08, 0.8, 0 \rangle$
5	Patch all while isolate	$\langle 0, 0, 0 \rangle$	$\langle 0.8, 0.8, 0 \rangle$
6	Random shutdown	$\langle 0.35, 0.63, 0.73 \rangle$	$\langle 0.94, 0.94, 0.94 \rangle$

This demonstration discusses mitigating impacts caused by vulnerabilities through patching or drastic actions of shutdowns. Notwithstanding, some vulnerabilities might not be patchable, or a patch cannot be applied to a system. Moreover, if a global adversarial impact origins from raised intrusion alerts, patching is not an option (which is directly incorporated in our approach—patching a non-affected node only leads to OI). In these particular situations a different approach must be taken, e.g., shutting down central nodes, prohibiting connections to “cut off” a path of an attacker, or manual inspections of then isolated nodes. Finding points of interest to execute these actions is an interesting problem, as millions of possible combinations exist. In the following section we discuss how these response plans are generatable automatically by using the probabilistic graphical model and a reduction to a graph theory problem.

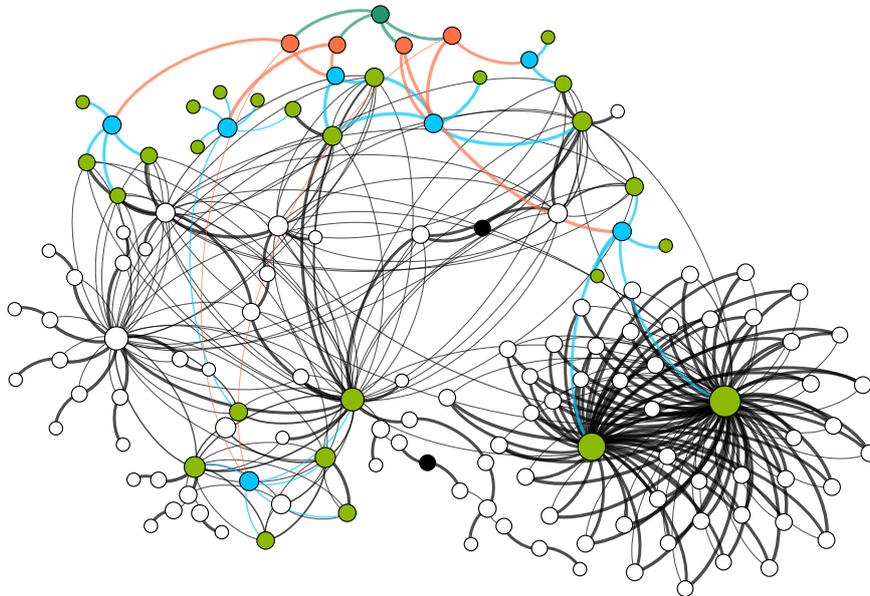


Figure 5: Resource dependency model extracted from roughly one month of traffic captures in Acea (represented in dark green) [2], where related critical devices are highlighted in green, business functions in blue, and business processes in orange. This model was validated and verified to be reasonable by the company’s IT experts. $n_N = 344, n_E = 754$.

5 Generation and Selection of Response Plans

Up to now we have discussed how one obtains understandable and transparent assessments of response plans from multi dimensional perspectives, considering the mitigation of attack surfaces, as well as potential negative side-effect of response plans themselves while focusing on the accomplishment of missions. The presented assessment is completely independent of how these response plans are generated and every assessment can be interpreted by its own.

Still, given multiple assessments, where one response plan is not clearly dominating, e.g., it is hard to decide between response plans 4 and 5 from Table 1. Moreover, if some nodes must be isolated, multiple options exist and it is hard to intuitively decide where specific mitigation actions should be placed to achieve a desired goal. In the upcoming two subsections we discuss how operators can be assisted in solving both problems based on a multi dimensional unweighted best compromise and graph theoretical problems.

5.1 Selection of Response Plans

Given multiple sets of response plans, where one response plan is not clearly dominating in all dimensions, a tradeoff must be found, i.e., what is the best compromise considering all dimensions. We describe in [3] an approach to select semi-optimal response plans based on an unweighted multi-dimensional optimization. By doing so one finds the best compromise in all dimensions, i.e., an operator must not come to a biased interpretation what is preferred, but is assisted in finding a best compromise. Notwithstanding, if an operator has a bias towards optimizing one dimension, e.g., the goal is to keep the long-term adversarial impact low at all costs (OI), one can exclude non-preferred dimensions from this optimization.

