Mission-oriented Security Model, Incorporating Security Risk, Cost and Payout

Sayed M. Saghaian N. E., Tom La Porta, Trent Jaeger,
Z. Berkay Celik, and Patrick McDaniel

Department of Computer Science and Engineering,
The Pennsylvania State University
{sms676,tlp,tjaeger,zbc102,mcdaniel}@cse.psu.edu

Abstract. One of the most difficult challenges facing network operators is to estimate risk and allocate resources in adversarial environments. Failure to properly allocate resources leads to failed activities, poor utilization, and insecure environments. In this paper, we explore an optimization-based approach to allocating resources called a *mission-oriented security model*. This model integrates security risk, cost and payout metrics to optimally allocate constrained secure resources to discrete actions called missions. We model this operation as a Mixed Integer Linear Program (MILP) which can be solved efficiently by different optimization solvers such as MATLAB MILP solver, IBM-CPLEX optimizer or CVX solver. We further introduce and explore a novel method to evaluate security risk in resource planning using two datasets—the Ponemon Institute cost of breach survey and CSI/FBI surveys of security events. Data driven simulations are used to validate the model robustness and uncover a number of insights on the importance of risk valuation in resource allocation.

1 Introduction

Operators of modern networks must allocate resources such as computation, bandwidth, storage capacity, or personnel to achieve operational goals in the face of adversaries. Consider the deployment of a file system service within a LAN–one could simply deploy an unauthenticated server, or use industry grade cryptography, multifactor authentication, multiple backups, and log every packet and filesystem event. These two deployments represent points in the spectrum in the cost/security/reliability space. In the absence of context, both deployments are equally appropriate. The key to secure operations is to make such a decision by understanding the needs and risks of the environment.

Indeed, today's operators make decisions about what is appropriate for an environment simply from intuition and experience—and often unconsciously assess and weigh the risks of the environment [3]. In this work, we seek to formalize this decision-making process. We develop a mathematically rigorous decision-making model to explore different policies by evaluating their effectiveness when dealing with different risk characteristics. Using this model, we explore the interplay between utilization, risk, and security and develop new insights on how to allocate resources in the face of adversaries with varying goals and strategies.

After reviewing operational best practices and work in decision theory, we have identified several essential elements of an operationally aware decision process: *risk*, *cost*, *payout* and *missions*. We provide intuitive definitions here (see Section 3 for formal definitions). Intuitively, the risk is a valuation of the harms that may occur, the damage they would cause, and the probability that the harms will occur. We refer to cost as the *security* cost. The cost is the required amount to spend to achieve some level of security. The payout is the value of performing an action (e.g., the "value" of an action to the environment)–which enables the decision process to prioritize actions by attempting to maximize profit.

Lastly, we refer to a mission as a series of actions that leads to an objective. A mission is defined by intent as well as how it is executed. With mission-oriented security, we consider the overall security of a system, not a single algorithm. Based on the risk level of a mission and the potential damage due to the risk, a mission might be allowed to continue even though there is a chance of being exposed to attacks or being compromised. When referring to a mission-oriented security model, we are referring to a mathematical model to make such decisions.

Note that often when considering the security of a system, people instinctively behave in a risk-averse fashion. However risk-averse approaches may sacrifice the total profit of the system, particularly when the probability of undesirable event occurrence or the damage due to these events are *relatively* small. On the other hand, a system that does not account or prepare for risk may suffer serious consequences, particularly when the probability of undesirable events is relatively large, or when the damage to the infrastructure or outcomes is extensive. Trading the contextual risk against the payout is a key to making good (optimal) decisions on resource allocations. We explore how this can be formally modeled throughout.

In general, there can be two different types of decision processes: deployment and operational. As an example of a deployment decision, suppose we have three missions in an enterprise: providing email service, telnet access, and wireless access. Assume all these services are at risk for various attacks. Remediation measures might include adding two-factor authentication for email service, adding a firewall and VPN for telnet access, or adding a VPN for wireless access. However, due to the limited number of available staff and infrastructure, we can only perform two out of the three remediations. The dilemma is that given attack characteristics corresponding to each of the missions, which of these missions should be assigned to our limited resource (staff). This kind of configuration or deployment decision is closely related to security planning.

As an example of an operational decision, assume a server in networked environment notes growing evidence of attacks (e.g., an adversary has found the IP address behind the firewall of a server in our server farm). Furthermore, there is a set of servers in our system which is more secure but slower. However, the cost of migrating processes to the secure server and the limited number of secure servers prevent us from moving all processes to those servers. The dilemma is which processes to migrate to the secure server, and if a process is not migrated, should we keep running that process or terminate it.

