A multi-objective decision support framework for simulation-based security control selection

Elmar Kiesling, Christine Strauß
Secure Business Austria
Vienna, Austria
Email: {ekiesling,cstrauss}@sba-research.org

Christian Stummer
Bielefeld University
Bielefeld, Germany
Email: christian.stummer@uni-bielefeld.de

***Abstract*—In this paper, we report on our ongoing research on simulation-based information security risk assessment and multi-objective optimization of investment in security controls. We outline a methodological framework that accounts for characteristics of the organization, its information infrastructure, assets to be protected, the particular threat sources it faces, and the decision-makers' risk preferences. This framework comprises (i) ontological modeling of security knowledge, (ii) dynamic attack graph generation techniques, (iii) probabilistic simulation of attacks by goal-driven threat agents, (iv) metaheuristic identification of efficient portfolios of information security controls, and (v) interactive decision support. These components facilitate novel techniques to infer possible routes of attacks and generate attack graphs based on attackers' motivation, objectives, capabilities, and available modes of entry and to use this inferred knowledge to simulate attacks on an organization's modeled infrastructure. The method supports decision makers evaluating potential security control investments in striking a balance between monetary and non-monetary criteria regarding risks, costs, and benefits. We are currently in the process of developing a prototypical implementation of the framework that will be used to evaluate the approach through application case studies.**

***Keywords*-Security and protection; systems analysis and design; computational modeling; simulation; decision support systems; human factors**

I. Introduction

In the face of the complex nature of information security (IS) problems and the large array of available physical, technical, operational, and organizational controls that aim to improve it (e.g., virus scanners, firewalls, intrusion detection and prevention systems, two-factor authentication, access control systems, improvements in network configuration, encryption, patch management, security awareness training, account/password management policies, etc.), IS managers struggle to identify the most appropriate means to counteract the threats their organizations face [1]. Due to a lack of standard methodologies for selecting the "best" overall combination of controls to implement [2], IS investment decisions tend to be driven by fear [3] and immediate needs [4]. This typically results in an ad-hoc approach that neglects positive and negative synergies, leads to an inefficient allocation of scarce resources, and may be the reason why IS has traditionally been seen as an expense that brings little tangible benefits [5].

A possible way to tackle the information security investment decision making problem in a more proactive way is to rely on standards and best practice guidelines that provide prescriptive guidance based on approaches that have been proven to work in other organizations. Examples include ISO/IEC 27000-series of standards [6], the German IT base protection catalogues [7], and the French EBIOS standard [8]. Compliance with such standards represents a significant improvement over reactive approaches to IS. However, these standards can only offer general, high-level recommendations without grounding the reasoning in the organizations particular IT infrastructure and environment [9]. Furthermore, best practices cannot support organization-specific threat scenarios [10]. Finally, their adoption may compel organizations to follow a practice irrespective of its applicability to the actual risks the organization faces [11] rather than precisely targeting security investments to tackle specific risks based on business needs and strategic objectives. While following best practices ideally provides a sufficient security level, security managers usually also have to ensure that no excessive investments are made [12].

The main challenge IS managers face in this context is to strike an appropriate balance between risk exposure and the opportunity to mitigate risk through investments in security. This balance must be defined within the business's risk environment, which includes the characteristics not only of the firm, but also those of attackers [13]. Making decisions based on an abstract threat posed by a generic population of "hackers" fails to take into account that there are substantial differences in attackers motivations, goals, skills, and points of access. Furthermore, such an approach neglects the fact that a large proportion of attacks come from insiders [14], [15]. This issue is critical because different attackers will respond differently to the same countermeasures [16]. Taking into consideration the heterogeneity in attacker characteristics and behavior is therefore a prerequisite for an efficient response.

In order to find an appropriate balance for security investments, it is also typically necessary to consider multiple

978-0-7695-4775-6/12 $26.00 © 2012 IEEE
DOI 10.1109/ARES.2012.70

conflicting objectives, some of them monetary (such as minimizing the costs of IS controls or the losses caused by security incidents), others non-monetary (such as minimizing the negative impact of intrusions on consumer confidence). Decision makers must trade off these objectives to arrive at an optimal strategy with respect to their preferences. As a system's overall security (e.g., in terms of confidentiality, integrity, availability) depends on the combined effects of all implemented controls – which is generally not cumulative – a comprehensive evaluation rather than an assessment of individual investment options in isolation is necessary. This approach accounts for complex interactions between the individual controls' effects.

Finally, the problem needs to be cast in terms that both security managers responsible for implementing an IS strategy (e.g., CIOs, CISOs) and senior managers responsible for allocating resources for IS investments can relate to. Existing models and approaches for supporting IS investment decisions are typically based on either an engineering or a managerial perspective. Our aim is to bridge the gap between these perspectives and to develop an integrated decision support approach for the evaluation of IS investment portfolios. To this end, we follow an interdisciplinary research approach drawing on a variety of disciplines, including Management, Operations Research, Economics, and Computer Science.

