

Game Theory Meets Information Security Management

Andrew Fielder¹, Emmanouil Panaousis², Pasquale Malacaria²,
Chris Hankin¹, and Fabrizio Smeraldi²

¹ Imperial College London

{andrew.fielder,c.hankin}@imperial.ac.uk

² Queen Mary University of London

{e.panaousis,p.malacaria,f.smeraldi}@qmul.ac.uk

Abstract. This work addresses the challenge “how do we make better security decisions?” and it develops techniques to support human decision making and algorithms which enable well-founded cyber security decisions to be made. In this paper we propose a game theoretic model which optimally allocates cyber security resources such as administrators’ time across different tasks. We first model the interactions between an omnipresent *attacker* and a team of system administrators seen as the *defender*, and we have derived the *mixed Nash Equilibria* (NE) in such games. We have formulated general-sum games that represent our cyber security environment, and we have proven that the defender’s *Nash strategy* is also *minimax*. This result guarantees that independently from the attacker’s strategy the defender’s solution is optimal. We also propose *Singular Value Decomposition* (SVD) as an efficient technique to compute approximate equilibria in our games. By implementing and evaluating a *minimax solver with SVD*, we have thoroughly investigated the improvement that Nash defense introduces compared to other strategies chosen by common sense decision algorithms. Our key finding is that a particular NE, which we call *weighted NE*, provides the most effective defense strategy. In order to validate this model we have used real-life statistics from Hackmageddon, the Verizon 2013 Data Breach Investigation report, and the Ponemon report of 2011. We finally compare the game theoretic defense method with a method which implements a *stochastic optimization algorithm*.

Keywords: Information security management, game theory, cyber security.

1 Introduction

Due to the growth of cyber attacks against government agencies and companies, there is a need for more investment to protect networks and platforms which allow the exchange and the storage of important information. Large organizations are more likely to employ a security manager whose main task is to lead *Information Security Management* to mitigate cyber security risks. However many companies do not have dedicated cyber security personnel, but the critical tasks required to

mitigate these risks must still be performed. In the latter case these tasks need to be carried out by the system administrators who must allocate part of their time for this. This is determined by other crucial non-security tasks. Since time is a scarce resource for system administrators, the available time to perform cyber security tasks should be ideally optimized. Hence there is a need for a method that gets as an input a time allowance and identifies how much time should be spent on tasks related to the defense of different parts of the infrastructure. The importance of this challenge has also been highlighted by Alpcan and Başar [4] (p. 9). This is a critical challenge given the fact that an attacker usually aims at exploiting one vulnerability of the system whilst the defending party must protect as many cyber security *targets* as possible. System administrators seek to maximize the mitigation of the greatest number of attacks against the various assets while an attacker aims to maximize the probability of penetrating the different “security layers” and complete their attack against a preferable target. This problem falls under the entire class of problems ISO 27001 examines with regards to balancing expenditure to the information security risk.

In this paper, we consider the *global asset* as the collection of all the data assets of an organization. We additionally consider a *target* to be a piece of hardware that has access to a subset of this asset. In our model the process of allocating resources is to improve the performance of the defense of a system. Our model defines the defense in terms of *all the actions that can be performed to improve the defense of a given target*. We assume that each target can be attacked by the exploitation of a vulnerability. The adversary follows a unique path to penetrate into the organization’s system and compromise a target. Examples of targets, vulnerabilities, as described by our model, and threats against these targets, can be found in 2013 Data Breach Investigation report [1] published by Verizon. By modeling a cyber security scenario in the form of targets that attacks can be performed against and the set of all actions that can protect that target, we are able to create a more compact and computationally feasible model removing the interdependencies between actions. According to our model, the attacker *probabilistically* chooses a target and plays with probability 1 all the actions required to attack this target. Their decision is based on the different profits they will obtain when they attack this target successfully and the baseline defense applied to each of the attack paths towards the targets. However, we assume that the attacker is not aware of the resource allocation strategy chosen by the defender nor conducts surveillance to identify such strategy.

The research question that this paper attempts to address is “*given that different targets have different weights how do we optimally allocate cyber security resources to defend such targets?*”. The defense of a target is indicated by two security levels, *baseline* and *best practice*. The baseline shows the probability of an attack against a target being mitigated, to guarantee that the basic security functionality of the systems is maintained, when no additional administrators time has been allocated. The *best practice* defense denotes the percentage of attacks that are mitigated assuming that the system is currently running the most up-to-date security software as well as having had all patches applied and any

potentially exploitable data removed. These tasks could be carried out by allocating additional system administrators' time. An example of baseline defense and resource allocation towards the provision of extra security is the following. We can assume that an unmonitored *spam filter* stops a certain percentage of junk e-mail, for example 60% which corresponds to the baseline defense. By having an administrator monitoring and upgrading the spam filters, more attacks can be prevented, and for instance 90% of the spam is mitigated. Thus under any resource allocation strategy, in terms of man hours, where an administrator is assigned to perform a spam filter based action, we have an improvement of 50% in stopping junk emails over the baseline defense.

