Simulation-based optimization of IT security controls: Initial experiences with meta-heuristic solution procedures

Elmar Kiesling, Andreas Ekelhart, Bernhard Grill[*], Christine Strauß[†], Christian Stummer[‡]

February 11, 2013

Extended Abstract

Today's most severe information security threats are no longer random, opportunistic attacks such as viruses, but targeted attacks that combine multiple attack vectors to achieve particular goals. Managers responsible for ensuring the security of complex information systems are therefore confronted with growing numbers of increasingly sophisticated attacks. In line with this development, the focus of information security research has broadened from individual technical vulnerabilities toward threats that emerge from the dynamic interaction of multiple attack vectors and their deliberate exploitation by sophisticated attackers. To cope with such threats, a comprehensive approach to secure information systems by selecting a set of appropriate security controls (i.e., countermeasures), while trading off multiple cost and benefit objectives, is necessary. For a related early approach, cf. [1].

In our ongoing research project MOSES3 (*"Multi-objective decision support for efficient security safeguard selection"*), we tackle this challenge and develop an optimization framework based on the idea of simulating deliberate attacks on a modeled systems with various sets of controls (i.e., "control portfolios") in place. To this end, we combine rich conceptual modeling of security knowledge, explicit behavioral modeling, attack graph generation techniques, and discrete event simulation to evaluate individual "control portfolios" (i.e., a modeled system configurations with a particular set of security controls in place). A comprehensive conceptual overview of the framework is provided in [2].

To simulate attacks, it is necessary to model the required knowledge in a well-structured and reusable manner. To this end, we use a knowledge base c consisting of (i) an attack and control model and (ii) a system model. The attack and control model captures complex causal relationships and may be shared among multiple organizations. The system model, by contrast, defines the system to be analyzed and contains the set of constituent tangible and intangible assets, including hardware components, networks, data, employees, policies etc.

We are currently experimenting with two alternative approaches that allow us to use reasoners and/or query languages. The first approach is to capture the knowledge in OWL 2 ontologies [3] modeled in Protégé [4], reusing concepts from the information security ontology introduced in [5]. The main advantages of this approach are that it allows us to transform attack patterns into SPARQL queries on the system model and that the modeled knowledge may easily be shared by a community of users and that existing reasoner engines can be used. The second approach is to formalize the knowledge in Prolog, using the SWI-Prolog [6] implementation. This approach offers substantial performance advantages and is hence more suitable for optimization purposes.

To specify an optimization problem, we define attack scenarios that consist of (i) an attacker model and (ii) a definition of the attacker's objectives. While attackers are typically classified based on a natural language description in the literature (e.g., external, internal, government, secret services etc. [7]), we take advantage of our formal models to allow for more specific attacker definitions that include particular objectives (e.g., access a particular data asset) as well as attacker attributes such as available equipment, skills, knowledge, and points of entry. Based on this attack scenario definition and the abstract attack patterns modeled in the knowledge base, it is possible to derive sequences of attack actions that enable the attacker to achieve the particular objective while accounting for individual attacker characteristics. In this context, we build upon and extend the existing literature on graph-based attack modeling [8, 9, 10].

The attack simulation is implemented in Java and based on an explicit behavioral model to capture the dynamic interaction between the simulated attacker and the system. It requires efficient means for maintaining and processing a timeline of events. To this end, we use the scheduling mechanisms provided by MASON [11], a fast discrete-event simulation core written in Java. During simulation runs, several types of events (for attacker actions, detective control actions etc.) are used to model the dynamics. Random distributions are used prevalently

[*]Andreas Ekelhart, Bernhard Grill, and Elmar Kiesling are with Secure Business Austria (email: akelehart|ekiesling|bgrill@sba-research.org)

[†]Christine Strauß is with the University of Vienna (email: christine.strauss@univie.ac.at)

[‡]Christian Stummer is with Bielefeld University (email: christian.stummer@uni-bielefeld.de)

within the simulation to capture uncertainty and variability. It is therefore necessary to simulate a large number of attacks to evaluate a single "control portfolio", which is characterized by a string of binary design variables that indicate whether or not a particular control is applied to a particular asset. A control portfolio's fitness values are estimated by aggregating over a number of samples taken. In this context, different aggregation functions may be used for different criteria. The common approach of averaging across simulation runs may be complemented with minimum and maximum values, which is highly relevant in risk evaluations where worst case values are frequently more interesting than average values.