In this paper, we propose a decision analytic approach and outline our ongoing research by introducing a framework that comprises (i) ontological security knowledge, (ii) dynamic attack tree generation techniques, (iii) stochastic simulation optimization, and (iv) interactive decision support. This approach facilitates an analytical process that systematically evaluates portfolios (i.e. bundles) of IS controls along multiple dimensions in terms of risks, costs and benefits, and takes into consideration the organization's characteristics, information infrastructure and assets, as well as properties of threat sources (including goal-oriented attackers) and candidate controls. The proposed approach not only provides security managers with assessments of individual controls and bundles thereof, but also is capable of optimizing these bundles with respect to multiple criteria concerning risk, costs, and benefits. In addition, it allows decision makers to interactively explore the space of efficient solutions when deciding which portfolio of controls to implement.

The remainder of this paper is structured as follows: Section II outlines our research approach, Section III presents the conceptual model that serves as a basis for the framework, Section IV introduces the methodological framework, and Section V concludes this paper.

II. Research approach

We set out four guiding principles for our research. First, we based our research on a probabilistic view that conceives IS investment decisions as activities that are inherently characterized by uncertainty and variability. While uncertainty results from the analyst's limited understanding (e.g., regarding existence and severity of vulnerabilities, effectiveness of controls, etc.), variability is the result of truly random processes (e.g., the time it takes to crack a password using a brute-force attack) [17]. Our probabilistic approach furthermore recognizes that "security" ought to be understood as a continuum marked by gradual differences, whereas "total security" is a purely hypothetical situation [18]. In line with this view, the aim of the proposed research is not to identify a single best optimal investment strategy, but to allow the decision maker to choose from a set of efficient security investment portfolios according to his/her preferences regarding risks, costs, and benefits.

Second, our research follows a bottom-up simulation approach. This approach recognizes that security or the lack thereof is the result of complex causal interactions. In contrast, existing risk analysis methods are typically based on a top-down "divide and conquer" approach and analyze the security of an information system by disassembling it into its parts, analyzing these parts statistically and assembling the understanding gained into an understanding of the whole system [19]. This reductionist approach is well suited for mechanistic reliability assessments, but fails to identify interactive combinations of failures. As soon as active elements such as human attackers are involved, failures are not uncorrelated and depend in subtle ways on system state. To capture these complex causal interactions and the problem's human nature, we approach the problem by explicitly modeling and simulating attacks on the system with and without IS controls in place. This approach evaluates how the choice of an IS investment portfolio affects the overall system, taking interacting vulnerabilities into account. By aiming at the selection of a diversified and balanced portfolio of IS controls (cf. [20], [21]), this comprehensive method enables rational information security decisions.

Third, we base our research on the assertion that guarding against threats requires an understanding of their source, which in most cases is ultimately human in nature. Schneier [22] states that the term "security" is essentially meaningless if the question "secure from whom?" is not addressed. Nevertheless, most IS risk management approaches neglect the fact that human threat sources vary widely with respect to their motives, capabilities, resources, trust status, risk preference, and objectives for an attack, although these properties determine their attack campaign and, ultimately, the risk they pose to an information system. To properly assess risks and select controls accordingly, IS managers must explicate their assumptions regarding adversaries and their intentions. Whereas methods for the modeling and analysis of occasional stochastic risks such as the risk of hardware component failure have become quite mature, modeling goal-driven attacks originating from human threat agents is a challenge because their occurrence does not follow any reasonable statistical pattern. Moreover, these human threats tend to be rather victim-specific [23] and they may also respond differently to the same countermeasures depending on their motivation and abilities. For these reasons, the selection of IS controls should be based on an explicit attack source model. Hence, we model human attackers as rational agents that maximize expected utility, weighting perceived monetary and psychic costs and benefits when choosing from available actions. These costs and benefits will vary greatly depending on the attackers motivation [24]. We consider an economic approach toward the modeling of attackers behavior most promising and base our attack agent model upon the solid theoretical foundation of Becker's economic theory of action in his classic work on the economics of crime [25]. This approach assumes that attackers respond to incentives imposed by the implementation of controls that either increase the effort required for or lower the probability of a successful attack (e.g., physical locks, firewalls, or improved password policies), increase the risk of being detected (e.g., intrusion detection systems), or reduce the expected benefits from a successful attack (e.g., reduce a targets value to the attacker).

Fourth, we base our research approach on the view that IS risk modeling should necessarily be a decision-driven activity, because "some policy will be inevitably chosen, even if that policy is to do nothing" [11]. While IS resource allocation decisions (if made systematically at all) are most commonly based on a single indicator criterion such as the return on investment, this approach conceals inherent trade-offs between security goals and other objectives (e.g., security vs. productivity, cost, etc.), as well as trade-offs among security goals, which often conflict fundamentally [26], [27]. Moreover, relevant metrics may be difficult to convert into monetary terms and it is often impracticable to aggregate measures from different dimensions by expressing all benefits and costs of controls as financial values [5]. We therefore follow a multi-criteria decision analysis (MCDA) approach, which can cope with situations in which multiple, conflicting objectives need to be considered simultaneously (for an overview cf. [28]).