In this paper we use game theory to model the cyber security environment and challenge of resource allocation as described above. Any resource allocation problem must consider the strategies available to the adversary and the cyber security team of the organization thus making game theory an appropriate tool to model such an environment and provide an effective set of solutions. According to [4] (p. 40), “*A game-theoretic framework for defensive decision-making has a distinct advantage over optimization as it explicitly captures the effects of the attacker behavior in the model, in addition to those of the defensive actions. Plain optimization formulations, on the other hand, focus only on the optimization of defensive resources without taking attackers into account.*”

The remainder of this paper is organized as follows. Section 2 discusses related work from the state-of-the-art on the intersection of game theory and security. In Section 3, we propose our game theoretic model while in Section 4 we undertake comparisons against alternative methods. Finally, Section 5 concludes this paper by summarizing its main contributions and highlighting future work to be undertaken to further improve the performance and the usability of our model.

2 Related Work

A significant amount of work on the intersection of economics and information security has been published by Grossklags. The authors in [3] discuss the factors that influence the decision process of individuals regarding their information security concerns. They contribute to the formulation of information policies and technologies regarding personal information security and privacy. The work published in [7] examines the weakest target game which refers to the case where an attacker is always able to compromise the system target with the lowest level of defense and not to cause any damage to the rest of the targets. The game theoretic analysis the authors have undertaken shows that the game leads to a conflict between pure economic interests and common social norms. While the former are concerned with the minimization of cost for security investments, the latter imply that higher security levels are preferable. A very thorough work has been published in [5] where the authors model security interactions by choosing different games such as weakest-link or best-shot to represent practical security scenarios. In these games, decision parameters are a protection level determining the security level chosen for the player's resources and a self-insurance level which mitigates damages when a successful attack occurs.

Security problems have been increasingly studied by using Stackelberg games to model the strategic interaction between a defender and an attacker. Some physical security games such as [10] assume the existence of targets that might be covered or uncovered during an attack. In [10], Kiekintveld et al. make the assumption that it is always preferable for the defender to defend as many targets as possible to obfuscate vulnerabilities over a period of time. The resulting players' payoffs depend only on the attacked target and whether this is covered by the defender or not. An important contribution is the work conducted by Korzhyk et al. in [11], showing how a *leader* (defender) should derive their strategy given that the security game could be either a Nash or a Stackelberg game. They also examine the case of a *follower* (attacker) who can attack more than one target. The authors show that Nash and Stackelberg strategies are the same in the majority of cases only when the follower attacks just one target.

Another recent contribution within the field of physical security has been published in [9], where the authors address the problem of finding an optimal defensive coverage. Because of the uncertainty about the attacker's payoffs they define such coverage as the one maximizing the worst-case payoff over the targets in the potential attack set. In a cyber security game we can consider that a defender may not know the payoffs for an attacker and they may be able to infer some bounds for those payoffs. Work published by Lye and Wing [13] uses a non-zero sum stochastic game model where an adversary is attacking an enterprise network and a network security administrator is defending the different network assets (e.g. web server, file server, workstation). A thorough survey of research contributions within the broad field of security and privacy in computer networks modeled by game-theoretic approaches has been carried out in [14].

3 Game Theoretic Formulation

We define a cyber security game as a game-theoretic model that captures essential characteristics of resource allocation decision making (i.e. system administrators' time) to prevent data loss and defend system and network assets of an organization. This is a two-player, non-cooperative static game between a system security team, defender \mathcal{D} , which defends an organization's information assets and data against external or insider adversaries who are modeled as an omnipresent attacker \mathcal{A} . Our model follows the definitions of the work done by Korzhyk et al. in [11]. The attacker \mathcal{A} might choose to attack any target from the available set of targets $T = t_1, t_2, \dots, t_n$ whilst \mathcal{D} aims to mitigate attacks by defending targets using resources. In our game the targets are data assets accessible by specific hardware components and system administrators are the limited resources. We consider a schedule as *a unique allocation of system administrators to perform tasks that will update the targets defense from baseline to best practice*. We define the set of available defense schedules as \mathcal{S} . The utility of \mathcal{D} when t_i is attacked and has no resources assigned to it equals the *baseline* defense utility $U_{\mathcal{D}}^{bl}(t_i)$. On the other hand, if \mathcal{D} has assigned some resource to defend t_i then their *best practice* defense utility equals $U_{\mathcal{D}}^{bp}(t_i)$. Finally, the utility in case where no attack has taken place equals zero.

Likewise for \mathcal{A} , we define their utility values as $U_{\mathcal{A}}^{bl}(t_i)$ and $U_{\mathcal{A}}^{bp}(t_i)$, respectively. Moreover, we denote the difference between best practice and baseline utilities of a target t_i for both players as $\Delta U_{\mathcal{D}}(t_i)$ and $\Delta U_{\mathcal{A}}(t_i)$. We observe that by applying a given resource to a target, the utility of \mathcal{D} for this target increases and the utility of \mathcal{A} decreases; namely $\Delta U_{\mathcal{D}}(t_i) > 0$ and $\Delta U_{\mathcal{A}}(t_i) < 0$. Similarly to the aforementioned example of a spam filter, let's assume that the baseline defense probability of a spam filter is 60% while the best practice defense probability equals 90%. If the damage caused to the organization due to a successful attack against this spam filter is given by $l = -100$ then the baseline utility value equals $U_{bl} = -40$ and the best practice equals $U_{bp} = -10$.