In the first part of this work, we define a cyber mission-oriented security model. We incorporate security risk, cost, and payout into our model, and consider a resource allocation problem where the objective is to jointly optimally allocate a secure constrained resource to missions and to decide whether to stop missions that do not receive the secure resources.

This formulation is the first step toward a larger body of analyses. To make the problem mathematically tractable, we explore three simple agility maneuvers: (i) assigning a mission to specially secured resources, (ii) continuing a mission, or (iii) stopping a mission. Our study can be extended to include more agility maneuvers such as reconfiguration, suspending a mission, etc.

We model the resource allocation problem as a Mixed Integer Nonlinear Program (MINLP) [6] where the objective function is to maximize the total profit gained from all missions. MINLP is a class of Nonlinear Programming which consists of both integer and continuous variables where the objective function or the constraints contain nonlinear terms. Our nonlinear programming problem can be linearized and transformed into a Mixed Integer Linear Program (MILP) by exploiting McCormick envelopes. Typically, a branch and bound algorithm is used to solve MILPs. Optimization solvers such as MATLAB MILP solver, IBM-CPLEX optimizer, or CVX commercial solver can efficiently provide a solution to a MILP. In this paper, we adopt MATLAB MILP solver to obtain an optimal resource allocation strategy. The main contributions of this paper are:

1. We introduce a mission-oriented security model that integrates security risk, cost and payout metrics, and develop a mathematical framework to optimally allocate constrained secure resources to missions.
2. We evaluate three different policies for allocating the constrained resources in facing with different risk profiles by using two datasets—the Ponemon Institute cost of breach survey and CSI/FBI survey of security events.
3. We investigate the sensitivity of our proposed framework with respect to under/over-estimating the probability of undesirable events.

We begin in the following section with a review of several key related work.

2 Related work

In the sensor network domain, [11] proposed a framework to optimally allocate constrained resources (sensors) to missions such that the total profit is maximized when there exist uncertainties in users demands and their achievable profits. They assumed missions are always profitable and hence, should always be continued. However, in this work, missions might not always be profitable due to risks, and hence, we might have to terminate them. Furthermore, missions payouts and damages due to risks depend on if and how much of their request to use secure special resources is satisfied.

One challenge in today's development of security technology is misaligned incentives of different parties. For example, potential failure or security breach rises when a person or a software guarding a system faces a lower failure or compromise cost. The goal in information security economics [1, 2] is to combine concepts from game theory, microeconomic theory and risk assessment with cryptography concepts to develop security technology. To protect a given set of information, [7] presented an economic model to derive the optimal amount to invest in security.

Unlike dependability and reliability of a computer system, there is not much work on quantitatively evaluating the security of a computer system. To evaluate the security of a system quantitatively, [16] surveyed model-based methods for evaluating dependability of a computer system. Furthermore, they discuss extending these methods to evaluate the security of a computer system.

A risk assessment method is a *qualitative* approach in which the likelihood of occurrence of undesirable events or their impacts are described qualitatively using terms like low, medium or high. Conversely, when the likelihood of occurrence of undesirable events and their consequences are expressed numerically, the risk evaluation method is called *quantitative* risk assessment [4]. Given a qualitative assessment of risk and complying with the ISO/IEC 27005 standard on risk management [12], European Network and Information Security Agency (ENISA) ranks different risks using the likelihood of a threat times its impact by assuming a scale from 1 to 5 (where 1 represents very low, whereas 5 denotes very high) for each of the likelihoods and impacts [5]. Clearly, the scaling values are relative values, and it is not clear how the resulted quantities for risk can be combined with other parameters such as cost or payouts in these methods.

Note that this work is highly related to the field of systems resilience [13] and agility [3, 15]. Agility refers to an approach to achieving system resilience in which a defender reconfigure the system or the operation in response to a potential attack or perceived risk. In this work, we focus on formulations of the decision problem to reconfigure in response to various risks and threats.

3 Problem Statement

Figure 1 illustrates our mission-oriented security model. Missions are running on their own local server. However, these servers are under risk because of various potential cyber attacks. To mitigate these risks, the system may assign some of these missions to special secure resources that are immune from cyber attacks. In the beginning of each time frame, the set of available missions are competing for scarce secure resources. We want to determine how to allocate resources and whether to stop or continue the missions for the current time frame. In the next time frame, previously stopped missions with potentially different risk profiles will compete with a set of newly arrived missions. We do not allow preemption.