III. Conceptual model

We base our research upon a conceptual model that revolves around the five key concepts *asset*, *threat source*, *threat*, *vulnerability*, and *control*. In doing so, we will rely on established concepts from sources such as the security and dependability taxonomy by Avizienis et al. [29] or the concepts used in the formalization of security knowledge by Fenz and Ekelhart [30]. Figure 1 illustrates all top-level concepts and their relations.

An *asset* represents something of value to an organization that may require a level of protection [31]. We use this broad definition that includes all types of tangible objects (e.g., building, room, server, network component, etc.) and intangible elements (e.g., software, data, etc.). A *threat* potentially exploits an existing vulnerability through a physical, technical, or administrative weakness and affects security properties (e.g., confidentiality, availability, integrity) of *assets* upon realization.

In order for *threats* to materialize, a *threat source* is required. This source may be either a *stochastic threat source* (e.g., hardware failure, non-targeted malware, etc.) or a *threat agent* (any type of internal or external attacker) in our conceptual model.

A *threat agent* is characterized by its *capabilities* (e.g., technical skills, resources, etc.), *access* to assets (e.g., an internal attacker will typically have more effective modes of access), motivation (e.g., curiosity, peer recognition, control, revenge, financial gain, etc.) and particular *objectives* (e.g., access a confidential information asset, inflict physical damage to an asset, disrupt availability of an asset, etc.), which affect an *asset*'s *security property* (e.g., confidentiality, availability, integrity).

By realizing a *threat* through a successful *attack action* (e.g., exploit a software vulnerability, obtain password through social engineering, etc.), an *attack agent* can obtain access to additional *assets* (e.g., root access to firewall), which may in turn enable it to execute additional *attack actions* (e.g., on systems behind that firewall). A *threat* can give rise to follow-up threats (e.g., break-in gives rise to unauthorized access or asset damage) and *assets* may require other *assets* (e.g., software is dependent upon the server it runs on). These relations reflect that attackers can achieve objectives by chaining *attack actions* and attacking *assets* indirectly.

The final key concept in our framework is *control*, which can be implemented by an asset to mitigate *vulnerabilities* through preventive, corrective, or detective measures. A control may affect other controls and their combination may yield positive or negative synergies (e.g., reduced costs if two controls from the same vendor are implemented, increased or reduced efficiency if two controls are deployed in combination, etc.), which is reflected by the relation "affects".

IV. Methodological framework

Figure 2 provides an overview of the proposed method and the related components that support each step of the iterative decision process. Step 1 consists in modeling the system and its environment as well as in identifying threats and potential controls by creating instances of concepts in an ontological knowledge base. This step can be supported by an ontology populated with an extensive set of initial information and may involve questionnaires, on-site interviews, document reviews, and automated scanning tools.

The resulting operational model of the system and its threat environment is then used in Step 2 to simulate attacks

