Intrusion Detection and Response: A Game Theoretic Perspective

Yihua Liao and V. Rao Vemuri

Department of Computer Science, University of California, Davis, One Shields Ave, Davis, CA 95616, USA
{yhliao, rvemuri}@ucdavis.edu

Abstract. While intrusion detection technologies have benefited from decades of study, there has been a lack of research into decision making tools required to evaluate the cost-effectiveness of intrusion detection systems (IDSs) and employ them properly. In this paper, we propose a game theoretic methodology for cost-benefit analysis and design of IDS. Game theory provides a natural setup for modeling the strategic interdependence between an IDS and an attacker. We use a simple two-person, nonzero-sum game to address the fundamental decision problems such as "does an organization need an IDS?" and "how should an IDS operator respond to the IDS alarms?". The solutions based on the game theoretic analysis integrate the cost-effectiveness and technical performance tradeoff of the IDS and identify the best defense and attack strategies. We also discuss the extensions of our method as well as the challenges of game theoretic modeling of security systems.

Key words: intrusion detection, cost-effectiveness, game theory, dominant strategy, Nash Equilibrium

1 Introduction

In May 2003, the Gartner Information Security Hype Cycle report declared that intrusion detection systems (IDSs) are "a market failure" and will be obsolete by 2005 due to the problems such as excessive false positives and false negatives, high operational cost and taxing incident-response process [1]. The report has stirred fierce debate within the IDS vendor as well as research communities. While it is debatable whether the Gartner's prediction for IDS is short-sighted or not, it is clear that cost-effectiveness will be one of the deciding factors in IDS' future.

IDS began in the 1980s as a promising paradigm for detecting hackers and malicious insiders that exploit security vulnerabilities or flaws in computer systems [2]. For the last two decades, most research efforts have been devoted to improve the technical effectiveness of IDS. That is, to what degree does an IDS detect and prevent intrusions into the target system, and how good is it at reducing false positives? In practice, however, no IDS will ever be 100% accurate in

detecting attacks. False positives and false negatives will be inevitably produced. Moreover, the reduction of one type of error (false positive or false negative) is usually accompanied by an increase in the other type.

Cost-effectiveness is an important, yet often overlooked aspect of IDS. When an organization makes an investment decision on a security mechanism such as IDS, risk assessment and cost-benefit analysis is essential. This includes assessment of the organization's assets and values, identification of threats and vulnerabilities, cost-benefit tradeoff evaluation, and so on. The major cost factors that ought to be taken into consideration are the operational cost of IDS, the expected loss due to intrusions and the cost of manual or automatic response to an intrusion [3]. Even when the adoption of IDS technology is justifiable, the IDS operator still faces the challenge of employing the IDS properly and determining the best response strategies against various types of attacks in order to minimize the cost of maintaining the IDS while protecting the system assets.

Our research aims to provide a game theoretic methodology for analysis and design of IDS and improve the effectiveness of IDS technology by modeling the interaction between IDS and attackers in a game playing context. Game theory offers a natural setup for adversarial situations, where multiple players with different objectives compete and interact with each other. As a powerful strategic decision making tool, game theory has been applied in many fields, including economics, political science, etc. [4]. The game theoretic modeling of security systems such as IDS makes it possible to bring the full spectrum of well-developed game theory techniques to bear on the information security problems.

In this paper, we propose to use a simple two-person, nonzero-sum game to model and analyze the IDS and attacker behavior in a general environment. Attacking and defending of the protected system are formalized in terms of a set of strategies for attacker and IDS, respectively, whereas risk and objectives are formalized in terms of payoff (or utility) functions. Each player strives to maximize his payoff function by selecting a feasible strategy. The solutions based on the game theoretic analysis naturally integrate the cost-effectiveness and technical performance tradeoff of the IDS and identify the "best" defense and attack strategies. Specifically, our results provide valuable insights in answering the following fundamental questions:

- Under what condition would an attack most likely occur?
- When is an IDS useful? When does its use become counterproductive?
- When an IDS is deemed useful, what should be its technical specification?
- What's the best response strategy when the IDS raises an alarm? Ignore it or respond to it?
- If the IDS operator can only respond to a subset of the alarms, what percentage is optimal?