The normal form of this game is described as follows. \mathcal{A} 's pure strategy space is the set of targets, while their mixed strategy is denoted by $\mathbf{A} = \langle a_i \rangle$, where a_i represents the probability of attacking a target t_i . A schedule (or pure strategy) of \mathcal{D} is a feasible unique assignment of resources to cover (i.e. provide best practice defense) the different targets. Assuming that one resource is adequate to provide best practice defense for a target, a pure strategy is represented by a tuple $\mathbf{s} = \langle s_i \rangle \in \{0, 1\}^n$, where s_i equals 1 when t_i is defended by best practices; or 0 when only baseline defense is in place. Due to limited resources we define the feasible schedules by $\mathcal{S} \subseteq \{0, 1\}^n$. The number of schedules available to m administrators in the case of n targets equals $\frac{n!}{(n-m)!m!}$. For example for a scenario with 4 targets, when the number of system administrators equals two then the available schedules are:

$$\begin{aligned} \mathcal{S} = \{ & \langle \mathbf{1}, \mathbf{1}, \mathbf{0}, \mathbf{0} \rangle, \langle \mathbf{1}, \mathbf{0}, \mathbf{1}, \mathbf{0} \rangle, \langle \mathbf{1}, \mathbf{0}, \mathbf{0}, \mathbf{1} \rangle, \\ & \langle \mathbf{0}, \mathbf{1}, \mathbf{1}, \mathbf{0} \rangle, \langle \mathbf{0}, \mathbf{1}, \mathbf{0}, \mathbf{1} \rangle, \langle \mathbf{0}, \mathbf{0}, \mathbf{1}, \mathbf{1} \rangle \} \end{aligned}$$

In this paper we assume homogeneous resources, namely each resource can apply best practice defense equally for each of the targets, allowing all the possible resource allocation schedules to be played. The mixed strategy $\mathbf{D} = \langle d_s \rangle$ of the defender is a probability distribution over the different schedules, where d_s is the probability of playing a schedule $s \in \mathcal{S}$. We define the *coverage* induced by the strategy \mathbf{D} to be the vector $C = \langle c_i \rangle$, where the probability c_i of applying the best practice defense for a target i is given by $c_i = \sum_{s \in \mathcal{S}} s_i d_s$. Going back to our previous example with 4 targets and 2 system administrators, assuming a mixed strategy $\mathbf{D} = \langle 0.3, 0.2, 0.15, 0.15, 0.1, 0.1 \rangle$ for the defender then the coverage vector equals $C = \langle 0.65, 0.55, 0.45, 0.35 \rangle$. Following [11], we define the utility functions of both players as follows.

Definition 1 (Utilities of the Cyber Security Game). *When a strategy profile $\langle \mathbf{D}, \mathbf{A} \rangle$ is played, the utility values of both players are given by the following:*

$$U_{\mathcal{D}}(\mathbf{D}, \mathbf{A}) = \sum_{i=1}^n E_{D,C}(t_i) = \sum_{i=1}^n a_i (c_i U_{\mathcal{D}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{D}}^{bl}(t_i)) \quad (1)$$

$$U_{\mathcal{A}}(\mathbf{D}, \mathbf{A}) = \sum_{i=1}^n E_{A,C}(t_i) = \sum_{i=1}^n a_i (c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i)) \quad (2)$$

If both players act rationally the game theoretic solution, upon their simultaneous moves is the Nash Equilibrium (NE).

Definition 2 (NE of the Cyber Security Game). A pair of mixed strategies (\mathbf{D}, \mathbf{A}) forms an NE if the following are satisfied:

- I. \mathcal{D} plays a best-response that is $U_{\mathcal{D}}(\mathbf{D}, \mathbf{A}) \geq U_{\mathcal{D}}(\mathbf{D}', \mathbf{A}), \forall \mathbf{D}'$
- II. \mathcal{A} plays a best-response that is $U_{\mathcal{A}}(\mathbf{D}, \mathbf{A}) \geq U_{\mathcal{A}}(\mathbf{D}, \mathbf{A}'), \forall \mathbf{A}'$

In a game it is possible that there are many NE. However in presence of unknown attack distributions not all Nash defenses perform equally. Therefore we are interested in the particular NE that most favors the targets with the highest defender's utility. We define the *NE's weight* as:

$$\sum_{1 \leq i \leq n} (c_i U_{\mathcal{D}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{D}}^{bl}(t_i)).$$

Definition 3. The weighted NE has the maximal weight among all NE.

In the case that there exists more than one NE with maximal weight any of these equilibria can be chosen as the weighted NE. Notice the NE's weight differs from the defender's utility in that the attacker's strategy is ignored. In the remaining of this paper, unless otherwise stated, by NE we refer to the weighted NE.

3.1 Characterization of Nash Defense

In a real scenario attackers and defenders often have different preferences and criteria for evaluating the financial impact of a successful attack, as such it is unrealistic to expect that cyber security games are necessarily zero-sum. Thus, in our cyber security game, players typically weigh the same outcomes in different ways. In our model we have made the assumption that both players assess the different targets in a proportionally equivalent manner.