Both, AI and OI are impact assessments of proposed response plans. Still, due to their nature, an AI and an OI assessment follow 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 the surface for an attack. It is the novel advantage of the proposed approach of being able to combine both assessments while not being forced to define a preference-metric over them. We believe it is not practical to find a preference towards one dimension (e.g., to be solely biased towards AI_2). Further, defining a cost function (e.g., biasing by 30% towards OI_1 and 70% towards AI_2) is not practical as well.

We, therefore, define semi-optimal response plans in [3] as follows.

Definition 8 (Semi-optimal response plans). *Let $\vec{R}P^d$ be a vector of proposed response plans, associated with a linearly scaled impact assessment of dimension d . Let $\dot{R}P^d \subseteq \vec{R}P^d$ denote the set of optimal proposed response plans in terms of dimension d . Let $\hat{R}P^d$ denote the assessment of the theoretical optimal response plan and let $\check{R}P^d$ denote the assessment of the theoretical worst response plan in terms of dimension d . Then, let $\dot{R}P_\epsilon^d \subseteq \vec{R}P^d$ represent the set of semi-optimal response plans in terms of dimension d and easing factor $\epsilon \in [0, 1]$ representing the allowed deviation ϵ of the theoretical response plan range $|\hat{R}P^d - \check{R}P^d|$ from the evaluated optimal response plan $\dot{R}P^d$. Thus, $\dot{R}P_0^d = \dot{R}P^d$ and $\dot{R}P_1^d = \vec{R}P^d$. ▲*

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

Definition 9 (Smallest semi-optimum). *Let \vec{d} be the vector of all impact dimensions. Then, the smallest semi-optimal set of response plans $\dot{R}P$ is the set*

$$\dot{R}P = \min_{\epsilon} \left(\left\{ \bigcap_{d \in \vec{d}} \dot{R}P_\epsilon^d \right\} \neq \emptyset \right). \quad \blacktriangle \quad (1)$$

As both OI and AI assessments represent absolute metrics, $\check{R}P^{OI} = 1$ and $\hat{R}P^{OI} = 0$ (likewise for AI). This procedure assists an expert in deciding on semi-optimal response plan sets in every dimension, without enforcing a bias towards one explicit dimension. Please note, that by doing so a semi-optimal response plan with a compromise in some dimensions is chosen, in which some plans might dominate in certain dimensions, but which is directly evident. This is highly beneficial for applications where response plans must be chosen where no preference can be made between AI and OI, e.g., highly critical infrastructures, where any impact of any form must be avoided in any form. The following example demonstrates this approach on Table 1 delivering highly interesting insights to mission defense and situational awareness.

Example 3 (Defending is not always the best solution). *By Definition 9 on Table 1 the best option to execute response plan 1—no action. In response plan 1, “only” a compromise by 75% must be made in long-term AI from the optimal long-term AI; but OI is optimal in all dimensions. If doing nothing is not an option, the next semi-optimal set is response plans 1,3,4 and 5, where 5 dominates 3, leaving 1,4, and 5 as possible candidates. This example greatly shows the huge tradeoff that is often missed when considering a defense: mitigating the potential attack sources is as worst as doing nothing; only in a long-term perspective an advantage is obtained, by the potential sacrifice of a mission in, at least, a mid-term perspective. This is exactly what our approach does—raise awareness for the good and the bad sides of diminishing attack surfaces.* ◆

This example greatly demonstrates the benefits of our approach, i.e., assessments are directly understandable, consider transitive effects (no mission critical are devices threatened adversarially), and consider the negative effects of responses as well in a non-holistic approach.

Another dimension to consider is the workload to execute each response plan: If the monetary cost of executing mitigation actions is crucial, the minimization can directly include the sum of costs associated with each mitigation action in a response plan as another dimension.

As assessment and selection are decoupled from generation, both *do not require reference values from all possible response plans*. This is highly beneficial for applications, where response plans origin from multiple sources, such as automatic generations, expert intuitions, or mandated operational procedures. The decoupled evaluation delivers an independent validation of each proposal without requiring reference results. In the following section we show how an automatic generation can benefit from the obtained probabilistic graphical models.

