Game-Theoretic Algorithms for Optimal Network Security Hardening Using Attack Graphs

(Extended Abstract)

Karel Durkota, Viliam Lisý
Dept. of Computer Science
FEE, CTU in Prague
{durkota,lisy}@agents.fel.cvut.cz

Christopher Kiekintveld
Dept. of Computer Science
Univ. of Texas at El Paso, USA
cdkiekintveld@utep.edu

Branislav Bošanský
Dept. of Computer Science
Aarhus University, Denmark
bosansky@cs.au.dk

ABSTRACT

In network security hardening a network administrator may need to use limited resources (such as honeypots) to harden a network against possible attacks. Attack graphs are a common formal model used to represent possible attacks. However, most existing works on attack graphs do not consider the reactions of attackers to different defender strategies. We introduce a game-theoretic model of the joint problem where attacker's strategies are represented using attack graphs, and defender's strategies are represented as modifications of the attack graph. The attack graphs we use allow for sequential attack actions with associated costs and probabilities of success/failure. We present an algorithm for an computing attack policy that maximizes attacker's expected reward and empirical results demonstrating our methods on a case study network.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection

General Terms

Algorithms, Economics, Performance, Experimentation

Keywords

attack graphs; optimal attack policy; and-or graph; honeypots; game theory; network security

1. INTRODUCTION

Networked computer systems support a wide range of critical functions in both civilian and military domains. Securing this infrastructure is costly and there is a need for new automated decision support systems that aid human network administrators to detect and prevent attacks.

We focus on network security hardening problems in which a network administrator (defender) reduces the risk of attacks on the network by introducing honeypots (fake hosts or services) as intrusion detection sensors into their network [7]. Deciding how to optimally allocate these resources to reduce the risk of undetected attacks on a network is a challenging decision for the defender. We use game theory to model this adversarial interaction and to determine the best way to use honeypots against a well-informed attacker. We introduce a novel game-theoretic model for network hardening using honeypots extends Stackelberg security games [8] by adopting a compact representation of strategies for attacking computer networks called *attack graphs*.

Attack graphs (AGs) can represent a rich space of attacker actions sequences for compromising a specific computer network. AGs can be automatically generated based on known vulnerability databases [5] and they are widely used in the network security to identify the minimal subset of vulnerabilities to be fixed to prevent all known attacks [6], or to calculate security risk measures (e.g., the probability of a successful attack) [4]. We use AGs for computing the attacker's optimal attack plan to each defender's honeypot allocation action to measure defender's action effectivness.

We present a novel game-theoretic model of security hardening based on attack graphs and a case study analyzing the hardening solutions for sample networks.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

2. NETWORK HARDENING GAME

We model the network hardening problem as a Stackelberg game, where the defender acts first, taking actions to harden the network by adding up to k honeypots (HPs). The attacker is the follower who selects an optimal attack plan based on (limited) knowledge about the defender's strategy. In particular, we assume that the attacker learns the number and type of deployed HPs; however he does not know which specific hosts are HPs and which are real.

A game instance is based on a specific computer network (e.g., in Fig. 1a). A network has a set of host types, such as firewalls, workstations, etc. Two hosts are of the same type if they run the same services and have the same connectivity in the network (i.e., a collection of identical workstations is modeled as a single type). The same host types present the same attack opportunities, so they are represented only once in an attack graph. During an attack, a specific host of a given type is selected randomly with uniform probability. E.g., adding more HPs of a specific type increases the likelihood that the attacker who interacts with this host type will choose a HP instead of a real host. If the attacker interacts with a HP during an attack, he is immediately detected and the attack ends. The attacker is rational and maximizes the expected utility taking into account the probabilities of interacting with HPs, his actions' costs and success proba-

bilities, and rewards from successful attacks. Installing and maintaining HPs has a cost for the defender depending on the host type that is duplicated. He minimizes his total expected loss which consists of the expected loss for being attacked and the cost for deploying the HPs into the network. The Stackelberg equilibrium is found by selecting the defender's pure action that minimizes the expected loss under the assumption that the attacker responds with an optimal attack [8]. If the attacker is indifferent between multiple attacks, it is typical to break ties in favor of the defender [8].

Computation of the equilibrium relies on computing the optimal attack policy as described in the following section.

2.1 Attack Graphs and Attack Policies

An attack graph (AG) captures all known ways that the computer network can be compromised. It is an and-or graph consisting of the fact nodes (OR), logical statements about the network (i.e., access to a database), and actions nodes (AND), that change the statements from *false* to *true*.