In the following we have proven that *the Nash defense is a minimax strategy for the defender which implies that it minimaximizes the attacker's utility in the cyber security game*. This follows the work reported in [11]. We represent the set of all defender's Nash strategies of a non-zero sum cyber security game G as S_{NE}^G and the set of all defender's minimax strategies as S_{MM}^G . We define a zero-sum cyber security game \hat{G} in which the baseline and best practice utilities of the attacker for each target are equivalent to their utilities in G , therefore:

$$\hat{U}_{\mathcal{A}}^{bp}(t_i) = U_{\mathcal{A}}^{bp}(t_i), \forall t_i \in T, \text{ and } \hat{U}_{\mathcal{A}}^{bl}(t_i) = U_{\mathcal{A}}^{bl}(t_i), \forall t_i \in T \quad (3)$$

Since \hat{G} is a zero-sum game, we know that $\hat{U}_{\mathcal{D}}^{bp}(t_i) = -\hat{U}_{\mathcal{A}}^{bp}(t_i)$ and $\hat{U}_{\mathcal{D}}^{bl}(t_i) = -\hat{U}_{\mathcal{A}}^{bl}(t_i)$. We define a function μ which maps the attacking probabilities of each target in \hat{G} to the attacking probabilities in G as defined in [11]:

$$\mu(a_i) = \hat{a}_i = \lambda a_i \frac{\Delta U_{\mathcal{D}}(t_i)}{\Delta U_{\mathcal{A}}(t_i)}, \forall t_i \in T \quad (4)$$

where $\mu(a_i)$ is the probability of a target t_i to be attacked in \hat{G} , therefore $\mu(A) = \hat{A}$. Respectively the set of defender's Nash strategies in G is represented by S_{NE}^G and the set of defender's minimax strategies by S_{MM}^G .

Proposition 1. $\langle D^*, A^* \rangle$ is NE profile in G iff $\langle D^*, \mu(A^*) \rangle$ is NE profile in \widehat{G} .

Proof. To prove proposition 1, we have combined Lemmas 1 and 2 presented in appendix. We have that A^* and $\mu(A^*)$ are the attacker’s best responses, in G and \widehat{G} respectively, to the same defender’s best response D^* . Therefore the tuple of strategies $\langle D^*, A^* \rangle$ is NE profile in G iff $\langle D^*, \mu(A^*) \rangle$ is NE profile in \widehat{G} . This implies that D^* is a Nash defense in G iff it is also a Nash defense in \widehat{G} . \square

Theorem 1. A defender’s Nash strategy in the non-zero sum cyber security game is also a defender’s minimax strategy in the same game.

Proof. The minimax theorem states that for zero sum games NE and minimax solution coincide. Thus D^* is a Nash defense in \widehat{G} iff D^* is a minimax defender’s strategy in \widehat{G} namely $D^* \in S_{MM}^{\widehat{G}}$. This means D^* minimaximizes the utility of the attacker in the zero-sum game where the defender’s strategy is as in G :

$$D^* = \operatorname{argmin}_D \max_{\widehat{A}} U_{\mathcal{A}}(D, \widehat{A}).$$

We also have that $U_{\mathcal{A}}(D, \widehat{A})$ equals:

$$\sum_{i=1}^n \widehat{a}_i (c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i)) = \sum_{i=1}^n \lambda a_i \frac{\Delta U_{\mathcal{D}}(t_i)}{\Delta U_{\mathcal{A}}(t_i)} (c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i))$$

From the above $U_{\mathcal{A}}(D, \widehat{A}) \leq U_{\mathcal{A}}(D, \widehat{A}') \iff U_{\mathcal{A}}(D, A) \leq U_{\mathcal{A}}(D, A')$ because the left hand side of the \iff is the right hand side multiplied by the constant(s) $\lambda \frac{\Delta U_{\mathcal{D}}(t_i)}{\Delta U_{\mathcal{A}}(t_i)}$. Hence minimum and maximum of the ordering are preserved, therefore a minimax in \widehat{G} is also a minimax in G . Using proposition 1 we conclude that a defense strategy D^* is NE equilibrium in G and \widehat{G} therefore it is a minimax in \widehat{G} and, by the above argument, it is also a minimax in G . \square

This result guarantees that *independently from the attacker’s strategy the defender’s solution is optimal.*

3.2 Minimax Solver with Singular Value Decomposition

Given the high number of possible targets within an organization the game formulation is subject to the “state explosion problem” [6]. Therefore there is a need for a method which will compute Nash defenses for large number of targets. This must be computationally fast and provide reasonable precision in the calculation of the equilibrium. A possible approach is to look for an *abstraction* of the payoff matrices where only the most relevant features of such matrices are kept and minor features are discarded. We achieve this by using *Singular Value Decomposition* (SVD) to reduce the rank of the individual payoff matrices. To the best of our knowledge, there are no general results about game solutions for rank reduced payoff matrices. The closest work, undertaken by Kannan and Theobald [8], performs rank reduction of the matrix given by the sum of the payoff matrices of the two players.

Based on the Ponemon report 2013 [12], we have illustrated the payoff matrices in Table 1. We notice that these matrices have a *particular underlying structure*

which is mostly captured in the very first singular values. In other words, there are few dominant singular values, usually two or three. By ignoring the singular values after the second or third largest, we obtain reasonably close solutions to the original game solutions and a dramatic speed up of the computation. As an example for a game with 8 targets, 1 security administrator, and payoff matrices illustrated in Table 3, the singular values are $\langle 792, 108, 91, 42, 22, 12, 10, 3 \rangle$. This means the first component already provides a reasonable approximation to the original matrix and the components after the third component can be treated as “noise”. The speed-up in performance is mainly caused by the fact that the rank reduction, in our games, results in a large number of strategies becoming dominated. For example in a game with 120 strategies and rank 10, around $3/4$ become dominated when the rank is reduced to 6. By comparing the equilibria found in a 10 targets, 2 system administrators game and its SVD rank 2 abstraction (Table 2), we found that there is *a performance improvement of more than 1000 times* while the approximate solution only slightly deviates from the precise solution.