Informally, the problem is to determine which missions are assigned to the secure resources, to determine the amount of special resources assigned to the selected missions, and to decide whether to continue or stop the missions. In

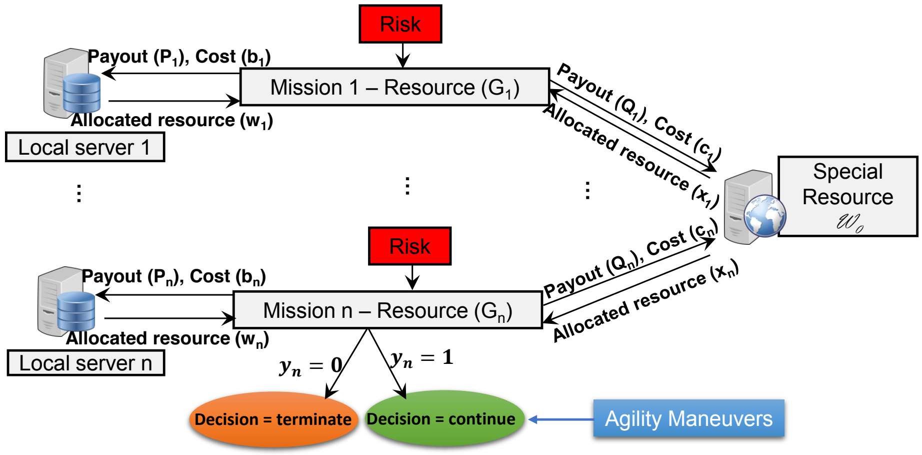
Risk
Payout (P$_1$), Cost (b$_1$)
Mission 1 – Resource (G$_1$)
Allocated resource (w$_1$)
Local server 1
Payout (Q$_1$), Cost (c$_1$)
Allocated resource (x$_1$)
Special Resource $\mathscr{W}_0$
Risk
Payout (P$_n$), Cost (b$_n$)
Mission n – Resource (G$_n$)
Allocated resource (w$_n$)
Local server n
Payout (Q$_n$), Cost (c$_n$)
Allocated resource (x$_n$)
$y_n = 0$
$y_n = 1$
Decision = terminate
Decision = continue
Agility Maneuvers

Fig. 1: Mission-oriented security model

other words, we find the optimal strategy that maximizes the total profit achieved from the execution of all the missions.

Formally, assume there are n missions $M_1, \ldots, M_n$ where for each mission, M_i requires G_i amount of resource. Initially, each mission is running on its own local server which has sufficient resources to support the mission. In a risk-free environment, M_i achieves payout P_i from its local resource/server if the local server fully allocates the entire requested amount of resources (G_i) to the mission. Further, the unit cost of allocating resources to M_i by its local server is b_i.

Now due to security risks, each mission will request to use a global special resource/server that has only $\mathscr{W}_0$ amount of a special (secure) resource. Each mission M_i is exposed to l different types of attack; each attack type happens with probability of α^i_j (where superscript i denotes the i-th mission, while subscript j indicates the j-th attack type corresponding to M_i) and results in damage (loss of profit) of d^i_j. These types of attack are independent. They are potentially derived from different probability distributions.

For each mission M_i, risk is modeled as a set of triples [4]:

$$R_i = \{(s^i_1, \alpha^i_1, d^i_1), \ldots, (s^i_l, \alpha^i_l, d^i_l)\} \tag{1}$$

where s^i_j is the j-th type of attack corresponding to M_i, α^i_j is the probability that this attack type happens, and d^i_j is the damage due to that event.

Let c_i be the unit cost of allocating the special resources to the i-th mission, M_i. If mission M_i is fully allocated G_i amount of the special resource (mission gets all the resources it needs), it achieves payout Q_i with no loss due to risk. If we don't allocate any special resource to mission M_i, then mission M_i is exposed to the entire potential damage as a result of risk. If a portion of the required special resource is assigned to mission M_i, only the remaining portion is exposed to risk.

Table 1: Symbols and parameters used in our resource allocation framework.

Symbols	**Descriptions**
M_i	i-th mission
s_j^i	j-th type of attack corresponding to M_i
α_j^i	probability that the event s_j^i happens
d_j^i	the damage as a result of occurrence of s_j^i
$\mathscr{W}_0$	amount of available special resource
G_i	total requested resources by mission M_i
b_i	unit cost of allocating local resources to M_i by its local server
c_i	unit cost of allocating the special resource to mission M_i
P_i	achieved payout by M_i from its local resource if the local server allocates G_i to M_i, when there is no risk to mission M_i
Q_i	achieved payout by M_i from the special resource if it is assigned G_i amount of special resource
w_i	allocated resources from local resource to M_i
x_i	allocated resources from special resource to M_i
y_i	$y_i = 0$ indicates to stop mission M_i $y_i = 1$ indicates to continue mission M_i

Table 1 provides the description of symbols used in the proposed framework for optimally allocating constrained secure resources to missions.