The combinatorial design space and the expensive simulation-based evaluation procedure make the simulation-optimization problem particularly challenging computation-wise. The optimization is implemented on top of the simulation core using Opt4J [12], a flexible framework for implementing and testing meta-heuristics. Initial experiences with this framework are promising. We are currently in the process of testing evolutionary solution procedures for the combinatorial optimization problem of identifying (an approximation of) the set of Pareto efficient security control portfolios with respect to multiple objectives, including both costs (e.g., implementation costs, running costs etc.) and benefits (e.g, reduced likelihood of successful attacks, reduced impact of attacks etc.).

Ultimately, we aim to provide interactive decision support using visualization methods such as those described in [13] or [14] to allow security managers and other stakeholders to explore the identified solution space and, thus, help them in making more profound decisions to improve IT security in their organizations.

In our talk, we introduce the problem setting, describe our approach, and illustrate the application by means of sample scenarios. In particular, we will report on initial experience with a standard meta-heuristic solution procedure, particular challenges for meta-heuristic optimization in the application domain, and discuss our ideas. As this is ongoing work, suggestions are not only welcomed, but highly appreciated.

Keywords: IT security management, simulation, combinatorial optimization, multi-objective meta-heuristic solution procedures, interactive decision support

Acknowledgements
The work presented in this abstract is developed within the project "MOSES3" funded by the Austrian Science Fund (FWF): P23122-N23. The research is carried out at Secure Business Austria, a COMET K1 program competence center supported by FFG - Austrian Research Promotion Agency.

References

[1] 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.

[2] E. Kiesling, C. Strauß, and C. Stummer, "A multi-objective decision support framework for simulation-based security control selection," in *Proceedings of the Seventh International Conference on Availability, Reliability and Security (ARES 2012)*, Prague, 2012, pp. 454–462.

[3] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph, "OWL 2 web ontology language primer," *W3C recommendation*, vol. 27, pp. 1–123, 2009. [Online]. Available: http://www.w3.org/TR/2009/REC-owl2-primer-20091022/all.pdf

[4] Stanford Center for Biomedical Informatics Research, "Protege ontology editor and knowledge acquisition system," 2009. [Online]. Available: http://protege.stanford.edu/

[5] 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.

[6] J. Wielemaker, T. Schrijvers, M. Triska, and T. Lager, "SWI-Prolog," *Theory and Practice of Logic Programming*, vol. 12, Special Issue 1-2, pp. 67–96, 2012.

[7] A. Panchenko and L. Pimenidis, "Towards practical attacker classification for risk analysis in anonymous communication," in *Proceedings of the 10th IFIP TC-6 TC-11 international Conference on Communications and Multimedia Security*, ser. CMS'06. Springer, 2006, pp. 240–251.

[8] P. Ammann, D. Wijesekera, and S. Kaushik, "Scalable, graph-based network vulnerability analysis," in *Proceedings of the 9th ACM Conference on Computer and Communications Security*. ACM, 2002, pp. 217–224.

[9] X. Ou, W. F. Boyer, and M. A. McQueen, "A scalable approach to attack graph generation," in *Proceedings of the 13th ACM Conference on Computer and Communications Security*, ser. CCS '06. ACM, 2006, pp. 336–345.

[10] R. E. Sawilla and X. Ou, "Identifying critical attack assets in dependency attack graphs," in *Proceedings of the 13th European Symposium on Research in Computer Security*, ser. ESORICS '08. Springer, 2008, pp. 18–34.

[11] S. Luke, C. Cioffi-Revilla, L. Panait, and K. Sullivan, "MASON: a new multi-agent simulation toolkit," in *2004 SwarmFest Workshop*, 2004.

[12] M. Lukasiewycz, M. Glaß, F. Reimann, and J. Teich, "Opt4J: a modular framework for meta-heuristic optimization," in *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, ser. GECCO '11. New York, NY, USA: ACM, 2011, pp. 1723–1730.

[13] 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.

[14] J. Gettinger, E. Kiesling, C. Stummer, and R. Vetschera, "A comparison of representations for discrete multi-criteria decision problems," *Decision Support Systems*, vol. 54, no. 2, pp. 976–985, Jan. 2013.