4 Model Comparison

To evaluate our game theoretic model we want to compare the performance of the *Nash resource allocation method (Minimax solver with Singular Value Decomposition)* which is given by the NE of the cyber security game, against alternative methods. The methods we have selected for comparison are based on approximations to two common sense approaches that one might consider, called *Uniform* and *Weighted* as well as a stochastic optimization algorithm called *Acceptable Coverage*.

The *Uniform* defense distribution gives no preference to any target that a defender wants to defend meaning that all targets are given equal probability to be defended. This method is a naive approach that assumes no knowledge on the part of the defender to decide how much each target is valued or how likely a target is to be attacked.

The *Weighted* defense distribution creates a probability distribution based on the value of each target that is being defended. According to this approach, the time a defender allocates to protect a target is proportional to its importance. A distribution of time across schedules is calculated such that each target is covered and the sum of the probability for playing each of the particular sets of defenses equals the intended coverage for each target. For example if a target t_1 should be defended with probability 0.2, then all schedules that include a defense of t_1 should sum to 0.2. This results in more time being scheduled to perform tasks which improve the defense of the targets that are either more vulnerable or more valuable to the organization. While this does not necessarily represent the best possible cyber security decision, this method identifies at least an average decision maker, that provides a reasonable defense by which to measure the improvement that the Nash defense introduces.

Acceptable Coverage. We compare the minimax solution to a *stochastic optimization algorithm with uncertainty in the attacker’s payoffs*. This approach is similar to the one published in [9] where the authors define a defensive coverage as the one maximizing the worst-case payoff over the targets in the potential attack set¹. For a given target we have two intervals from which the attacker gets their utility based on the applied defense; baseline or best practice. These are represented as $[U_{\mathcal{A}}^{\min,bl}, U_{\mathcal{A}}^{\max,bl}]$ and $[U_{\mathcal{A}}^{\min,bp}, U_{\mathcal{A}}^{\max,bp}]$, respectively.

When uncertainty about the attacker’s payoff is introduced, the concept of attack set gets more involved because we do not know what the maximum expected payoff for the attacker is. One is then led to define a potential attack set for a coverage as *the set of all targets that could give the attacker the maximum expected value*, for any attacker’s payoffs that can be extracted from the intervals of the payoffs. Given a coverage C the attacker can be *guaranteed* a payoff of at least the maximum of the minimum values over all targets; $R(C) = \max_{t_i} E_{A,C}^{\min}(t_i)$ where $E_{A,C}^{\min}(t_i) = a_i(c_i U_{\mathcal{A}}^{\min,bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{\min,bl}(t_i))$. $R(C)$ defines the potential attack set associated to the coverage C as $A(C) = \{t_i | E_{A,C}^{\max}(t_i) \geq R(C)\}$.

Following these ideas we look for a defender’s strategy giving the *best utilities guaranteed*. As such the defender will choose a strategy providing best coverage, where a coverage is better than another coverage if the former guarantees both a higher minimum expected utility for the defender and a lower maximum expected utility for the attacker compared to the latter. More formally we define the coverage ordering as follows:

$$C \geq C' \iff \min_i(E_{D,C}(t_i)) \geq \min_i(E_{D,C'}(t_i)) \wedge \max_j(E_{A,C}^{\min}(t_j)) \leq \max_j(E_{A,C'}^{\min}(t_j))$$

Given a set of coverages maximal with respect to the above ordering a coverage is *acceptable* if it is the maximal element, with respect to the defender’s utility, of this set. If a coverage is acceptable then we have considered all attacker’s payoffs that can be extracted from the payoffs intervals otherwise there would be a coverage above it, contradicting maximality.

Comparisons. We have undertaken simulations to identify how effective a Nash defense performs against distributions created by the *Uniform*, *Weighted* and *Acceptable Coverage* methods. For the purposes of testing we consider an organization like an online store whose assets include users payment details, website configurations and data related to operational procedures. The most valuable target would be a database server with the most of clients’ details. This could be compromised by an SQL injection attack. The second most valuable target would be the store’s website that an attacker could deface by modifying data of the web server the website is located at. The rest of the targets could consist of the mail server and different workstations. The device of each user, seen as a target, has a value related to the privileges that a user obtains. To undertake our comparisons, the values of utilities for each target have been taken from a report on data breaches [12], where the utility per target is given by mid-value for each

¹ In security games without uncertainty, the attack set is defined as the set of all targets that give the attacker the maximum expected payoff, given some coverage.