The rest of the paper is organized as follows. In Section 2 we review some related work. Section 3 is a brief introduction to game theory. Section 4 describes details of our game model. In Section 5, we discuss the extensions of our method as well as the challenges of game theoretic modeling of security systems. Section 6 presents conclusions.

2 Related Work

The economics of information security is an emerging and growing research area [5] [6]. For example, Gordon and Loeb [7] presented an economic model that determines the optimal amount to invest to protect a given set of information assets. Iheagwara [8] examined the effect of implementation methods, management methods and intrusion detection policy on the return of investment.

The application of game theory to the domain of computer security has recently been a topic of interest. Lye and Wing [9] constructed a stochastic game to model the interactions between an attacker and the administrator in a typical network environment. Liu and Zang [10] presented a general incentive-based method to model attacker's intent, objectives and strategies (AIOS) and employed game theoretic techniques to infer AIOS. Different game models (stochastic or Bayesian games) were proposed based on the accuracy of intrusion detection and the correlation among attack actions. Alpcan and Basar [11] argued that game theory can provide the much needed mathematical framework for analysis, decision and control processes for information security and intrusion detection. They designed a security warning system for distributed IDS using Shapley values. A two-person finite game was used to model security attacks in a multiple sensor environment. Cavusoglu and Raghunathan [12] took a game theoretic approach to determine the optimal configuration (detection and false positive rates) of an IDS and compared it with the decision theory approach. While the game models in this paper are similar to those in [11] and [12], our focus is on the modeling of general IDS and its insights in IDS' cost-effectiveness. In addition, our incentive-based payoff functions more accurately represent the interactions between the IDS and attacker.

Decision theory is often employed to facilitate risk management and cost-benefit analysis [13] [14]. For example, Gaffney and Ulvila [13] used a decision analysis to evaluate and configure IDS. Decision theory assigns prior probabilities (usually fixed and exogenous) to handle the uncertainty of the environment (e.g., the prior probability of an intrusion). As pointed out in [12], decision theory based approaches are inadequate for IDS problems because they don't allow the defense system's decisions to influence the attacker's behavior. In contrast, game theory brings the attacker into the model and thus makes itself more attractive for handling the strategic interdependence between the IDS and attacker.

3 Review of Game Theory

A *game* is a formal representation of a situation in which a number of individuals with different objectives compete and interact with each other. In general, a game consists of the following elements:

- *Players*: who are involved?
- *Rules*: who moves when? What do they know when they move? What can they do (i.e., what strategies do they have)?

– *Outcomes*: for each possible set of actions by the players, what is the outcome of the game?
– *Payoffs*: what are the players' utility functions over the possible outcomes?

A classic example is the game of *Matching Pennies*. In this game, player 1 puts a penny down, either heads up or tails up. Player 2, not knowing player 1's choice, also puts a penny down. If the two pennies match (either both heads up or both tails up), player 1 pays 1 dollar to player 2. Otherwise, player 2 pays 1 dollar to player 1. The Matching Pennies game can be represented in the *extensive form* depicted in Figure 1. Due to its treelike structure, the extensive form is also known as a *game tree*. Matching Pennies is a two-player zero-sum game, in which case the sum of the payoffs is always zero. In general, however, most games of interest are non-zero-sum.

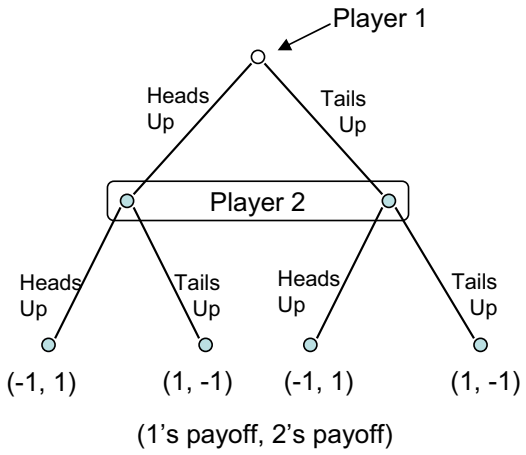

Fig. 1. Extensive form for Matching Pennies.