We assume payouts (P_i's and Q_i's) and potential damage are proportional (linear) to the amount of allocated resources. Formally, let w_i and x_i be the amount of allocated resources from the local resources and the special secure resources to mission M_i, respectively. Then, we define:

Partial payout from local server = $\frac{w_i}{G_i} P_i$,
Partial payout from special server = $\frac{x_i}{G_i} Q_i$,
Partial potential damage = $\frac{G_i - x_i}{G_i} \alpha_j^i d_j^i$,
where $w_i, x_i \leq G_i$.

We only have $\mathscr{W}_0$ amount of the special resource, and we would like to optimally allocate it amongst multiple missions in response to risk. We formalize the problem as follows:

1. Allocate resources (w_i's, x_i's) to missions optimally such that the total profit (payout minus cost minus damage due to risk) from all missions is maximized.
2. Decide to stop the missions ($y_i = 0$) or let them continue ($y_i = 1$):
   (a) if a mission stops, then it gains no profit and no damage.
   (b) if a mission continues, then it achieves partial payouts minus partial cost minus partial potential damage from risk.

4 Problem Formulation

We define profit as payout minus cost minus expected potential damage due to risk. We model the problem as a Mixed Integer Nonlinear Program (MINLP), where the objective function is to maximize the total profit gained from all the missions. Using our assumption regarding the linearity (in respect to the allocated resource) of partial payouts and damage, and by using techniques in Nonlinear Programming, we linearize our nonlinear programming problem and transform the problem into a Mixed Integer Linear Program (MILP). An optimal solution to the MILP can be found by using a branch and bound algorithm.

In Section 4.1, we construct a MINLP that formulates the problem. Using McCormick envelopes, we then linearize this MINLP in Section 4.2. The resulted MILP can be solved by using different solvers such as MATLAB MILP solver.

4.1 Objective Function

We consider a binary decision variable y_i for each mission M_i; $y_i = 0$ indicates the mission is terminated, and the mission receives no profit and no damage. Whereas $y_i = 1$ indicates the mission M_i is continued. Further, let w_i and x_i be the amount of the local and the special resource allocated to the i-th mission M_i, respectively. Then, the optimization problem is formulated by (2):

$$\begin{aligned}
\operatorname*{Min}_{w,x,y} \quad & \sum_{i=1}^{n} \overbrace{(b_i w_i + c_i x_i)}^{\text{cost}} \quad - \\
& \sum_{i=1}^{n} y_i \left[\underbrace{\left(\frac{w_i}{G_i} P_i + \frac{x_i}{G_i} Q_i \right)}_{\text{payout}} - \sum_{j=1}^{l} \quad \underbrace{\frac{G_i - x_i}{G_i} \alpha_j^i d_j^i}_{\text{potential risk damage}} \right] \qquad (2) \\
\text{s.t.} \quad & \sum_{i=1}^{n} x_i \leq \mathscr{W}_0 \\
& w_i + x_i \leq G_i \qquad \text{for} \quad i = 1, \ldots, n \\
& w_i, x_i \geq 0 \qquad \text{for} \quad i = 1, \ldots, n \\
& y_i \in \{0, 1\} \qquad \text{for} \quad i = 1, \ldots, n
\end{aligned}$$

By construction and inherent nature of MINLP, this formulation returns a strategy that leads to the highest expected total profit.

4.2 Solving The Minimization Problem

We linearize the objective function in (2) by exploiting *McCormick envelopes*, and reformulate the problem as a *Mixed Integer Linear* Problem (MILP). McCormick envelopes are a convex relaxation technique that can be exploited to linearize MINLPs. In particular, if the objective function is a bilinear function, applying McCormick envelopes linearizes the objective function. The function $f(x, y)$ is a bilinear function if it is linear with respect to each variable x and y individually.

Let $f(x, y) = yg(x)$, where y is a binary variable and $g(x)$ is a linear function with lower bound of L and upper bound of U, i.e. $L \leq g(x) \leq U$. Then, one can linearize $f(x, y)$ by defining a new variable z and exploiting McCormick envelopes. The function $f(x, y)$ is equivalent to:

$$z$$

$$\text{s.t.} \qquad Ly \leq z \leq Uy \tag{3}$$

$$g(x) + U(y - 1) \leq z \leq g(x) + L(y - 1) \tag{4}$$

We note that if $y = 0$, then $f(x, y) = 0$. Furthermore, from (3), z must be 0 too. Moreover, if $y = 1$, then $f(x, y) = g(x)$. On the other hand, from (4), z must be $g(x)$ as well. Hence, $f(x, y)$ is equivalent to z with constraints (3) and (4).