5.2 Generation of Response Plans

In theory, there exists a hypothetical and extremely large, finite set of possible responses. These are built by considering all potential combinations of mitigations on each and every node. Naturally, it is completely intractable to evaluate all of them. Fortunately, as we do not follow a holistic approach, evaluating all is *not* required, which is a significant advantage compared to some related approaches. In our approach, as mentioned earlier, only a subset of promising responses must be evaluated. Moreover, the proposal of “promising” response plan can be based on greedy heuristics and is allowed to produce “false positive,” i.e., “bad” response plans, as such sub-optimal response plan will be assessed with a high OI and/or AI. Informally this means that we can generate as much response plans as we want, based on any heuristic, and the probabilistically correct assessment will remove the bad ones. In the following proposition we propose such a heuristic.

Proposition 1 (Response Plan Generation). *As identified in Example 1 and 2 patching can eliminate vulnerabilities in a long term. Let us name the set of external shock events that can be eliminated completely “curable.” Let \vec{V} be the set of nodes in a resource dependency model R affected by curable shock events. However, not all AI-causing shock events are directly “curable,” let \vec{A} be the set of nodes in R affected by adversarial shock*

events which are not curable. Then let $\vec{M}A_P$ be a set of mitigation actions instructing a patch, i.e., cure, of every node in \vec{V} . Let $\vec{M}A_{SV}$, $\vec{M}A_{SA}$ be sets of mitigation actions instructing a shutdown, i.e., deactivation, of every node \vec{V} , \vec{A} . Let $\vec{M}A_{IV}^i$, $\vec{M}A_{IA}^i$ be a set of mitigation actions isolating a every node \vec{A} , \vec{I} for up to i abstract timeslices.

In certain situations, nodes must be isolated, or a “path” must be cut in advance, e.g., by strategically placed firewall rules. Essentially, every edge in R is a candidate, leading to an infeasible amount of possibilities. Still, the best choice for an isolation is given directly from a resource dependency model R . For every edge $e \in R$ we define a minimal contribution probability p_{min} as follows: Let e origin from some node $X \in R$, then let p_{min} be the product of all local conditional probability fragments of the shortest path between X and the mission/company in a mission dependency model M . This follows the idea that if edge e is prohibited, at least a p_{min} probability of OI is caused on the mission. Then, let $\vec{M}A_{CV}$, $\vec{M}A_{CA}$ be sets of mitigation actions instructing firewall rules, i.e., connection prohibitions, which separate all critical nodes in M from directly affected nodes \vec{V} , \vec{A} with an expected low OI. The set of to-be-prohibited edges in $\vec{M}A_{CV}$, $\vec{M}A_{CA}$ is defined by the minimum cut of graph R partitioning all mission nodes $MN \in M$ from all nodes \vec{V} , \vec{A} based on p_{min} .

Every set $\vec{M}A$ and combination of multiple sets then represent possible and promising response plans. For example, $\vec{M}A_P$ and $\vec{M}A_{IV}^2$ are likely to represent one of the best response plans for proactive removal of known vulnerabilities, as evaluated in Table 1. Further, $\vec{M}A_{CA}$ is likely to represent one of the best responses to ongoing attacks. If the total number of individual mitigation actions becomes too large, e.g., by combinations of many $\vec{M}A$ sets, a randomly sampled subset is used, probably delivering valuable response plans. ▲

This proposition shows how response plans as sets of mitigation actions are proposed based on a probabilistic graphical model. Every proposed response plan is then evaluated as described above delivering qualitative results on which a decision can be grounded by relying on the validated parameters in the probabilistic graphical model. The nomenclature used in this proposition directly shows the broad applicability of the complete approach also to non-cyber-security related domains, such as healthcare or military applications. In the discussed example, one obtains 128 response plans by all complete combinations of mitigation action sets. As evaluated in [1], a single assessment is obtained in the range of milliseconds, allowing for near-realtime analysis in changing environments, where sets of external shock events quickly change.

6 Discussion and Related Work

Our approach is based on a probabilistic graphical model composed of three sub-models: a mission dependency model, a resource dependency model and a set of external shock events with associated local impacts. All three models are designable independently by different experts and incorporate a potential disagreement between different experts. For example, an identified web-server in a mission dependency model might be operationally unimportant, as an underlying database server or computational cluster is much more important. Due to an identified dependency of web-server on the computational cluster or database in the resource dependency model, both views are directly covered. Nevertheless, a resource dependency model must be learned, for which we propose an approach, but which may fail if exchanged information amounts do not correspond to actual information dependencies between devices, for example if enormous amounts of irrelevant data are transferred for no significant reason. In that particular situation a resource dependency model must be corrected manually. By periodically relearning the resource dependency model, it adapts to slowly changing environments and can be used in dynamic environments. If environments are changing rapidly, a differential

analysis is required, which is subject to future work. Mission dependency models are acquirable directly from experts, as we show in [2] for the ARETi use case or by automatic analyzes from BPMN models. Given multiple mission dependency models, one congruent model must be obtained which poses a problem of semantic normalization and merging, which we deeply discuss in [10].