A central concept of game theory is the notion of a player's *strategy*. A strategy is a complete contingent plan, or decision rule, that specifies how the player will act in every possible circumstance in which he might be called upon to move. For example, in Matching Pennies, both players have two possible strategies: play heads (H) or tails (T). Then the game can be represented in terms of strategies and their associated payoffs. This representation, often depicted as a matrix, is known as the *normal* (or *strategic* form). The normal form of Matching Pennies is presented in Figure 2. Each row of the matrix represents a strategy of player 1, and each column a strategy of player 2. Within each cell, the first entry is player 1's payoff for the corresponding strategy profile; the second is player 2's.

Up to this point, we have assumed that players make their strategic decisions with certainty. However, a player could randomize when faced with a choice. We

	H	T
H	-1, 1	1, -1
T	1, -1	-1, 1

Fig. 2. Normal form of Matching Pennies.

now call the deterministic strategy a *pure strategy*. By contrast, a *mixed strategy* for a player is simply a probability distribution over his pure strategies. For example, a mixed strategy for player 1 in Matching Pennies is to play heads with the probability of 30%, and tails of 70%, instead of playing heads or tails all the time.

It is not so obvious to predict how each player should play Matching Pennies in order to maximize his own payoff. Consider the game illustrated in Figure 3 instead. In this game, player 1 has three pure strategies (U, M and D) and player 2 has two pure strategies (L and R). Note that, no matter how player 1 plays, R gives player 2 a strictly higher payoff than L does. In formal language, strategy L is *strictly dominated*. Thus, a "rational" player 2, who uses only those strategies that are best responses to some beliefs he might have about the strategies of his opponent, should not play L. Furthermore, if player 1 knows that player 2 will not play L, then U is a better choice than M or D. This process of elimination is called *iterated dominance*. It reduces the strategy sets of the players and thus simplifies the game.

	L	R
U	5, 1	6, 2
M	8, 4	3, 6
D	9, 6	2, 8

Fig. 3. Game example.

Unfortunately, most games (e.g., Matching Pennies) are not solvable by iterated dominance. The *Nash equilibrium* solution provides the optimal response to other players' strategies for each player. In a Nash equilibrium, none of the players has an incentive to deviate unilaterally from the equilibrium as long as the other players don't deviate. It can be proved that every finite game has a mixed strategy Nash Equilibrium. Solving the Nash Equilibrium for a 2×2 matrix game is trivial, although it can be costly for higher-dimensional matrix

games [4]. The Nash Equilibrium for Matching Pennies is the mixed strategy in which each player randomizes between his two pure strategies, assigning equal probability to each.

So far we have assumed that players know all relevant information about each other, including the structure of the game and payoffs that each receives. Such games are known as games of *complete information*. However, this assumption may be invalid in practice. How to weaken or entirely dispense with this assumption and solve games of *incomplete information* has been a challenging research topic in game theory. One widely used approach is to transform incomplete information about players into *imperfect* information about nature's moves. A game of this sort is known as a *Bayesian game*.

4 Game Theoretic Modeling

We use a two-person non-zero-sum game model to formulate the strategic interdependence between a general IDS and an attacker. The IDS can be host-based or network-based in an organizational environment [1]. The organization can range from small enterprises to government agencies. Intrusions are identified through anomaly detection, misuse detection or hybrid techniques [2]. The attacker can be a skillful intruder from outside, a malicious insider, or even a script kiddie.

Before we delve into the game modeling details, we shall introduce the parameters and identify the cost and payoff factors related to both players of the game.

4.1 Parameters and Cost/Payoff Factors

Table 1 summarizes the parameters used in our models. The parameters are always positive.

The cost and payoff factors associated with an IDS are:

- Operational cost (OC). This includes the cost of purchasing the IDS, the amount of time and computing resources needed to process the audit data stream and the personnel cost involved in administering and maintaining the IDS. Considering the voluminous audit data an IDS processes, the average operational cost for each audit event (the unit of analysis) should be nominal.
- Damage cost (DC), the amount of damage to the organization by an attack when IDS is not available or ineffective. It can be measured by the degradation of a set of security measurements associated with the organization's security metrics [10]. Different types of attacks can incur various levels of damage. Here we use DC to represent the expected damage cost by a generic attack.

[1] Our model can be easily extended to the case when there are multiple IDSs within the organization. The game formulation will remain the same.