By defining $D_i := \sum_{j=1}^{l} \alpha_j^i d_j^i$, $B_i := \frac{Q_i}{G_i} + \frac{1}{G_i} D_i$ and $A_i := \frac{P_i}{G_i}$, and by using McCormick envelopes, we can linearize (2) and obtain the following MILP:

$$\begin{aligned}
\underset{w,x,y,z}{\text{Min}} \quad & \sum_{i=1}^{n} (b_i w_i + c_i x_i - z_i + D_i y_i) \\
\text{s.t.} \quad & \sum_{i=1}^{n} x_i \leq \mathscr{W}_0 \\
& \text{for} \quad i = 1, \ldots, n: \\
& w_i + x_i \leq G_i \\
& y_i \in \{0, 1\} \\
& w_i, x_i, z_i \geq 0 \\
& z_i \leq (P_i + Q_i + D_i) y_i \\
& z_i \geq A_i w_i + B_i x_i + (P_i + Q_i + D_i)(y_i - 1) \\
& z_i \leq A_i w_i + B_i x_i
\end{aligned} \tag{5}$$

To solve MILP, a branch and bound algorithm is typically used in different solvers such as MATLAB MILP solver, IBM-CPLEX optimizer, or CVX commercial solver. In this paper, we use MATLAB MILP solver.

To deploy our proposed decision making under security risk in practice, a method to assess risk quantitatively is required. We next formally define security risk and propose a novel quantitative risk assessment method.

5 Risk Evaluation

In this section, we first briefly review the definition of security risk and the expected damage from risk (Section 5.1). We then construct the elements of our proposed novel probabilistic risk assessment method based on the relative probability of the j-th attack type (Section 5.2.1) and the probability of experiencing a successful attack (Section 5.2.2).

5.1 Risk Definition

Probabilistic risk assessment [4] is an analysis that measures and evaluates risk systematically. It estimates the consequences of undesirable events and predicts the likelihood of such events. These approaches often use expert opinion or historical data to assess the likelihood of undesirable events and their consequences. This method of risk assessment aims to address three questions [14]:

1. What vulnerabilities are exposed? What are undesirable events? What are types of attacks experienced by a mission, system or enterprise?
2. What is the probability of occurrence of those events?
3. What are the consequences or potential impact of occurrence of those undesirable events? What is the damage due to those events?

Historical data can be used as test data to evaluate security systems as we discuss later. Furthermore, they form the basis of our proposed risk assessment framework. One might argue that historical data cannot be a valid sampling set to estimate future probabilities and consequences as they only refer to series of known past events. However, as reference [10] points out, although studies by the Computer Emergency Response Team Coordination Center (CERT/CC) show that risks do change with technological advances and human factors, the changes are small and infrequent, i.e., risk is largely stable over time.

Discussed previously, we model the risk of each M_i as a set of triples in (1). In a probabilistic risk assessment approach, risk is related to two parameters. The first parameter is the likelihood of occurrence of undesirable events such as experiencing cyber attacks, or vulnerability exposure. The other parameter is the consequence of the occurrence of these events.

Following ISO/IEC 27005 [12], we define the expected loss due to the event s_j^i, the j-th type of attack corresponding to mission M_i as: $\mathscr{L}_j^i := \alpha_j^i d_j^i$. However, there is very limited statistical information from studies on attacks experienced by enterprises. As a result, it is very difficult to quantify α_j^i's in practice. To be able to quantify risk, we define potential loss due to the j-th type of attack corresponding to mission M_i as [21]:

$$\mathscr{L}_j^i := p_\tau^i p_j^i d_j^i \tag{6}$$

where p_τ^i is the probability that mission M_i experiences a successful attack during time interval τ and p_j^i is the *relative* probability of occurrence of the j-th type of attack, or threat. Table 2 provides the description of symbols used in the proposed risk assessment method.

The relative probability of the j-th attack type is defined as the likelihood that a mission is under the j-th type of attack given an intense assumption that the mission is in fact under an attack.

From (6), we can compute the total expected damage from risks for each of the missions from: $\mathscr{L}^i := \sum_{j=1}^{l_i} p_\tau^i p_j^i d_j^i$, where l_i is the number of attack types corresponding to the i-th mission, M_i.

Table 2: Symbols and parameters used in the risk assessment method.

