Computers & Industrial Engineering 90 (2015) 352–360

Contents lists available at ScienceDirect

Computers & Industrial Engineering

journal homepage: www.elsevier.com/locate/caie

How to deal with an intelligent adversary ☆

Andrey Garnaev [a,b,*], Melike Baykal-Gursoy [c,d], H. Vincent Poor [e]

[a] WINLAB, Rutgers University, North Brunswick, NJ 08901, USA
[b] Department of Computer Modelling and Multiprocessor Systems, Saint Petersburg State University, St. Petersburg 198504, Russia
[c] Department of Industrial and Systems Engineering, Rutgers University, Piscataway, NJ 08854-8018, USA
[d] Center for Advanced Infrastructure and Transportation, Rutgers University, Piscataway, NJ 08854-8018, USA
[e] Department of Electrical Engineering, Princeton University, Princeton, NJ 08544, USA

ARTICLE INFO

Article history:
Received 3 May 2015
Received in revised form 29 August 2015
Accepted 6 October 2015
Available online 10 November 2015

Keywords:
Network protection
Equilibrium
Bayesian game

ABSTRACT

Traditionally, the design of network protection strategies is based on the answers of a protector and an adversary to the question "How?": how should the protector allocate its protection resources, and how should the adversary allocate its attacking resources? This paper considers a more sophisticated adversary, who, planning its malicious activities, considers two questions: "What for?" and "How?". Namely, what is the motivation for the attack? and how to attack based on the chosen motivation? To study this problem, a simple game-theoretic network protection model is considered, in which the adversary decides whether to intrude on the network to inflict maximal damage or to perform a reconnaissance mission, and based on this decision an intrusion strategy is designed. The solution to this game shows that such an adversary may try a feint to draw the protector's efforts away from the nodes that the adversary intends to attack. Taking into account this feature of the adversary's behavior allows improvements in the reliability of a protection strategy.

1. Introduction

Computer networks have come to serve a critical societal role, but this has created a new type of terrorism, namely, cyber-terrorism. So many critical activities, such as commerce, finance, energy, education and health care are online, that gaining control of or disrupting such online systems, by Rainie, Anderson, and Connolly (2014), can sow panic, cause damage or even lead to loss of life. For example, by Magnuson (2014), cyber-attacks on electric utilities can be devastating, since "taking down an electric grid, especially one that serves a major city, could do real damage to the economy and may indirectly cost lives". Testifying to the House of Representatives Intelligence Committee on cyber threats, Admiral Rogers (see, Zengerle, 2014) said that a few countries have the ability to invade and possibly shut down computer systems of U.S. power utilities, aviation networks and financial companies, and these capabilities can be used by nation-states, groups or individuals to take down these critical activities. Cyber threats are only one of the challenges homeland security has to meet. Despite trillion-dollars investments (see, Mueller & Stewart, 2011), the resources are still inadequate to respond to an increasing number of old and new threats as adversaries (criminals or terrorists) create new non-trivial methods of attack. For this reason, the National Research Council (see, NRC, 2008) has emphasized the importance of modeling terrorists as intelligent adversaries, and has proposed three possible techniques to assess the impact of an intelligent adversary, one of which is game-theoretic modeling. The problem of security involves many different aspects, see, for example, a recent review of Hausken and Levitin (2012), where 129 published research papers on different aspects of security were classified according to the system structure, defense measures, attack tactics and circumstances involved.

Numerous researchers have used game theory to study resource-allocation decisions for network protection, see for example, Manshaei, Zhu, Alpcan, Basar, and Hubaux (2013), Guikema (2009) and Baykal-Gursoy, Duan, Poor, and Garnaev (2014) that provide references of research contributions that analyze and solve security problems in networks via game-theoretic approaches. In these works, the main setting is one in which the protector and the adversary seek answers to the same question, "How?" Namely, how to best allocate protection resources? how to best allocate attacking resources? In this paper we examine network protection from a different point of view, and, consider a more sophisticated adversary, who plans an attack or an intrusion by asking two questions: "What for?" and "How?". Namely, what is the motivation for

☆ This material is based upon work supported by the National Science Foundation under Grant Numbers CMMI-1436288 and CMMI-1435778.

* Corresponding author at: WINLAB, Rutgers University, North Brunswick, NJ 08901, USA.
E-mail addresses: garnaev@yahoo.com (A. Garnaev), gursoy@rci.rutgers.edu (M. Baykal-Gursoy), poor@princeton.edu (H. Vincent Poor).

http://dx.doi.org/10.1016/j.cie.2015.10.001
0360-8352/

intruding on the network? and how to intrude based on the chosen motivation? Of course, answers to these questions might lead to completely different adversarial behavior, than answering only the question "How?"

Admiral Rogers (see, Zengerle, 2014), in his testimony, pointed out that in addition to some countries already having the ability to shut down valuable U.S. computer systems, some digital attackers have also been able to penetrate such systems and perform "reconnaissance" missions to determine how the networks are put together. Such adversaries, planning their intrusions, had to answer the question: What is the purpose of the intrusion: to shut down the system or to perform "reconnaissance"?, and then to act according to the answer.

As an example of other purposes for intrusion, see Levitin, Hausken, Taboada, and Coit (2012), where a problem to store information securely if an adversary may steal or destroy the information was considered. Non-dominated solutions to this information security problem were found based on a multiple objective genetic algorithm.