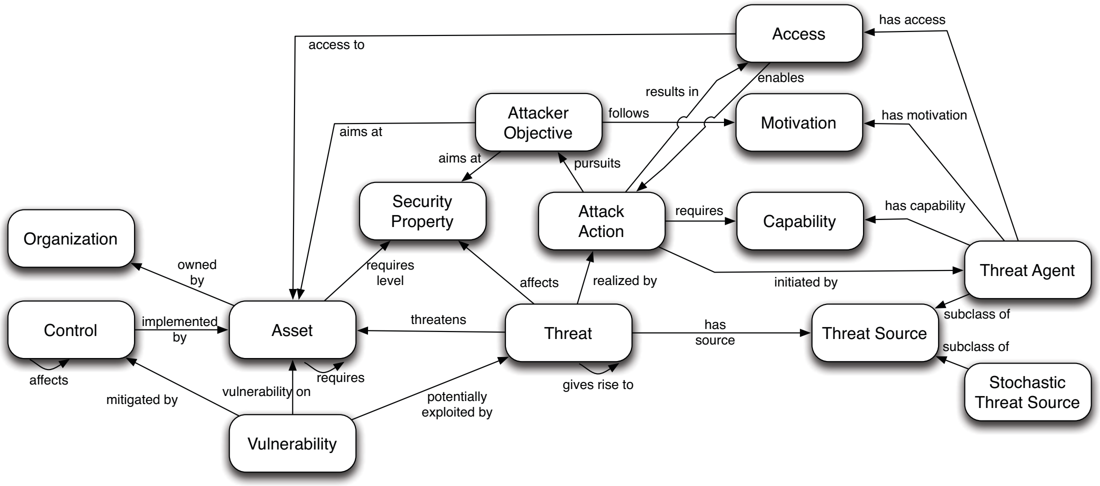

Figure 1. Conceptual model

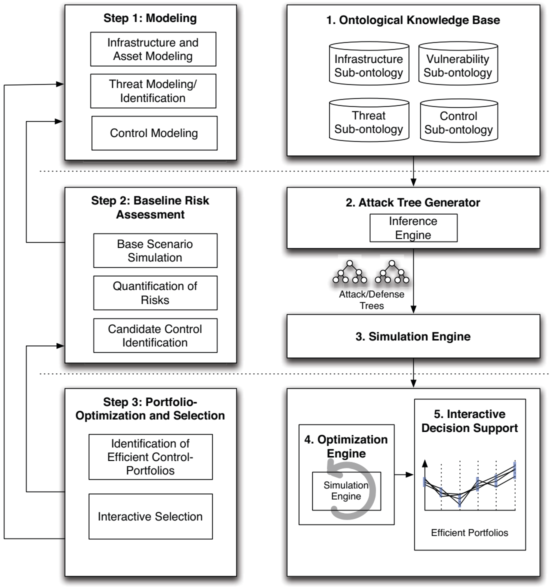
Step 1: Modeling
Infrastructure and Asset Modeling
Threat Modeling/ Identification
Control Modeling
Step 2: Baseline Risk Assessment
Base Scenario Simulation
Quantification of Risks
Candidate Control Identification
Step 3: Portfolio-Optimization and Selection
Identification of Efficient Control-Portfolios
Interactive Selection
1. Ontological Knowledge Base
Infrastructure Sub-ontology
Vulnerability Sub-ontology
Threat Sub-ontology
Control Sub-ontology
2. Attack Tree Generator
Inference Engine
Attack/Defense Trees
3. Simulation Engine
4. Optimization Engine
Simulation Engine
5. Interactive Decision Support
Efficient Portfolios

Figure 2. Methodological framework

and obtain a baseline risk assessment of the current state that serves as a basis for the nomination of candidate controls.

Step 3 identifies efficient bundles of controls and provides interactive decision support for selecting a preferred portfolio to implement. The remainder of this section describes each of the components in the proposed methodological framework.

1) Ontological knowledge base: To formalize the conceptual model presented in Section III and capture the knowledge required for simulating attacks in a structured, reusable, and machine-processable manner, we are developing an ontological knowledge base.

In the context of IS, ontologies were initially proposed as a potential solution to the problem of vaguely defined terminology [32]. To this end, lightweight ontologies have been used for organizing terminology in a hierarchy of concepts and relations. In our framework, by contrast, we opt for heavyweight ontologies that are additionally enriched with axioms to fix the semantic interpretation of concepts and relations [33]. These ontologies not only provide terminological rigor, but also lend themselves to the use of inference engines to obtain new knowledge based on existing facts and rules, as well as the use of reasoner engines to maintain consistency [34]. In the context of our research, this facilitates the processing of knowledge contained in the ontology, e.g., to infer attacker-specific attack trees.

The ontology's top-level concepts will correspond to the elements of the conceptual model (cf. Section III). We will refine these top-level concepts in the following sub-ontologies: (i) an infrastructure sub-ontology will store tangible and intangible assets, (ii) using a threat sub-ontology, we will capture abstract types of threats and concrete threats to specific assets, as well as abstract and specific (technical or human-behavioral) vulnerabilities and their relations, (iii) a threat source sub-ontology will store information on threat sources (stochastic threat sources and threat agents) and their properties, and (iv) IS controls will be organized in a controls sub-ontology that defines and relates them to the infrastructure concepts that implement them, the infrastructure concepts they can be applied to, and the abstract or specific threats that they mitigate.

An interesting aspect of the proposed approach is that parts of the knowledge base, such as the threat and control sub-ontologies, may be stored in a shared repository and reused by multiple organizations. Whereas domain experts are required for the definition and maintenance of the centralized part of the ontology, organizations could then use this existing knowledge set about threats, vulnerabilities and controls and automatically relate it to their own systems (modeled as instances of infrastructure sub-ontology concepts). Our research will therefore result in a ready-to-use but extendable knowledge set to be filled with concrete data about the particular threat model and infrastructure, which will be modeled as assets stored as instances of the already modeled asset classes.

2) Attack tree generator: Todays serious attacks are complex, multi-stage scenarios that coordinate the effects of multiple single-point attack actions to reach goals otherwise not obtainable [35]. For this reason, proper means to capture the causal interdependencies between attack actions are necessary. To this end, attack trees provide a convenient formalism to systematically categorize the different ways in which a system can be attacked [36]. The term was introduced in the field of IS (cf. [37]) as a concept derived from fault trees [38], which have been used for fault assessment of critical systems for decades. The basic idea of the approach is simple – an attack scenario is represented in a tree-based structure in which the root node represents the attackers goal and paths from leaf nodes to the root represent different ways of achieving this goal. Nodes that lie on the path between leaves and the root node are sub-goals. Children of a node are refinements of this (sub-)goal and can be conjunctive (aggregation) or disjunctive (choice).