Table 1. Financial cost of an organization based on the findings of the Ponemon 2011 report [12]

target ID	t_1	t_2	t_3	t_4	t_5
cost (\$)	10k	30k	50k	75k	150k
target ID	t_6	t_7	t_8	t_9	t_{10}
cost (\$)	250k	350k	450k	750k	3m

Table 2. Times and defender utilities for a 10 targets 2 system administrators games using SVD

Rank	Time (s)	Value
10 (no SVD)	204	-615.744
8	60	-614.854
6	1.8	-617.732
4	0.4	-616.124
2	0.1	-616.651

Table 3. Probability of Attack per Target based on Statistics from Hackmageddon [2]

ID	t_1	t_2	t_3	t_4	t_5
8 t.	0.011	0.016	0.016	0.039	0.061
10 t.	0.005	0.005	0.011	0.016	0.016
ID	t_6	t_7	t_8	t_9	t_{10}
8 t.	0.127	0.332	0.398	-	-
10 t.	0.038	0.060	0.125	0.327	0.392

Table 4. Probability of Attack per Target based on Statistics from Verizon [1]

ID	t_1	t_2	t_3	t_4	t_5
8 t.	0.060	0.060	0.060	0.080	0.133
10 t.	0.021	0.021	0.021	0.021	0.063
ID	t_6	t_7	t_8	t_9	t_{10}
8 t.	0.222	0.222	0.244	-	-
10 t.	0.008	0.125	0.208	0.208	0.229

of the ranges of damage reported, and are presented in Table 1. We consider two different sizes of online stores in terms of reputation, number of clients and amount of data held. The larger store has 10 targets while the smaller one has 8. We assume that the organization is limited to the number of system administrators that they have available to them and cannot cover all the targets at one time. We test the defense of a system with up to 3 administrators. These choices seem reasonable for such organizations in our proof-of-concept comparisons.

The utilities for the attacker in this scenario are set at 20% of the value that the defender has for each given target. This is to represent that an attacker will still receive a large payoff for successfully attacking a target. However while the defender may have long term damage, the attacker's profit will generally be given by the immediate impact that their actions have. For the comparisons we have used 2 specific attack distributions given by the Hackmageddon [2] and Verizon data sets [1], where any of the attack distributions, shown in Tables 3 and 4, represents an unknown attacker that is attempting to breach a target.

While we have identified the range of damages that can be expected from a successful attack, the values themselves do not necessarily reflect a single attack, so we consider *varying the utilities both players perceive for each target*. By varying the utility of each target, we are capable of identifying if the Nash method performs better than either of the other methods independently of the cost of each target. We vary the utilities as presented in Table 1 by adjusting the values perceived by the defender, where the experimental bounds allow for a deviation of up to 10% from the stated utilities for each target. The baseline defense has been fixed at a 50% breach rate for each target while the best practice

defense reduces that rate to 20%. These numbers have been decided artificially for the model validation.

Performance Evaluation. We define the performance of a solution as the *average amount of damage expected from any single attack*, where the expected damage is calculated as the result of an attack on a given target. In each single attack one or more targets are defended following a schedule determined by the specified defender's strategy. If a target that is being attacked is not defended, then the damage equals the *baseline* damage. If the target is defended, the loss equals the damage with the *best practice* defense. To measure the performance of a solution we have created a simulated environment, in Python, used to perform attack sampling. For all comparisons performed, a sample size of 100,000 attacks was used. Such a sample is referred to as a *run* in the results. In the following we present the results, where 25 runs have been performed in each case and the average of the percentage improvement and the standard deviation seen across the runs have been plotted. We have set the rank values for SVD to 5 in both 8 and 10 target scenarios. These values allow us to minimize the runtime for each of the comparisons while maximizing the accuracy of the results.

The percentage improvement seen in comparison to the uniform defense shows a minimum improvement of approximately 15%. Comparably, the smallest average improvement for Nash defense over the *Weighted* is around 7%. In addition to this, the maximum improvement seen in the Nash solution over the *Uniform* is approximately 50%, where the maximum improvement over a *Weighted* defense does not exceed 25%. The improvements between different numbers of administrators for defending a single system identify the impact that the addition of an administrator has in improving the defense. We see that with 8 targets for both attack distributions the addition of more administrators increases the improvement seen in the defense of the system by more, over the common sense approaches. A large difference in improvement indicates that adding an additional administrator will have a greater impact on the defense of the system than in the case of a lower level of improvement. This growth however saturates, as when the number of administrators tends towards the number of targets the improvement seen tends to zero. This happens because all targets can be covered thus the expected damage is minimized across the whole system for all defender strategies.

When comparing the Nash and *Acceptable Coverage (AC)*, we see that unlike the comparisons to the common sense approaches, the Nash defense does not always perform better. In contrast to the approximate performance as measured against the common sense approaches, the average improvement does not exceed 10% for either method. With the 8 target scenario, the results show that the AC performs better for 1 administrator regardless of the data set used; specifically the AC method appears to be approximately 2% to 3% better. However, when there are 2 system administrators, the Nash defense performs between 2% and 6% better depending on the data set. With 3 system administrators the results show that on average, for a Hackmageddon attacker, the AC method performs better, but the Nash defense is preferable with a Verizon attacker. It should be