Table 1. List of parameters.

notation	description
α	false alarm rate, the probability of an alarm, given no intrusion
β	false positive rate, the probability of no alarm, given an intrusion
δ	intrusion detection rate, $\delta = 1 - \beta$
ψ	probability of an attacker conducting an attack
ρ	probability of responding to an alarm
OC	IDS operational cost
DC	damage cost of an attack
RC	IDS response cost
RD	reward (i.e., payoff) of IDS for responding to an attack
LC	attacker's cost of launching an attack
PC	penalty to the attacker when the attack is detected
BA	benefit to the attacker when the attack is undetected

- Response cost (RC), the cost of acting upon an IDS alarm. Depending on the type of response mechanisms being used, the response cost includes the computing resources for terminating the offending connection or session, the downtime needed to repair and patch the computer systems, the labor cost of the response team, and so on.
- RD, the reward to the organization for responding to an attack. It can be measured by the improvement of the organization's security metrics after the response to the attack. In other words, RD is the potential damage cost caused by the attack if it went unnoticed otherwise [2].

Similarly, from an attacker's point of view, the cost and payoff factors include:

- LC, the cost for an attacker to launch an attack. It is the amount of time and resources needed to conduct the attack, which includes searching for system vulnerabilities, designing malicious code to exploit the vulnerabilities, etc.. It is reasonable for a vigilant security officer to assume that the cost of attack LC is small.
- Penalty cost (PC).This characterizes the risk for the attacker to be traced-back and punished. Quantitatively, it is the product of the probability of the attacker being held accountable and the penalty to the attacker when he is caught.
- BA, the benefit to the attacker when the attack is undetected. We use the attacker's incentive, quantified as the organization's damage cost DC, to represent BA (i.e., $BA = DC$), although it may not be the benefit he receives directly from the attack [10].

[2] For organizations such as law enforcement agencies, there is additional value for catching the attackers.

In this section we assume the values of the parameters are common knowledge known to both players of the game. Section 5 shows how we can weaken this assumption and design the game of incomplete information. We further assume that OC and LC are much less than the other cost and payoff factors.

4.2 IDS vs Attacker

The extensive form of the game is illustrated in Figure 4. The attacker's strategy is either to attack or not to attack the targeted organization. Accordingly, the organization may simply choose to have or not to have an IDS to defend against the attacker. Due to the imperfect technical performance of the IDS, it is possible that the IDS can raise a false alarm when there is no attack or generate no alarm when an attack is occurring. The corresponding probabilities and payoffs are shown in the shaded area in Figure 4. in this game we assume the IDS operator responds to every alarm the IDS generates (i.e., $\rho = 1$).

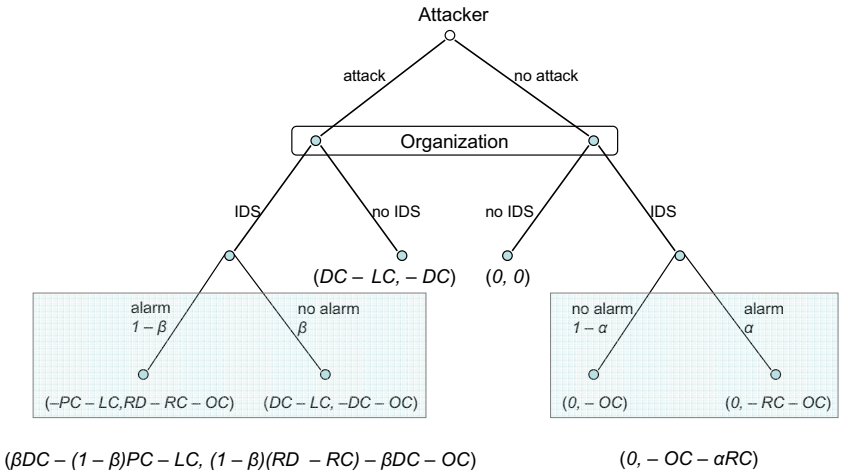

Fig. 4. Extensive form for Game A.

Figure 5 presents the normal form of the game. The payoffs for both players are determined as follows. If the attacker decides not to commit an attack, he receives no payoff. In contrast, the the organization has to pay for the operational cost (OC) along with the cost of false alarms (αRC) if an IDS is employed. When the attacker conducts an attack while the IDS is not in place or does not generate an alarm, the organization's loss is DC, whereas the attacker's payoff is the difference between BA and LC, which is the same as $DC - LC$. If the IDS successfully detects an attack, the organization gains RD at the cost of RC and OC. Meanwhile, the attacker bears the expected penalty cost PC in addition to

LC [3]. The organization's expected payoffs when having an IDS is determined by taking the sum of products of probabilities and payoffs for two scenarios (alarm or no alarm).