In the context of our research, extended attack trees provide a representational formalism for the causal structures relevant for particular attack objectives, taking adversaries entry points and capabilities into account. The ontological knowledge base described in the previous section stores vulnerabilities and threats which are related through preconditions and postconditions to represent complex causal interactions. This makes it possible to define potential attack actions without knowledge of how they will be used. As new attack actions are added to the knowledge base, they may combine in ways not originally realized. The attack tree generation component can then use an inference engine to harness this knowledge and create attack trees for particular attackers taking their respective goals (e.g., obtain full access to consumer database) and individual characteristics (e.g., insider/outsider, available points of entry, capabilities, etc.) into account. The chaining of attack actions through preconditions and postconditions is not a trivial task, since such attack graph construction techniques are generally plagued with a combinatorial search space [39]. Developing adequate techniques for tackling the complexity involved, e.g., through abstraction or selective limitation of the depth of enumeration, will therefore be an important area of research.

Individual nodes in the obtained attack trees are annotated with probabilities for successful execution of the respective attack action; these probabilities, which may depend upon individual attacker capabilities, are stored in the knowledge base. When controls are applied, they may affect a single or multiple nodes in the tree and either eliminate them, lower the probability of a successful attack, increase the probability of detecting an attempted attack or, in some cases, decrease the attackers success probability for one type of attack action while increasing it for another. The construction of individual attack trees for various types of attackers and attack objectives finally results in an "attack forest" used as input for the simulation component.

3) Simulation engine: The simulation engine constitutes the core component of the proposed architecture. It performs a probabilistic evaluation of IS safeguard portfolios by explicitly simulating attacks on the modeled infrastructure and ascertaining the consequences with regard to criteria defined by the decision maker. The simulation relies on information on infrastructure elements, threats, and vulnerabilities stored in the ontological knowledge base and uses the attack trees generated by the attack tree generator component described in the previous paragraph.

The adequate representation and simulation of attackers' behavior in line with their characteristics and objectives is an important aspect. In this context, we draw upon and intend to extend the existing body of literature on attacker behavior modeling (e.g., [23], [40], [41], [42]).

The simulation engine samples each candidate portfolio by performing a large number of replications. Sampling a portfolio finally yields distributions of outcomes measured along multiple dimensions; these dimensions may include cumulative monetary damage caused by successful attack actions, number of distinct attack routes which led to threat realizations, impact on confidentiality, availability, and integrity, etc. Mont et al. [43] suggest assurance, agility, security, compliance, productivity and empowerment as strategic outcomes of interest. Neubauer et al. [10] use effectiveness, maintainability, reliability, running costs, and initial costs as criteria for selecting ISO 27001 controls. In the context of our research, it will also be interesting to trade off the effectiveness of the portfolio against different types of attackers.

Depending on the type of objectives defined by the IS manager, each outcome distribution may need to be transformed into scalar objective values, e.g., by using average or median values, user-specified quantiles, or worst case realization values. Alternatively, stochastic dominance criteria may be used for comparing portfolios in some cases. Based on the performance indicators obtained, the optimization component can automatically determine the set of efficient portfolios of IS controls or an approximation thereof.

4) Optimization engine: Rather than constructing portfolios manually and evaluating them individually until a satisfying solution is found, an optimization approach allows decision makers to select portfolios implicitly by stating their preferences regarding different types of risks, costs, and benefits. Instead of dealing with the problem "locally" and focusing on questions such as whether to implement specific controls, the multi-objective optimization approach aims for overall efficiency and enables decision makers to deal with the problem in terms of trade-offs between high-level objectives.

Due to the combinatorial nature of the portfolio selection problem (i.e., one binary decision for each candidate control and possible deployment location) and the multitude of relevant criteria, multiobjective combinatorial optimization algorithms have to be used to identify efficient portfolios of controls. Our use of the term "efficient" corresponds to the standard definition of Pareto-optimality, i.e., the property that no other portfolio exists that is better in at least one criterion and at least equally good in all other criteria.

Mathematically, we are dealing with a problem that is NP-hard [28] and requires a considerable amount of computational effort for each portfolio evaluation performed. This implies that even though the approach is highly parallelizable and potentially almost ideally scalable, the problem still will not be amenable to algorithms enumerating and sampling all feasible portfolios for non-trivial problem instances. To compute approximations of the set of Pareto-optimal solutions, one must therefore resort to (meta-) heuristic approaches (for a survey cf. [44]). Techniques that have turned out to be promising in this problem context include multiobjective genetic algorithms (e.g., NSGA-II [45], SPEA2 [46]), multi-criteria variants of the Nested Partitioning method [47], and Pareto Ant Colony Optimization [48].

5) Interactive decision support: Given the potentially large number of efficient solutions, selecting a final portfolio of controls that best fits the IS manager's preferences is not a trivial task. The literature [49] distinguishes three categories of methodological approaches to support the decision maker in this final step and limit the considerable cognitive burden involved: (i) filtering methods, (ii) clustering methods and (iii) search-based procedures.

The first class of approaches reduces the set of efficient portfolios by discarding the most redundant points while retaining solutions that are most dissimilar. An application of such an approach in the context of IS control is [50], which is based on a k-tree data structure [51].

Next, clustering methods can be applied to form groups of similar portfolios. Once a manageable number of clusters have been identified, a representative portfolio of IS controls from each cluster is presented to the decision maker, who can then choose the most preferred of these solutions and, in a second step, examine a neighborhood around this point.