Symbols	Descriptions
p_{τ}^{i}	probability that M_i experiences a successful attack during time interval τ when no secure resource is assigned to this mission
p_j^i	relative probability of occurrence of the j-th attack
l_i	number of attack types corresponding to M_i
$\mathscr{L}_j^i$	potential loss due to the event s_j^i
$\mathscr{L}^i$	total expected damage from risk for mission M_i
$\mathscr{L}_{LogN}^i$	total expected damage from risk for mission M_i when TBC is modeled by lognormal distribution

5.2 Risk Assessment

To calculate the expected loss due to risk, we find the relative probabilities of attack types (Section 5.2.1) and the probability that missions experience a successful attack during a time interval (Section 5.2.2).

5.2.1 Computing Relative Probabilities The relative probability of the j-th attack type can be computed based on available historical data. Given historical data from a previous time period, the relative probability of the j-th attack type corresponding to mission M_i is computed from:

$$p_j^i = \frac{perc_j^i}{\sum_{k=1}^{l_i} perc_k^i} \tag{7}$$

where $perc_k^i$ is the percentage of times that mission M_i experienced attack type k during the last time period. Alternatively, depending on the available data set, $perc_k^i$ can be defined as the percentage of enterprises that experienced the k-th type of attack during the last time period. Moreover, in case that an enterprise is supplied with its security expert's opinion rather than past data, $perc_k^i$ can be expressed based on the expert opinion, and be defined as the percentage of experts that think the enterprise will experience the attack type k.

5.2.2 Computing Probability of Experiencing an Attack The probability of experiencing a successful attack during time interval τ can be modeled by different probability distributions. Let random variable X represent the time between two consecutive attacks. Further, let $q^i(\tau)$ denote the probability that no successful attack has occurred in the time interval τ for mission M_i. Then:

$$q^i(\tau) = \Pr(X > \tau)$$
$$p_{\tau}^i = 1 - q^i(\tau) = F_X(\tau) \tag{8}$$

where F_X is the Cumulative Distribution Function (CDF) corresponding to the random variable X [17].

Lognormal Distribution: Authors in [9] conducted a large-scale study on the required time to compromise a computer system. According to their analysis on detected cyber intrusions on $260,000$ computer systems over a period of three years, they find that lognormal distribution is the best fit to model the Time Between Compromises (TBC).

Assume X has a lognormal distribution with parameters μ^i and σ^i: $X \sim ln\mathscr{N}(\mu^i, \sigma^i)$. Then,

$$p_\tau^i = \frac{1}{2}\left[1 + erf\left(\frac{ln\ (\tau) - \mu^i}{\sigma^i\sqrt{2}}\right)\right] \tag{9}$$

where erf is the error function. The parameters of lognormal distribution (μ^i and σ^i) can be estimated from the historical data.

6 Data and Simulation Results

We evaluate the mission-oriented security model through an example of seven missions competing for constrained secure resources. We first calculate a numerical value for the probability of experiencing a successful attack during 30 days (Section 6.1). We then present datasets of the Ponemon Institute cost of breach survey and CSI/FBI surveys of security event that we exploit to characterize the risk of missions (Section 6.2). We evaluate the performance of our proposed model and compare it with the performance of two other policies (Section 6.3). We further investigate the effect of different factors on the performance of the three different approaches (Section 6.4). Finally, we provide a sensitivity analysis that studies the effect of incorrectly estimating the probability of occurrence of attacks on the performance of our method (Section 6.5).

6.1 Computing the probability of experiencing a successful attack on the local servers (p_τ)

We consider the probability of attack arrival over a month by setting the parameter τ in (9) to 30 days.

The study in [9] found that the Maximum Likelihood estimate of the lognormal parameters are $\hat{\mu} = 3.719$ and $\hat{\sigma} = 1.065$. Hence, assuming TBC has a lognormal distribution, from (9), the probability that M_i experiences a successful attack when its resource requirements are satisfied from its local server is computed as:

$$p_\tau^i = \frac{1}{2}\left[1 + erf\left(\frac{ln\ (30) - 3.719}{1.065\sqrt{2}}\right)\right] = 0.3827$$

Recall that mission M_i will not experience any successful attack if all of its resource requirements are satisfied from the secure server. If part of its resource requirements is satisfied from its local server, the damage due to a successful attack will only be proportional to the amount of resources allocated from the local server, i.e., only local server use induces damage.