	IDS	No IDS
Attack	$A_{11} = \beta DC - (1-\beta)PC - LC$, $B_{11} = (1-\beta)(RD - RC) - \beta DC - OC$	$A_{12} = DC - LC$, $B_{12} = -DC$
No Attack	$A_{21} = 0$, $B_{21} = -OC - \alpha RC$	$A_{22} = 0$, $B_{22} = 0$

Fig. 5. Normal form of Game A.

We shall consider the attacker's strategies first. As shown in Figure 5, the attacker's payoff $A_{12} = DC - LC$ is greater than $A_{22} = 0$, based on our assumption that $LC \ll DC$. "Attack" would be the attacker's dominant strategy if A_{11} is also greater than $A_{21} = 0$. That is,

$$\begin{aligned} A_{11} &= \beta DC - (1-\beta)PC - LC \\ &\simeq \beta DC - (1-\beta)PC \\ &> 0 \end{aligned}$$

This is equivalent to

$$\frac{DC}{PC} > \frac{1-\beta}{\beta},$$

or

$$\frac{DC}{PC} > \frac{\delta}{1-\delta}.$$

We define $\lambda = DC/PC$, which essentially represents the benefit-to-risk ratio for the attacker.

Remark 1. If the attacker's benefit-to-risk ratio λ is greater than $\delta/(1-\delta)$, "attack" is his dominant strategy. In other words, "no attack" is the strictly dominated strategy. A rational attacker would always choose to attack the organization and thus maximize his payoff regardless of the organization's decision.

This result is not surprising. Intuitively, the greater potential damage cost or the lower risk would motivate the attacker more to commit the attack. On the other hand, the higher the attack detection rate, the more risk to the attacker and the less likely the attack occurring. This is illustrated in Figure 6. For a fixed λ value, an IDS with the intrusion detection rate less than $\delta_{threshold} = \lambda/(1+\lambda)$ ($\lambda > \frac{\delta}{1-\delta}$ is equivalent to $\delta < \frac{\lambda}{1+\lambda}$) would not play a deterrent role for the attacker at all. For example, when $\lambda = 10$, $\delta_{threshold} = 0.91$. Even when $\lambda = 5$,

[3] Our model can be extended to accommodate the case in which the attack is partially in progress when the IDS raises an alarm and the organization only recovers a portion of the damage.

$\delta_{threshold}$ is still as high as 0.83. On the other hand, for a fixed δ value, a malicious attacker who wants to maximize his payoff would rather conduct an attack with $\lambda > \delta/(1-\delta)$ than do nothing! This implies that the effective way to discourage the occurrence of an attack is to not only improve the attack detection rate but also increase the punishment for the attackers. In addition, it is interesting to note that false alarms cost nothing to the attacker. Therefore α does not come into play when determining the attacker's dominant strategy.

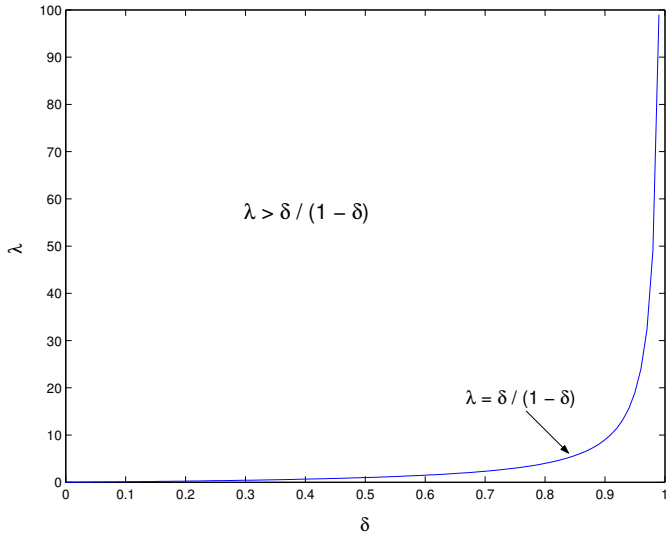

Fig. 6. Attacker's dominant strategy is "attack" when $\lambda > \delta/(1-\delta)$.