We admit that the approach and definitions of locally created impacts of external shock events may seem simplistic and may simplify the complete problem. However, as we show in Section 4 and [2] both approaches deliver great results in real world use cases. Moreover, the obtained well-defined model and mathematically grounded approach has significant advantages as we discussed throughout this paper. As identified in the beginning, a large variety of related work exists, but often suffer from problems of not considering mission impact relations, negative side-effects of responses, requiring the finely granular modeling of exact attack paths and only provide intransparent solutions from holistic approaches. In the following we discuss various approaches, their benefits and drawbacks.

Viduto et al. present in [11] a promising approach to ponder between risks posed by vulnerabilities against a degree of financial investments for responses using a multi-objective tabu search. However, they only consider direct impacts on nodes and assume that in the absence of a vulnerability and in the absence of an exploit no harm may be caused at all. The use of relative metrics does not allow an operator to understand metrics directly and can only support a holistic approach without a considering mission impact relations and negative side effects of the proposed responses.

Considering indirect effects and missing information is partially incorporated by Foo et al. in [12] by the use of “spread” channels in a network. They notably identify that only considering local reactions to raised alarms is not sufficient and try to maintain subfunctions of provided services. However, a novel propagation algorithm employed for “spreading” impacts is not mathematically grounded and can solely provide responses in a holistic approach and only consider implications of proposed responses to a limited extent.

Considering the negative side effects of response plans is considered by Toth et al. in [13] in a similar approach to ours by modeling dependencies of a network. However, no approaches for automatically obtaining these dependencies is provided and the use of novel propagation algorithms does not provide a directly understandable assessment.

Various approaches, e.g., [14, 15, 16] consider attack-countermeasure-trees which provide a mixture of attack graphs and associated defense actions and show to be scalable in large domains. These works extend approaches based on attack graphs, such as [17], [18]. Based on attack graphs, [19, 20] consider a novel and interesting problem of defending cloud-based virtual networks, where virtual instances might compromise underlying systems. However, all of these approaches are based on an explicit and detailed representation of all attackers’ paths through a network up to their goal. We believe that it is highly impractical to model these attack trees manually involving all possible ways an attacker can take, and it is an infeasible process to constantly adjust these models in dynamic and changing environments. Moreover, an automatic generation of such attacker models, such as [7, 8, 21] suffer from a significant problem, as they are fully dependent on knowing all vulnerabilities. In effect, these approaches assume that in the absence of a vulnerability absolutely no harm may be caused, even if all surrounding nodes are highly compromisable. Exaggeratedly said, these approaches assure that no threat to a mission exists, even if all nodes in a network are directly compromisable by vulnerabilities except the mission critical business resources. Furthermore, to the best of our knowledge, none of these approaches consider the negative side effects of the proposed responses.

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 a response. [22] presents an interesting and well-formalized approach for situations where too few budget exists to fully implement all mitigation actions to known attack surfaces. [23] and [24] 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 [24, 16, 25, 11]. 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. While the local assessment of these costs is easily to perform, assessing their global negative side-effect costs is a highly pertinacious challenge and impossible to define by an expert manually for every device. In our approach the complete assessment of global negative side-effects is directly incorporated based on well-defined mathematical principles.

The degree of employed uncertainty in our approach, e.g., considering vulnerabilities as first-step sources of impacts and inferring all possible consequences is highly beneficial, as missing information is directly incorporated, and a situational awareness is still raised if partial “links” in vulnerability-chains are missing. Moreover, by the use of a well-defined probabilistic graphical model all parameters are directly understandable, easy to parametrize and can be validated. Therefore, returned results are immediately validated if the parameters are validated, as the assessment is based on a formally correct probabilistic inference problem. Often, approaches, e.g., [14] or [25], involve extremely large amounts of parameters which are not directly understandable and hard to parametrize, requiring a deeply trained expert in an employed framework to do so. [19] considers a probabilistic approach as well to determine the likelihoods of explicit attack paths. However, presented probability theory in [19] is not sound and voids fundamental principles of probabilistic inference in multiply connected graphs, as far as we can tell.