6.2 Data and Experiments Settings for Ponemon Institute and CSI/FBI Surveys

We consider seven different missions with risk characteristics derived from the Ponemon Institute and CSI/FBI surveys, and conduct 100 iterations in which attack characteristics are simulated based on p_τ and risk vectors in Table (3). The special resource has only $W_0 = 3500$ units of resource while the amount of required resources (G_i's) and payouts (P_i's and Q_i's) are drawn uniformly at random from ranges: $G_i \in U([1000, 1500])$ (for all missions).
$P_1 \in U([80000, 100000])$, $P_2 \in U([30000, 40000])$, $P_3 = U([40000, 60000])$,
$P_4 \in U([40000, 60000])$, $P_5 \in U([20000, 22000])$, $P_6 \in U([50000, 52000])$,
and $P_7 \in U([11000, 13000])$.
$Q_1 \in U([50000, 70000])$, $Q_2 \in U([15000, 25000])$, $Q_3 \in U([30000, 40000])$,
$Q_4 \in U([30000, 40000])$, $Q_5 \in U([12000, 15000])$, $Q_6 \in U([30000, 35000])$,
and $Q_7 \in U([7000, 10000])$.
Further, the unit cost of using the special resources is $c_i = 2$ for each mission, while the unit cost of the local resources is $b_i = 1$ for each mission.

6.2.1 Ponemon Institute Study The Ponemon Institute conducted a survey [18] on the annual cost of cybercrime from 237 organizations in six different countries: United States (M_1), Japan (M_2), Germany (M_3), United Kingdom (M_4), Brazil (M_5) and Australia (M_6). We use their results to assess risk corresponding to six missions. We consider a scenario in which the data of each country contributes to the evaluation of risk for the six different missions ($M_1, \dots, M_6$).

Eight different types of attack have been identified by the Ponemon Institute study. The eight types of attack categorized by Ponemon Institute and the percentage of users who experienced these attacks are: Malware (98%), Phishing and Social Engineering (70%), Web-based attacks (63%), Malicious code (61%), Botnets (55%), Stolen devices (50%), Denial of services (49%), and Malicious insiders (41%). Organizations were asked if they have experienced these types of attack. Using (7), we calculate the relative probability of each attack type (p^i_j's for $i = 1, \dots, 6$ and $j = 1, \dots, 8$). Furthermore, we use their results, and compute the monthly damage as a result of these types of attack for each mission. As an example, we can compute the relative probability of the Malware attacks (the first type of attack corresponding to missions $M_1, \dots, M_6$) from (7) as: $p_1 = \frac{0.98}{0.98+0.70+0.63+0.61+0.55+0.50+0.49+0.41} = 0.2012$. Furthermore, the damage due to Malware attacks for mission M_1 (United States) for the fiscal year of 2016 was reported to be 0.13x17.36 million = \$2, 256, 800. Therefore, on average, the monthly damage as a result of Malware attacks for mission M_1 was $d_1 = \frac{2{,}256{,}800}{12} = 188, 067$.

Table (3) summarizes risk parameters for each mission with the columns $M_1, \dots, M_6$. From this data set, Malware attacks were the most common types of attack experienced by users, while Malicious code had the most severe damage.

6.2.2 CSI/FBI Survey In a sequence of surveys [8, 19, 20] conducted yearly by the Computer Security Institute (CSI) and the Federal Bureau of Investigation

Table 3: Risk characteristic from Ponemon Institute and CSI/FBI surveys.

	Missions						
relative probabilities/damage	M_1	M_2	M_3	M_4	M_5	M_6	M_7
p_1^i	0.2012	0.2012	0.2012	0.2012	0.2012	0.2012	0.1005
d_1^i	188,067	139,833	124,133	78,108	79,050	57,333	118,919
p_2^i	0.1437	0.1437	0.1437	0.1437	0.1437	0.1437	0.1411
d_2^i	217,000	69,917	124,133	66,092	35,133	53,750	3,926
p_3^i	0.1294	0.1294	0.1294	0.1294	0.1294	0.1294	0.0024
d_3^i	173,600	146,825	111,067	102,142	87,833	46,583	29,375
p_4^i	0.1253	0.1253	0.1253	0.1253	0.1253	0.1253	0.0239
d_4^i	347,200	48,942	58,800	66,092	48,308	46,583	4,176
p_5^i	0.1129	0.1129	0.1129	0.1129	0.1129	0.1129	0.1077
d_5^i	43,400	34,958	13,067	18,025	8,783	7,167	2,605
p_6^i	0.1027	0.1027	0.1027	0.1027	0.1027	0.1027	0.1962
d_6^i	86,800	34,958	78,400	72,100	43,917	28,667	16,656
p_7^i	0.1006	0.1006	0.1006	0.1006	0.1006	0.1006	0.0359
d_7^i	231,467	90,892	104,533	138,192	79,050	68,083	27,383
p_8^i	0.0842	0.0842	0.0842	0.0842	0.0842	0.0842	0.1914
d_8^i	159,133	132,842	39,200	60,083	57,092	50,167	11,271
p_9^i	—	—	—	—	—	—	0.0861
d_9^i	—	—	—	—	—	—	4,684
p_{10}^i	—	—	—	—	—	—	0.0144
d_{10}^i	—	—	—	—	—	—	1,267
p_{11}^i	—	—	—	—	—	—	0.0502
d_{11}^i	—	—	—	—	—	—	17,877
p_{12}^i	—	—	—	—	—	—	0.0502
d_{12}^i	—	—	—	—	—	—	224,987
Expected damage ($\mathscr{L}_{LogN}^i$)	70987	34894	33638	28745	21674	17585	12238

(FBI), respondents from different industry organizations are asked the types of attack they experienced and the cost of those attacks to their organization. We use this dataset to evaluate security risk corresponding to another mission (M_7).