We next examine the organization's strategies. The negative B_{21} is always less than B_{22}. If B_{11} is also less than B_{12}, the organization's dominant strategy would be "no IDS". Therefore we have

$$\begin{aligned} B_{12} - B_{11} &= -DC - (1-\beta)(RD - RC) + \beta DC + OC \\ &\simeq (1-\beta)(RC - DC - RD) \\ &> 0 \end{aligned}$$

Clearly, B_{12} is greater than B_{11} if and only if $RC - DC - RD > 0$ (OC is relatively small and neglected and β is always less than 1).

Remark 2. Having an IDS becomes counterproductive when $RC > DC + RD$. It is more cost-effective for the organization not to use the IDS due to its high response cost.

Note that RC is compared to the sum of DC and RD (instead of DC alone). As long as $RC \leq DC + RD$, it is beneficial for the organization to employ an IDS. A similar analysis of the IDS' decisions upon an alarm (respond or not respond) would reveal the following:

Remark 3. When an IDS is deployed and it generates an alarm, it is more cost-effective to ignore the alarm and not to respond to it if $RC > DC + RD$.

In case there is no dominant pure strategy for neither player, we need to examine the Nash Equilibrium of the game. Let ψ and $1 - \psi$ be the probabilities for attacker's strategies "attack" and "no attack", respectively. Also let q and $1 - q$ be the probabilities for strategies "having IDS" and "no IDS" of the organization. In practice, an probability distribution over an organization's pure strategies can be interpreted as the extent to which the organization needs an IDS or the percentage of the time that the IDS should be available.

The Nash Equilibrium solution of the game is as follows:

$$\psi^* = \frac{\alpha RC}{\delta(DC + RD - RC) + \alpha RC},$$

$$q^* = \frac{DC}{\delta(PC + DC)} = \frac{\lambda}{\delta(1 + \lambda)}.$$

The organization's mixed strategy $(q^*, 1 - q^*)$ is the best response to the attacker's strategies. In fact, if the IDS is available with the probability of q^*, the attacker's expected payoff will be 0, whether he attacks or not. Similarly, if the probability for the attacker to attack the organization is ψ^*, the organization's expected cost is

$$V^* = -\frac{\alpha DC RC}{\delta(DC + RD - RC) + \alpha RC},$$

whatever its defense strategy is.

Remark 4. When an IDS is useful, the organization's best response to the attacker's strategies is to employ the IDS with the probability of $\lambda/[\delta(1+\lambda)]$. This way, the attacker's expected payoff will be 0, whatever he does.

It may seem counter-intuitive that q^* is proportional to $1/\delta$ and it increases monotonically with increasing λ. However, δ is expected to have a value greater than $\delta_{threshold} = \lambda/(1 + \lambda)$. The higher benefit-to-risk ratio for the attacker, the better the IDS ought to be in terms of intrusion detection rate, and the more it is inclined for the organization to have the IDS in order to catch the attacker and reduce his payoff. On the other hand, ψ^* increases with increasing false alarm cost of the IDS or decreasing δ, which implies an attack is more likely to happen if the IDS is less accurate.

Alternatively, q can be interpreted as the probability of responding to an alarm when an IDS is employed (i.e., ρ). Therefore the optimal value of ρ is q^*.

Remark 5. The optimal probability of responding to an IDS alarm is $\lambda/[\delta(1+\lambda)]$.

Table 2 lists a set of numerical examples of λ, δ, $\delta_{threshold}$ and q^* values.

Table 2. Numerical examples.

λ	δ	$\delta_{threshold} = \lambda/(1+\lambda)$	$q^* = \lambda/[\delta(1+\lambda)]$
0.1	80%	9.1%	11.4%
1	80%	50%	62.5%
5	85%	83.3%	98.0%
10	95%	90.9%	95.7%

5 Discussion

The game theoretic methodology of cost-benefit analysis is not limited to IDS. In fact, it can be easily extended to any security mechanism. It is important to bring adversarial attackers into the security models. The game theoretic formulation makes it possible to understand an attacker's intent and strategies from the attacker's perspective, which has been an often neglected facet of computer security research.