7 Conclusion

In this paper we present a probabilistic approach to mission defense, focusing on the assurance of missions, without sacrificing them for the sake of security. We consider that locally created impacts, e.g., attacks may lead to unforeseeable events and spreads of impact throughout a network through ways unforeseeable in advance by experts. Moreover, we consider that any response to an ongoing or proactively defended attack surface may impact a mission to the same extend as the attack itself. By reducing the problem onto a well-defined inference problem in probabilistic graphical models, all parameters are understandable by themselves, without a need to overlook a big picture and are validatable by individual experts from different expertise. Based on validated parameters, obtained results, i.e., assessments of adversarial and operational impacts, are validated as well and directly understandable without requiring reference values for comparison. With these novel advantages, we decouple assessments from a generation and selection of response plans. This means that every assessment is acceptable or deniable by itself, no matter where or who proposed the response by whatever approach. In contrast in a holistic approach, an operator is forced to accept “the relatively best value” found by some response plan generator, whose assessment does not bare any meaning and only provides the knowledge that it was somehow superior to others, but the worseness and betterness remains unknown.

Future work is dedicated to extend the approach to a completely dynamic model allowing for real-time and forensic analyses in rapidly changing environments based on novel advantages by [26] on artificial intelligence and probabilistic graphical models as discussed in [2].

Acknowledgments

This work was partly supported by the Seventh Framework Programme (FP7) of the European Commission as part of the PANOPTESec integrated research project (GA 610416).

References

- [1] A. Motzek, R. Möller, M. Lange, and S. Dubus, “Probabilistic Mission Impact Assessment based on Widespread Local Events,” in *NATO IST-128 Workshop: Assessing Mission Impact of Cyberattacks, NATO IST-128 Workshop, Istanbul, Turkey, June 15-17, 2015*, pp. 16–22, 2015.
- [2] A. Motzek and R. Möller, “Context- and Bias-Free Probabilistic Mission Impact Assessment,” tech. rep., Universität zu Lübeck, Institut für Informationssysteme, 2016. under review.
- [3] G. G. Granadillo, A. Motzek, J. Garcia-Alfaro, and H. Debar, “Selection of Mitigation Actions Based on Financial and Operational Impact Assessments,” in *ARES 2016: 11th International Conference on Availability, Reliability and Security, Salzburg, Austria, August 31 - September 2, 2016*, pp. in–press, 2016.
- [4] G. Jakobson, “Mission Cyber Security Situation Assessment using Impact Dependency Graphs,” in *FUSION 2011: 14th International Conference on Information Fusion, Chicago, Illinois, USA, July 5-8, 2011*, pp. 1–8, 2011.
- [5] F. G. Cozman, “Axiomatizing Noisy-OR,” in *ECAI 2004: 16th European Conference on Artificial Intelligence, including PAIS’2004: Prestigious Applicants of Intelligent Systems, Valencia, Spain, August 22-27, 2004*, pp. 979–980, 2004.
- [6] J. Pearl and S. Russell, “Bayesian Networks,” in *Handbook of Brain Theory and Neural Networks* (M. A. Arbib, ed.), pp. 157–160, MIT Press, 2003.
- [7] S. Jha, O. Sheyner, and J. Wing, “Two Formal Analyses of Attack Graphs,” in *CSFW 2002: 15th IEEE Workshop on Computer Security Foundations, Cape Breton, Nova Scotia, Canada, June 24-26, 2002*, pp. 49–63, IEEE, 2002.
- [8] X. Ou, S. Govindavajhala, and A. W. Appel, “MulVAL: A Logic-based Network Security Analyzer,” in *14th USENIX Security Symposium, Baltimore, Maryland, USA, July 31 - August 5, 2005*, 2005.
- [9] Acea SpA, “The Acea Group,” 2016.
- [10] A. Motzek, C. Geick, and R. Möller, “Semantic Normalization and Merging of Business Dependency Models,” in *CBI 2016: 18th IEEE Conference on Business Informatics, Paris, France, August 29 - September 1, 2016*, pp. in–press, 2016.
- [11] V. Viduto, C. Maple, W. Huang, and D. López-Pérez, “A novel risk assessment and optimisation model for a multi-objective network security countermeasure selection problem,” *Decision Support Systems*, vol. 53, no. 3, pp. 599–610, 2012.
- [12] B. Foo, Y. Wu, Y. Mao, S. Bagchi, and E. H. Spafford, “ADEPTS: Adaptive Intrusion Response Using Attack Graphs in an E-Commerce Environment,” in *DSN 2005: International Conference on Dependable Systems and Networks, Yokohama, Japan, 28 June - 1 July, 2005*, pp. 508–517, 2005.
- [13] T. Toth and C. Krügel, “Evaluating the Impact of Automated Intrusion Response Mechanisms,” in *ACSAC 2002: 18th Annual Computer Security Applications Conference, Las Vegas, NV, USA, December 9-13, 2002*, pp. 301–310, 2002.