To characterize the attacker's reaction to the set of honeypots, we compute a full *contingent attack policy* (AP), which defines an action for each situation that may arise during an attack as described in [1]. This allows identifying actions likely to be executed by a rational attacker as well as the *order* of their execution. The attacker chooses the optimal AP that maximizes his expected utility, which is an NP-hard problem [2]. We address this issue by translating AGs into an *Markov Decision Processes* and introducing several pruning techniques that reduce the computation considerably. First, we use a generalized version of the *Sibling-Class Theorem* from [2]. It states that in certain cases the optimal action order can be determined directly from the actions' success probabilities and costs, without any search. Second, we developed a heuristic to compute lower and upper bounds of the expected reward for the AG, which we use in a branch and bound manner to prune out the unpromising subtrees.

3. EXPERIMENTS

As an example of our result, we experimentally evaluated the proposed network hardening game on the network topology in Fig. 1a taken from [3]. This network consists of a server (srv), a vpn, a firewall (fw), a database (db) a group of 20 PCs (20grp) and a group of 4 PCs (4grp). The defender's expected loss (EL) without HPs is 1973. In Fig. 1b we present the optimal HP allocations computed by the game for different numbers of HPs k. For the first two HPs, it is best to duplicate the server and vpn—the network "door"—despite the fact that they are not the most valuable hosts. The lowest EL 617 is reached with 6 HPs by duplicating the database, server and vpn. Any additional HP only increases the EL, because their contribution in detecting the attack does not compensate their maintenance cost, so the cheapest option is selected to minimize the cost.

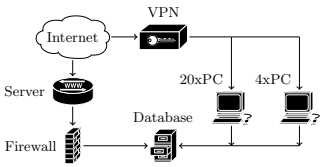

(a) Network

host	$k=1$	$k=2$	$k=6$	$k=7$
db	0	0	2	2
srv	1	1	2	2
vpn	0	1	2	2
4grp	0	0	0	1
EL	1524	1085	617	635

(b) HP allocations

Figure 1: (a) A small network topology. (b) The optimal HP allocations for 1,2,6 and 7 HPs and corresponding defender's expected loss (EL).

4. CONCLUSION

We introduce a game-theoretic model for the network hardening problem. The defender seeks an optimal deployment of honeypots into the network, while the attacker tries to attack the network and avoid the interaction with the honeypots. Our model provides a novel combination of using compact representation of the strategies of the attacker in the form of attack graphs, and using deception by the defender. By translating the attack graphs into MDPs and employing a number of pruning techniques, we are able to solve problems of realistic size and analyze the results for realistic case studies. We show that a few HPs can significantly reduce the defender's expected loss.

Our work has significant potential for further research. Since the majority of the required input data can be automatically acquired by standard network scanning tools, or extracted from existing vulnerability databases, the proposed model can be deployed in real-world networks and evaluated in practice. Secondly, our model can be further extended from the game-theoretical perspective and use additional uncertainty about the knowledge of the attacker, or model multiple types of the attacker using Bayesian variants of Stackelberg games.

Acknowledgments

This research was supported by the Office of Naval Research Global (grant no. N62909-13-1-N256) and Danish National Research Foundation and The National Science Foundation of China (under the grant 61361136003) for the Sino-Danish Center for the Theory of Interactive Computation.

REFERENCES

[1] K. Durkota and V. Lisý. Computing optimal policies for attack graphs with action failures and costs. In *STAIRS*, pages 101–110, 2014.

[2] R. Greiner, R. Hayward, M. Jankowska, and M. Molloy. Finding optimal satisficing strategies for and-or trees. *Artificial Intelligence*, pages 19–58, 2006.

[3] J. Homer, X. Ou, and D. Schmidt. A sound and practical approach to quantifying security risk in enterprise networks. *Kansas State University*, 2009.

[4] S. Noel, S. Jajodia, L. Wang, and A. Singhal. Measuring security risk of networks using attack graphs. *International Journal of Next-Generation Computing*, 1(1):135–147, 2010.

[5] X. Ou, W. F. Boyer, and M. A. McQueen. A scalable approach to attack graph generation. In *CCS*, pages 336–345, 2006.

[6] O. Sheyner, J. Haines, S. Jha, R. Lippmann, and J. Wing. Automated generation and analysis of attack graphs. In *IEEE S&P*, pages 273–284, 2002.

[7] L. Spitzner. *Honeypots: tracking hackers*. Addison-Wesley Reading, 2003.

[8] M. Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.