Our game model assumes complete information of the IDS and attackers. However, this assumption is somewhat unrealistic in practice. In particular, it is very difficult to estimate an attacker's payoff values. Significant research effort is needed to address the issue of accurate quantification of the attacker's payoff functions. Nevertheless, a qualitative game theoretic analysis can still bring unique and valuable insights to the understanding of the attacker's behavior and the decision making of security systems. Meanwhile, Bayesian games can be used to accommodate the uncertainty of the payoffs and handle the case of incomplete information. For instance, an attacker can be classified into three types: a skillful intruder from outside, a malicious insider, or a script kiddie. A prior probability is assigned to each type and the payoff values are identified for each sub-game associated with each type. Then the transformed game can be analyzed with standard techniques. Similarly, different attack types, such as denial of services, port scanning, masquerading, and so on, can be incorporated into the game model as well.

Finally, the interaction between attackers and security systems can be viewed as a repeated game. Both players can improve with the experience of playing the repeated games. How to incorporate game theory with learning is another important issue for our future work.

6 Conclusion

Cost-effectiveness is an important aspect of intrusion detection systems. In this paper, we have proposed a game theoretic methodology for cost-benefit analysis and design of IDS. A simple two-person game was used to model the strategic interdependence between a general IDS and an attacker. The solutions based on

the game theoretic analysis naturally integrate the cost-effectiveness and technical performance tradeoff of the IDS and provide valuable insights in intrusion detection and response.

Our game theoretic methodology can be applied to the cost-benefit analysis of any security mechanism. The main difficulty is the quantification of the players' payoffs. Our future work includes Bayesian game modeling of security systems and learning of repeated games.

7 Acknowledgment

This work is supported by the AFOSR grant F49620-01-1-0327 to the Center for Digital Security of the University of California, Davis.

References

1. Gartner Group: Hype Cycle for Information Security, 2003. http://www.gartner.com/pages/story.php.id.8789.s.8.jsp
2. J. McHugh, "Intrusion and Intrusion Detection," *International Journal of Information Security*, (1) 14–35, 2001.
3. W. Lee, W. Fan, M. Miller, S.J. Stolfo and E. Zadok, "Toward Cost-sensitive Modeling for Intrusion Detection and Response," *Journal of Computer Security*, (10) 5–22, 2002.
4. D. Fudenberg and J. Tirole, *Game Theory*, MIT, 1991.
5. K. Matsuura, "Information Security and Economics in Computer Networks: An Interdisciplinary Survey and a Proposal of Integrated Optimization of Investment," The 9th International Conference of Computing in Economics and Finance (CEF 2003), Seattle, July 2003.
6. The Third Annual Workshop on Economics and Information Security (WEIS04). http://www.dtc.umn.edu/weis2004/
7. L. Gordon and M. Loeb, "The Economics of Information Security Investment," *ACM Transactions on Information and System Security*, 5(4): 438-457, 2002.
8. C. Iheagwara, "The Effect of Intrusion Detection Management Methods on the Return on Investment," in press, *Computers & Security*, 2004.
9. K. Lye and J. Wing, "Game Strategies in Network Security," in Proc. of 15th IEEE Workshop on Foundations of Computer Security (FCS '02), Copenhagen, Denmark, July 25-26, 2002.
10. P. Liu and W. Zang, "Incentive-Based Modeling and Inference of Attacker Intent, Objectives, and Strategies," in Proc. of ACM Conference on Computer and Communications Security (CCS' 03), 179-189, 2003.
11. T. Alpcan and T. Basar, "A Game Theoretic Approach to Decision and Analysis in Network Intrusion Detection," in Proc. of 42nd IEEE Conference on Decision and Control, Maui, HI, December 2003.
12. H. Cavusoglu and S. Raghunathan, "Configuration of Intrusion Detection Systems: A Comparison of Decision and Game Theoretic Approaches," in Proc. of International Conference on Information Systems (ICIS), Seattle, December 2003.
13. J. Gaffney and J. Ulvila, "Evaluation of intrusion detectors: A decision theory approach," in Proc. of 2001 IEEE Symposium on Security and Privacy. IEEE Computer Society, Los Alamitos, CA, 50–61, 2001.

14. I. Balepin, S. Maltsev, J. Rowe and K. Levitt, "Using Specification-Based Intrusion Detection for Automated Response," in Proc. of the 6th International Symposium, RAID 2003, Recent Advances in Intrusion Detection, Pittsburgh, PA, September 8-10, 2003.