To gain insight into this type of situation, we suggest a simple game, in which an adversary can intrude on a network to corrupt its nodes, and design its intrusion plan based on the chosen motivation. We consider two basic motivations: (a) to inflict maximal damage, and (b) to perform reconnaissance. Note that, in Garnaev, Baykal-Gursoy, and Poor (2014), it was shown, that a protection strategy may depend essentially on the type of attack, and incorporating a priori knowledge of the attack's type, which is fixed but unknown to the protector, increases defense efficiency. In this paper, we extend this approach by allowing the adversary to be more sophisticated and skillful in designing the intrusion, namely, allowing the adversary to choose consciously its motivation for intrusion, and to optimize its intrusion accordingly. This allows us to incorporate a human factor into the adversary's strategy.

The main contributions of this paper are the following:

(a) Developing a game-theoretic resource allocation model for inflicting a maximal damage attack on a network and for an intrusion attack into network to perform a reconnaissance mission.
(b) Incorporating a human factor into the adversary's behavior allowing him to choose consciously one of the types of attack.
(c) Showing the difference in the principles that the intrusion strategy and the detection strategy have to be based on in order to be optimal. Namely, the intrusion strategy has to be based on a tactical decision making approach allowing sudden switching between strategies. Meanwhile, the protection strategy has to be based on a strategic decision making approach incorporating the possibility of such tactical adversary's decision making by a proper allocation of protection resources in advance.

The organization of this paper is as follows: in Section 2 and its four subsections, we first model two types of attack on a network by means of resource allocation games. In both games the type of attack is fixed, and known to the rivals. In Section 3 and its two subsections, we extend the model to allow for a sophisticated adversary to choose the type (motivation) of intrusion. In Section 4, discussions and conclusions are offered. In the appendix, the proofs of the obtained results are supplied.

2. Two types of attack

In this section and its four subsections, we describe two game-theoretic models describing two types of attack on a network: to inflict maximal damage and to perform a reconnaissance mission.

2.1. Strategies

The game is played on a network. Here we have in mind a computer or communication network consisting of N nodes. It is an abstract network composed of communication links and nodes that may contain data that need to be protected. As such, the network does not correspond to any specific topology. In the network two *agents* (players, rivals) are present. An agent who wants to minimize the effects of an attack is called the *protector* (say, it can be an intrusion detection system (IDS)). An agent who wants to intrude the network is called the *adversary*. We assume that each game is played in one time slot with a total duration Y, during which the intrusion has to be detected. If it is not detected, it could yield some serious consequences, say, loss of valuable data, or loss in the network's security due to successful "reconnaissance". During the time slot the adversary might intrude a single node, i.e., an adversary's strategy is a vector $\boldsymbol{x} = (x_1, \ldots, x_N)$, where x_i is the probability that the adversary intrudes node i, and $\sum_{i=1}^{N} x_i = 1$. The protector has a more sophisticated strategy, namely, during the time slot, the protector can switch from one node to another to scan. Thus, its strategy corresponds to the amount of time it has to spend scanning each selected node, i.e., a protector's strategy is a vector $\boldsymbol{y} = (y_1, \ldots, y_N)$, where y_i is the scanning time of node i, and $\sum_{i=1}^{N} y_i = Y$.

2.2. Value of node and detection probability

Each node of the network is characterized by a value C_i (say, the amount of stored valuable data). We assume that the damage to node i equals to the value of the stolen data, and that all data stored in the corrupted node can be stolen, if the scanning failed. We consider only the direct cost of an attack including data loss, or financial losses caused. In addition to the direct cost, as suggested by Kumar and Liu (2014), indirect losses might arise, and they could be significantly higher than direct losses, since a successful attack could impact negatively on consumer behavior and investor confidence.

For simplicity we assume that the probability of not detecting the adversary depends exponentially on the scanning time, namely, it is $\exp(-\lambda_i y_i)$, if node i is corrupted, with λ_i as a scanning characteristic of node i. Thus, detection probability is $1 - \exp(-\lambda_i y_i)$. See, also Stone (2007), Iida, Hohzaki, and Sato (1994), Sakaguchi (1973), Lewis (2009), Baston and Garnaev (2000), Garnaev and Trappe (2014), as examples of using exponential dependence in network protection games.

2.3. Game with the maximal damage attack

In this section, we consider the scenario in which the adversary intrudes on the network to inflict maximal damage. The payoff to the adversary is the total expected damage this can cause, i.e., $v_A^D(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i=1}^{N} C_i x_i \exp(-\lambda_i y_i)$. The payoff to the protector is $v_P^D(\boldsymbol{x}, \boldsymbol{y}) = -v_A^D(\boldsymbol{x}, \boldsymbol{y})$. Thus, this is a zero-sum game (see, Fudenberg & Tirole, 1991). We assume that the rivals know the nodes' values C_i, the scanning characteristics λ_i for every node i, and the duration of the time slot Y. Recall that $(\boldsymbol{x}_*, \boldsymbol{y}_*)$ is an equilibrium (saddle point) of such a game if and only if $v_A^D(\boldsymbol{x}, \boldsymbol{y}_*) \leqslant v_A^D(\boldsymbol{x}_*, \boldsymbol{y}_*) \leqslant v_A^D(\boldsymbol{x}_*, \boldsymbol{y})$ for any $(\boldsymbol{x}, \boldsymbol{y})$.

For the sake of simplicity, we assume that all nodes have different values, i.e., $C_i \neq C_j$ for $i \neq j$, and without loss of generality, we can assume that the nodes are arranged by their values in decreasing order

$$C_1 > C_2 > \cdots > C_N. \tag{1}$$