- [14] A. Roy, D. S. Kim, and K. S. Trivedi, “Attack Countermeasure Trees (ACT): Towards Unifying the Constructs of Attack and Defense Trees,” *Security and Communication Networks*, vol. 5, no. 8, pp. 929–943, 2012.
- [15] A. Roy, D. S. Kim, and K. S. Trivedi, “Cyber Security Analysis Using Attack Countermeasure Trees,” in *CSIIRW 2010: 6th Cyber Security and Information Intelligence Research Workshop, Oak Ridge, TN, USA, April 21-23, 2010*, p. 28, 2010.
- [16] A. Roy, D. S. Kim, and K. S. Trivedi, “Scalable Optimal Countermeasure Selection Using Implicit Enumeration on Attack Countermeasure Trees,” in *DSN 2012: IEEE/IFIP International Conference on Dependable Systems and Networks, Boston, MA, USA, June 25-28, 2012*, pp. 1–12, 2012.
- [17] T. Somestad, M. Ekstedt, and P. Johnson, “Cyber Security Risks Assessment with Bayesian Defense Graphs and Architectural Models,” in *HICSS 2009: 42nd Hawaii International Conference on Systems Science, Waikoloa, Big Island, HI, USA, January 5-8, 2009*, pp. 1–10, 2009.
- [18] S. Ossenbuhl, J. Steinberger, and H. Baier, “Towards Automated Incident Handling: How to Select an Appropriate Response against a Network-Based Attack?,” in *IMF 2015: 9th International Conference on IT Security Incident Management & IT Forensics, Magdeburg, Germany, May 18-20, 2015*, pp. 51–67, 2015.
- [19] C. Chung, P. Khatkar, T. Xing, J. Lee, and D. Huang, “NICE: Network Intrusion Detection and Countermeasure Selection in Virtual Network Systems,” *IEEE Trans. Dependable Sec. Comput.*, vol. 10, no. 4, pp. 198–211, 2013.
- [20] J. B. Hong, C. Chung, D. Huang, and D. S. Kim, “Scalable Network Intrusion Detection and Countermeasure Selection in Virtual Network Systems,” in *ICA3PP: International Workshops and Symposia on Algorithms and Architectures for Parallel Processing, Zhangjiajie, China, November 18-20, 2015*, pp. 582–592, 2015.
- [21] S. Jajodia and S. Noel, “Topological Vulnerability Analysis,” in *Cyber Situational Awareness - Issues and Research*, pp. 139–154, 2010.
- [22] R. Dewri, N. Poolsappasit, I. Ray, and D. Whitley, “Optimal Security Hardening Using Multi-Objective Optimization on Attack Tree Models of Networks,” in *CCS 2007: ACM Conference on Computer and Communications Security, Alexandria, Virginia, USA, October 28-31, 2007*, pp. 204–213, 2007.
- [23] S. A. Zonouz, H. Khurana, W. H. Sanders, and T. M. Yardley, “RRE: A Game-Theoretic Intrusion Response and Recovery Engine,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 25, no. 2, pp. 395–406, 2014.
- [24] S. Bistarelli, M. Dall’Aglia, and P. Peretti, “Strategic Games on Defense Trees,” in *FAST 2006: 4th International Workshop on Formal Aspects in Security and Trust, Hamilton, Ontario, Canada, August 26-27, 2006*, pp. 1–15, 2006.
- [25] N. Stakhanova, C. Strasburg, S. Basu, and J. S. Wong, “Towards Cost-Sensitive Assessment of Intrusion Response Selection,” *Journal of Computer Security*, vol. 20, no. 2-3, pp. 169–198, 2012.
- [26] A. Motzek and R. Möller, “Indirect Causes in Dynamic Bayesian Networks Revisited,” in *IJCAI 2015: 24th International Joint Conference on Artificial Intelligence, Buenos Aires, Argentina, July 25-31, 2015*, pp. 703–709, 2015.