Finally, search-based procedures start from an efficient portfolio and enable "movement" in the solution space toward more attractive alternatives until no "better" solutions can be found. We develop a search-based interactive decision support approach similar to that described in [52]. Following this approach, the IS manager iteratively establishes aspiration levels for objectives or modifies upper and lower bounds through a graphical interface and, thus, narrows the set of candidate portfolios. During this procedure, the user gains a better understanding of the solutions available, the structure of the problem, and the trade-offs between criteria before making a final choice. The merit of the approach lies not only in the numbers produced, but also in the insights that security managers gain during each refinement step of the assessment.

V. Conclusions

The need for comprehensive quantitative models of information security is frequently highlighted in the literature [53], [54], [55], [56]. Although there has been some progress in this direction, quantitative methods that supports decision makers in systematically optimizing the selection of information security investments based on a comprehensive assessment of their joint effectiveness in protecting a modeled system are still lacking.

Our research endeavor presented in this paper takes up this issue. It develops a framework built upon the foundation of heavyweight ontologies which serve as a knowledge base used to generate attack trees for individual attackers. One important advantage of this approach is that possible attack actions and patterns can be defined without knowledge of how they will be combined and used. We develop both an ontological framework for representing rich security knowledge and novel attack tree computation techniques to harness that knowledge through automated reasoning.

The simulation component uses the inferred knowledge to simulate attacks on the modeled infrastructure. It properly captures uncertainty, variability and complex interactions, and explicitly accounts for the organization's particular infrastructure and the particular threat sources it faces. Although opportunistic and random attacks (such as malware) can also be modeled using the proposed method, the salient contribution of our research lies in the modeling of attackers as goal-oriented agents. This characterization is consistent with theoretical and empirical studies [57] and highly relevant, given the increasing frequency of goal-oriented multi-stage attacks. Another key merit of this approach is the proper differentiation between outsider and insider threats.

The use of meta-heuristic procedures to identify efficient portfolios of IS controls through stochastic simulation optimization constitutes a novel application of these techniques. Given the combinatorial difficulty involved in manually constructing non-dominated subsets of candidate security controls taking multiple, partly conflicting, objectives into account, we expect that the use of optimization algorithms will provide significant benefits in this area.

To sum up, the bottom-line impact of our research on a multi-objective decision support framework for simulation-based security control selection comes from (i) the construction of an ontological knowledge base that captures the complex causal interactions between threats and vulnerabilities to enable automated inference of possible routes of attack, (ii) models and tools for simulating sophisticated multi-stage attacks carried out by goal-oriented threat agents, (iii) the development of new simulation optimization approaches, and (iv) an improved valuation of information security investment opportunities through the innovative application of multi-criteria decision support methods. We are currently in the process of implementing the individual components of the outlined framework and will evaluate the proposed approach through application case studies.

Acknowledgment

Funding from the Austrian Science Fund (FWF) under project number P 23122-N23 is gratefully acknowledged.

References

[1] W. H. Baker and L. Wallace, "Is information security under control? Investigating quality in information security management," *IEEE Security & Privacy*, vol. 5, no. 1, pp. 36–44, 2007.

[2] S. Bistarelli, F. Fioravanti, and P. Peretti, "Defense trees for economic evaluation of security investments," in *Proceedings of the First International Conference on Availability, Reliability and Security (ARES 2006)*, Vienna, Austria, 2006, pp. 416–423.

[3] T. Neubauer, M. Klemen, and S. Biffl, "Secure business process management: a roadmap," in *Proceedings of the First International Conference on Availability, Reliability and Security (ARES 2006)*, Vienna, Austria, 2006, pp. 457–464.

[4] M. T. Dlamini, J. H. P. Eloff, and M. M. Eloff, "Information security: the moving target," *Computers & Security*, vol. 28, no. 3-4, pp. 189–198, 2009.

[5] T. Neubauer, C. Stummer, and E. Weippl, "Workshop-based multiobjective security safeguard selection," in *Proceedings of the First International Conference on Availability, Reliability and Security (ARES 2006)*, Vienna, Austria, 2006, pp. 366–373.

[6] ISO/IEC, "ISO/IEC 27000-series: Information Security Management Systems standards," International Standardization Organization/International Electrotechnical Commission, 2011. [Online]. Available: http://www.iso.org/

[7] BSI, "IT-Grundschutz catalogues," Federal Office for Information Security (BSI), 2008. [Online]. Available: https://www.bsi.bund.de/gshb

[8] DCSSI, "Expression des besoins et identification des objectifs de securité (EBIOS)," General Secretariat of National Defense Central Information Systems Security Division (DCSSI), 2004.

[9] A. Baldwin, M. C. Mont, and S. Shiu, "Using modeling and simulation for policy decision support in identity management," in *IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY 2009)*, London, UK, 2009, pp. 17–24.

[10] T. Neubauer, A. Ekelhart, and S. Fenz, "Interactive selection of ISO 27001 controls under multiple objectives," in *Proceedings of the 23rd International Information Security Conference*, Milano, Italy, 2008, pp. 477–492.