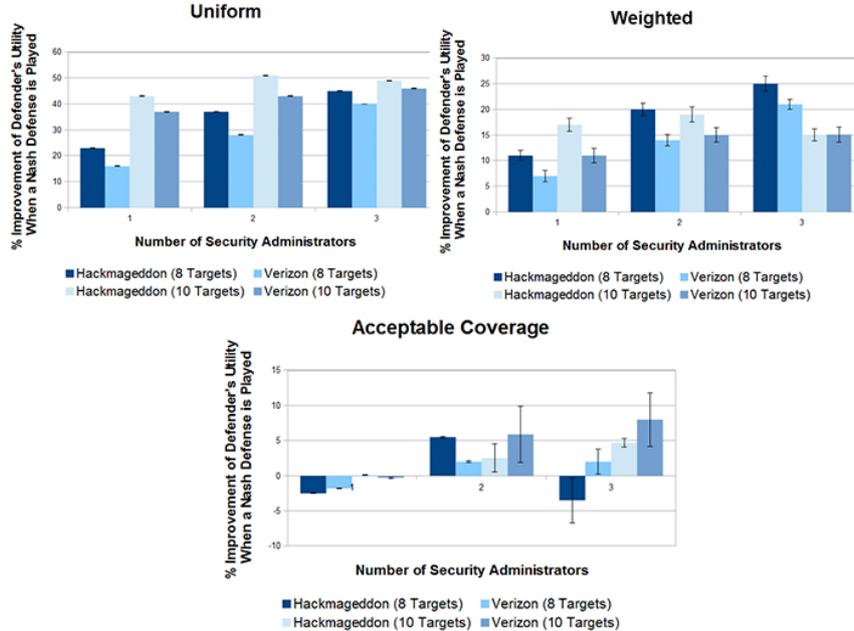


Fig. 1. Improvement of Nash defense over the different methods for both Hackmageddon (H) and Verizon (V) attackers

considered that the deviation of these results shows that in some circumstances the improvement of one method over the other can be less than 1%. For 10 targets and 1 administrator we see that with both datasets there is no improvement in the average performance of the Nash solution over AC, with a very small standard deviation. This indicates that the allocation of system administrators' time to targets is similar for both Nash and AC, where the randomized sampler gives the deviation in the results. For more than 1 administrator it appears to show a relatively large positive improvement for the Nash defense over the AC, where the range improvement is approximately 2% and 5% for 2 and 3 administrators respectively using a Hackmageddon attacker. For a Verizon attacker the improvements are close to 6% and 8% for 2 and 3 administrators. As has been noted above, there is a small deviation in the approximate scheduler used for the AC method, which may account for some of the improvement seen by the Nash defense and the reason for the large deviations in some of the results.

5 Conclusion

In this paper, we have introduced a cyber security model along with game theoretic tools to prove that common sense techniques are not as efficient at providing effective defense schedules as the ones computed by game theory. Our future

work includes interviews with system administrators to define levels of values for the different model components. We are also interested in validating the effectiveness of the model when considering an organization where the number of targets available to be attacked will range between 25 and 100. The efficiency of performance for SVD is important given that there is a trade-off in time and accuracy, where we have seen that even for small games it is computationally inefficient to derive the equilibria. As such an interesting extension to this work would be to measure the performance of SVD in terms of efficiency when large games are played. We have described an environment where the defense of a system is pro-active, but we don't consider the scheduling of time in a reactive manner in order to recover systems after a successful attack. In order to recover a system, the time of available system administrators needs to be assigned to this task, which limits the number of system administrators available to maintain the defense of the system. Therefore a further extension of the game considers the concept of optimally dividing the available resources between the recovery of the system and the maintenance of pro-active security. One limitation of the current model, that future work will address, is that in many cases the methods taken by an attacker to break into a system may have steps that are relevant to more than one target. In this case there are specific actions that can be performed by a system administrator that effectively cover multiple targets.

References

1. 2013 data breach investigations report by verizon, <http://www.verizonenterprise.com/DBIR/2013/>
2. Hackmageddon.com, <http://hackmageddon.com/> (accessed october 2013)
3. Acquisti, A., Grossklags, J.: Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behavior. In: Proc. of the 2nd Annual Workshop on Economics and Information Security (WEIS 2003), vol. 3 (2003)
4. Alpcan, T., Başar, T.: Network security: A decision and game-theoretic approach. Cambridge University Press (2010)
5. Grossklags, J., Christin, N., Chuang, J.: Secure or insure?: a game-theoretic analysis of information security games. In: Proc. of the 17th International Conference on World Wide Web (WWW 2008), pp. 209–218. ACM (2008)
6. Hankin, C., Malacaria, P.: Payoffs, intensionality and abstraction in games. In: Coecke, B., Ong, L., Panangaden, P. (eds.) Abramsky Festschrift. LNCS, vol. 7860, pp. 69–82. Springer, Heidelberg (2013)
7. Johnson, B., Grossklags, J., Christin, N., Chuang, J.: Nash equilibria for weakest target security games with heterogeneous agents. In: Jain, R., Kannan, R. (eds.) Gamenets 2011. LNICST, vol. 75, pp. 444–458. Springer, Heidelberg (2012)
8. Kannan, R., Theobald, T.: Games of fixed rank: A hierarchy of bimatrix games. *Economic Theory* 42(1), 157–173 (2005)
9. Kiekintveld, C., Islam, T., Kreinovich, V.: Security games with interval uncertainty. In: Proc. of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), pp. 231–238. International Foundation for Autonomous Agents and Multiagent Systems, Richland (2013)

10. Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Ordóñez, F., Tambe, M.: Computing optimal randomized resource allocations for massive security games. In: Proc. of the 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), pp. 689–696 (2009)
11. Korzhyk, D., Yin, Z., Kiekintveld, C., Conitzer, V., Tambe, M.: Stackelberg vs. nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research* 41(2), 297–327 (2011)
12. Ponemon Institute LLC. State of web application security, <http://www.ponemon.org/library/2011-cost-of-data-breach-united-states>
13. Lye, K.W., Wing, J.M.: Game strategies in network security. *International Journal of Information Security* 4(1-2), 71–86 (2005)
14. Manshaei, M., Zhu, Q., Alpcan, T., Basar, T., Hubaux, J.: Game theory meets network security and privacy. *ACM Transactions on Computational Logic*, 25:1–25:39 (2011)

Appendix

Lemma 1. *An attacker’s strategy A^* is best response to a defender’s strategy D in G iff $\hat{A}^* = \mu(A^*)$ is the attacker’s best response to D in \hat{G} .*

Proof. To prove this lemma we must prove that

$$U_{\mathcal{A}}^G(D, A^*) - U_{\mathcal{A}}^G(D, A') \geq 0 \Leftrightarrow U_{\mathcal{A}}^{\hat{G}}(D, \hat{A}^*) - U_{\mathcal{A}}^{\hat{G}}(D, \hat{A}') \geq 0, \forall A'$$

Solving (4) for a_i we have that $a_i = \frac{1}{\lambda} \hat{a}_i \frac{\Delta U_{\mathcal{A}}(t_i)}{\Delta U_{\mathcal{D}}(t_i)}$ and as we know that $\Delta U_{\mathcal{A}}(t_i), \Delta U_{\mathcal{D}}(t_i), \lambda \geq 0$ the following holds:

$$\begin{aligned} & U_{\mathcal{A}}^G(D, A^*) - U_{\mathcal{A}}^G(D, A') \geq 0 \\ \Leftrightarrow & \sum_{i=1}^n a_i^*(c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i)) - \sum_{i=1}^n a'_i(c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i)) \geq 0 \\ & \Leftrightarrow \sum_{i=1}^n (a_i^* - a'_i)(c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i)) \geq 0 \\ & \Leftrightarrow \sum_{i=1}^n \frac{1}{\lambda} \frac{\Delta U_{\mathcal{A}}(t_i)}{\Delta U_{\mathcal{D}}(t_i)} (\hat{a}_i^* - \hat{a}'_i)(c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i)) \geq 0 \\ \Leftrightarrow & \sum_{i=1}^n (\hat{a}_i^* - \hat{a}'_i)(c_i U_{\mathcal{A}}^{bp}(t_i) + (1 - c_i) U_{\mathcal{A}}^{bl}(t_i)) \geq 0 \Leftrightarrow U_{\mathcal{A}}^{\hat{G}}(D, \hat{A}^*) - U_{\mathcal{A}}^{\hat{G}}(D, \hat{A}') \geq 0 \end{aligned}$$

□

Lemma 2. *A defender's strategy D is best response to an attacker's strategy A in G iff D is also best response to the attacker's strategy $\hat{A} = \mu(A)$ in \hat{G} .*

Proof. To prove this lemma we must prove that $U_D^G(D^*, A) - U_D^G(D', A) \geq 0 \Leftrightarrow U_D^{\hat{G}}(D^*, \hat{A}) - U_D^{\hat{G}}(D', \hat{A}) \geq 0, \forall D'$. We prove this inequality as follows

$$\begin{aligned}
 & U_D^G(D^*, A) - U_D^G(D', A) \geq 0 \\
 \Leftrightarrow & \sum_{i=1}^n a_i (c_i^* U_D^{bp}(t_i) + (1 - c_i^*) U_D^{bl}(t_i)) - \sum_{i=1}^n a_i (c_i' U_D^{bp}(t_i) + (1 - c_i') U_D^{bl}(t_i)) \geq 0 \\
 & \Leftrightarrow \sum_{i=1}^n a_i (c_i^* \Delta U_D(t_i) + U_D^{bl}(t_i)) - \sum_{i=1}^n a_i (c_i' \Delta U_D(t_i) + U_D^{bl}(t_i)) \geq 0 \\
 \Leftrightarrow & \sum_{i=1}^n a_i (c_i^* - c_i') \Delta U_D(t_i) \geq 0 \Leftrightarrow \sum_{i=1}^n \frac{1}{\lambda} \frac{\Delta U_A(t_i)}{\Delta U_D(t_i)} \hat{a}_i (c_i^* - c_i') \Delta U_D(t_i) \geq 0 \\
 & \Leftrightarrow \sum_{i=1}^n \hat{a}_i (c_i^* - c_i') \Delta U_A(t_i) \geq 0 \Leftrightarrow \sum_{i=1}^n \hat{a}_i (c_i^* - c_i') \geq 0 \\
 \Leftrightarrow & \Delta \hat{U}_D(t_i) \geq 0 \sum_{i=1}^n \hat{a}_i (c_i^* - c_i') \Delta \hat{U}_D(t_i) \geq 0 \Leftrightarrow U_D^{\hat{G}}(D^*, \hat{A}) - U_D^{\hat{G}}(D', \hat{A}) \geq 0
 \end{aligned}$$

□