[11] K. J. S. Hoo, "How much is enough: a risk management approach to computer security," PhD Diss., School of Engineering, Stanford University, Stanford, CA, 2000.

[12] A. Jürgenson and J. Willemson, "Processing multi-parameter attack trees with estimated parameter values," in *Advances in Information and Computer Security, LNCS 4752*. Berlin: Springer, 2007, pp. 308–319.

[13] H. Cavusoglu, S. Raghunathan, and W. Yue, "Decision-theoretic and game-theoretic approaches to IT security investment," *Journal of Management Information Systems*, vol. 25, no. 2, pp. 281–304, 2008.

[14] J. D'Arcy, A. Hovav, and D. Galletta, "User awareness of security countermeasures and its impact on information systems misuse: a deterrence approach," *Information Systems Research*, vol. 20, no. 1, pp. 79–98, 2009.

[15] E. Shaw, K. G. Ruby, and J. M. Post, "The insider threat to information systems: the psychology of the dangerous insider," *Security Awareness Bulletin*, vol. 2, 1998.

[16] M. Rounds and N. Pendgraft, "Diversity in network attacker motivation: a literature review," in *Proceedings of the IEEE International Conference on Computational Science and Engineering*, Los Alamitos, CA, USA, 2009, pp. 319–323.

[17] F. O. Hoffman and S. H. Jana, "Propagation of uncertainty in risk assessments: the need to distinguish between uncertainty due to lack of knowledge and uncertainty due to variability," *Risk Analysis*, vol. 14, no. 5, pp. 707–712, 1994.

[18] K. Bauknecht and C. Strauss, "A framework to support decisions on appropriate security measures," in *Safety of Computer Control Systems*, H. Frey, Ed. Oxford: Pergamon, 1992, pp. 253–258.

[19] D. White, "Application of systems thinking to risk management: a review of the literature," *Management Decision*, vol. 33, no. 10, pp. 35–45, 1995.

[20] C. Strauss and K. Bauknecht, "Portfolio techniques to support risk management and security," in *Computer Security and Information Integrity*, S. Rautakivi and J. Saari, Eds. Amsterdam: North-Holland, 1991, pp. 9–28.

[21] T. Finne, "Information security implemented in: the theory on stock market efficiency, Markowitz's portfolio theory and porter's value chain," *Computers & Security*, vol. 16, no. 6, pp. 469–479, 1997.

[22] B. Schneier, *Secrets & Lies: Digital Security in a Networked World*. New York, NY: Wiley, 2000.

[23] A. Buldas, P. Laud, J. Priisalu, M. Saarepera, and J. Willemson, "Rational choice of security measures via multi-parameter attack trees," in *Proceedings of the first International Workshop on Critical Information Infrastructures Security (CRITIS 06), LNCS 4347*, 2006, pp. 235–248.

[24] N. Kshetri, "The simple economics of cybercrimes," *IEEE Security and Privacy*, vol. 4, no. 1, pp. 33–39, 2006.

[25] G. S. Becker, "Crime and punishment: an economic approach," *Journal of Political Economy*, vol. 76, no. 2, pp. 169–217, 1968.

[26] L. A. Gordon and M. P. Loeb, "The economics of information security investment," *ACM Transactions on Information and System Security*, vol. 5, no. 4, pp. 438–457, 2002.

[27] G. Elahi and E. Yu, "A goal oriented approach for modeling and analyzing security trade-offs," in *Proceedings of the 26th International Conference on Conceptual Modeling, LNCS 4801*, Auckland, New Zealand, 2007, pp. 375–390.

[28] M. Ehrgott and X. Gandibleux, "A survey and annotated bibliography of multiobjective combinatorial optimization," *OR Spectrum*, vol. 22, no. 4, pp. 425–460, 2000.

[29] A. Avizienis, J. C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, no. 1, pp. 11–33, 2004.

[30] S. Fenz and A. Ekelhart, "Formalizing information security knowledge," in *Proceedings of the 4th International Symposium on Information, Computer, and Communications Security (ASIACCS '09)*, Sydney, Australia, 2009.

[31] D. Stepanovka, S. Parkin, and A. van Moorsel, "A knowledge base for justified information security decision-making," in *Proceedings of the 4th International Conference on Software and Data Technologies (ICSOFT)*. Sofia, Bulgaria: INSTICC Press, 2009.

[32] M. Donner, "Toward a security ontology," *IEEE Security & Privacy*, vol. 1, no. 3, pp. 6–7, 2003.

[33] F. Fürst and F. Trichet, "Heavyweight ontology engineering," in *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops, LNCS 4277*. Montpellier, France: Springer, 2006, pp. 38–39.

[34] G. Goluch, A. Ekelhart, S. Fenz, S. Jakoubi, S. Tjoa, and T. Mück, "Integration of an ontological information security concept in risk aware business process management," in *Proceedings of the 41st Hawaii International Conference on System Sciences (HICSS)*, Hawaii, USA, 2008, pp. 377–385.

[35] S. J. Templeton and K. Levitt, "A requires/provides model for computer attacks," in *Proceedings of the 2000 Workshop on New Security Paradigms*, 2000, pp. 31–38.

[36] S. Mauw and M. Oostdijk, "Foundations of attack trees," in *Revised Selected Papers of the 8th Information Security and Cryptology (ICISC 2005), LNCS 3935*. Seoul, Korea: Springer, 2006, pp. 186–198.

[37] B. Schneier, "Attack trees," *Dr. Dobb's Journal*, vol. 24, no. 12, pp. 21–29, 1999.

[38] W. Vesely, F. Goldberg, H. Roberts, and D. Haasl, "Fault tree handbook (NUREG-0492)," Division of Systems and Reliability Research, Office of Nuclear Regulatory Research, U.S. Nuclear Regulatory Commision, Washington, DC, U.S., Tech. Rep., 1980.

[39] J. Dawkins, "Heuristics for scalable compound exposure analysis: a foundation for a comprehensive security risk assessment," PhD Diss., University of Tulsa, 2005.

[40] E. Jonson and T. Olovsson, "A quantitative model of the security intrusion process based on attacker behavior," *IEEE Transactions on Software Engineering*, vol. 23, no. 4, pp. 235–246, 1997.

[41] P. Liu, W. Zang, and M. Yu, "Incentive-based modeling and inference of attacker intent, objectives, and strategies," *ACM Transactions on Information and System Security*, vol. 8, no. 1, pp. 78–118, 2005.

[42] K. Sallhammar, B. E. Helvik, and S. J. Knapskog, "Incorporating attacker behavior in stochastic models of security," in *Proceedings of the International Conference on Security and Management (SAM '05)*, Las Vegas (NV), USA, 2005, pp. 55–68.

[43] M. C. Mont, Y. Beres, D. Pym, and S. Shiu, "Economics of identity and access management: a case study on enterprise business services," HP Laboratories, Technical Report HPL-2010-11, 2010.

[44] M. Ehrgott and X. Gandibleux, "Approximative solution methods for multiobjective combinatorial optimization," *TOP*, vol. 12, no. 1, pp. 1–63, 2004.

[45] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast elitist multi-objective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2000.

[46] E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: improving the strength pareto evolutionary algorithm," in *Evolutionary Methods for Design, Optimisation and Control*, K. Giannakoglou, D. Tsahalis, K. Papailiou, and T. Fogarty, Eds. International Center for Numerical Methods in Engineering, 2002.

[47] L. H. Lee, E. P. Chew, and S. Teng, "Integration of statistical selection with search mechanism for solving multi-objective simulation-optimization problems," in *Proceedings of the 38th Winter Simulation Conference*, Monterey, California, USA, 2006.

[48] K. F. Doerner, W. J. Gutjahr, R. F. Hartl, C. Strauss, and C. Stummer, "Pareto ant colony optimization with ILP preprocessing in multiobjective project portfolio selection," *European Journal of Operational Research*, vol. 171, no. 3, pp. 830–841, 2006.

[49] S. B. Graves and J. L. Ringuest, *Models & Methods for Project Selection*. Berlin: Springer, 2002.

[50] C. Strauss and C. Stummer, "Multiobjective decision support in IT-risk management," *International Journal of Information Technology & Decision Making*, vol. 1, no. 2, pp. 251–268, 2002.

[51] M. Sun and R. E. Steuer, "InterQuad: an interactive quad tree based procedure for solving the discrete alternative multiple criteria problem," *European Journal of Operational Research*, vol. 89, no. 3, pp. 462–472, 1996.

[52] C. Stummer, E. Kiesling, and W. J. Gutjahr, “A multicriteria decision support system for competence-driven project portfolio selection,” *International Journal of Information Technology & Decision Making*, vol. 8, no. 2, pp. 379–401, 2009.

[53] S. C. Patel, J. H. Graham, and P. A. S. Ralston, “Quantitatively assessing the vulnerability of critical information systems: a new method for evaluating security enhancements,” *International Journal of Information Management*, vol. 28, no. 6, pp. 483–491, 2008.

[54] J. J. C. H. Ryan and D. J. Ryan, “Expected benefits of information security investments,” *Computers & Security*, vol. 25, no. 8, pp. 579–588, 2006.

[55] O. Alhazmi, Y. Malaiya, and I. Ray, “Security vulnerabilities in software systems: a quantitative perspective,” in *Proceedings of the 19th Annual IFIP WG 11.3 Working Conference on Data and Applications Security, LNCS 3654*, S. Jajodia and D. Wijesekera, Eds., Storrs, CT, USA, 2005, pp. 281–294.

[56] D. Geer, K. S. Hoo, and A. Jaquith, “Information security: why the future belongs to the quants,” *IEEE Security & Privacy*, vol. 1, no. 4, pp. 24–32, 2003.

[57] M. Cremonini and D. Nizovtsev, “Understanding and influencing attackers’ decisions: implications for security investment strategies,” in *Proceedings of the Fifth Workshop on Economics of Information Security (WEIS 2006)*, Cambridge, UK, 2006.