The latest CSI/FBI survey that we found which contains damage lost per attack type, per respondent is their 2003 survey [19]. 12 different types of attacks were identified by this survey. These attack types and percentage of respondents that experience each of these types of attack are: Denial of service (42%), Stolen devices (59%), Active wiretap (1%), Telecom fraud (10%), Unauthorized access by insiders (45%), Virus (82%), Financial fraud (15%), Insider abuse of Net access (80%), System Penetration (36%), Telecom Eavesdropping (6%), Sabotage (21%), and Theft of Proprietary Info (21%).

Likewise, we calculate the relative probability of each attack type (p_j^7's for $j = 1, \ldots, 12$) for M_7 using (7). For example, the relative probability corresponding to DoS (the first type of attack corresponding to missions M_7) is $p_1 = 0.1005$. Moreover, the damage due to DoS attack was on average $d_1 = \frac{1{,}427{,}028}{12} = 118,919$ monthly. Refer to column M_7 of Table (3) for risk vectors of this mission.

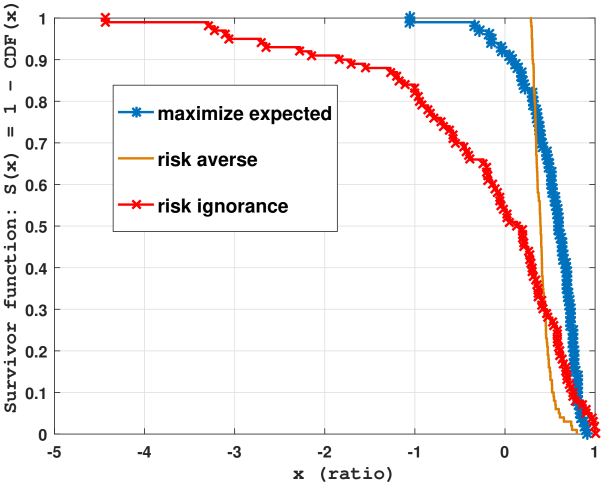

Fig. 2: Survivor functions of three policies

6.3 Performance Evaluation

6.3.1 Comparison of Different Policies: We compare our proposed MILP-based framework that maximizes the expected total profit with two other policies:

(a) **Risk-Averse Policy:** In this policy, we only continue missions whose request to use the secure resources can be fully satisfied. We stop the rest of the missions. Then, the secure resource allocation is a well-known *knapsack* problem where we have n objects with values of Q_i, weights of G_i and cost of c_i. The capacity of the knapsack is W_0. The optimal allocation of secure resources in this policy can be found by a greedy algorithm which serves the requests of the missions with a higher $\frac{Q_i}{c_i G_i}$ value first until the knapsack does not have enough capacity to store any other object.

(b) **Risk-Ignorant Policy:** In this policy, the risk is ignored. Hence, every mission is always continued, and the resource allocation is based on a simple Linear Program, a modified version of (2), where for each mission, $y_i = 1$ and the probability of occurrence of attacks $\alpha_j^i = 0$.

We present our results by showing the ratio of achieved total profit from these three policies over the maximum possible total profit. Based on the value of p_τ and the risk vectors shown in Table (3), we simulate the attack occurrences. The maximum possible total profit is achieved from a hypothetical scenario in which an oracle knows the attack occurrences ahead of time so that accordingly the assignment of the secure resources and the decisions to continue or stop missions are made with full knowledge. Note that in reality, we do not have such oracles, so the best we can achieve is to plan to maximize the expected total profit.

The average ratios are 0.52, -0.23 and 0.41 for our proposed approach, risk-ignorant, risk-averse policies, respectively. Moreover, the standard deviations of the ratios are 0.32, 1.11 and 0.1 for our proposed method, risk-ignorant, risk-averse policies, respectively. Fig. 2 shows the survivor functions (1-CDF) of the ratios for 100 iterations.

Insight 1. *We observe that the risk-ignorant method might result in a negative total profit with a probability of about* 0.50*, while our proposed approach leads to a positive total profit with a probability of about* 0.90*.*

Insight 2. *The risk-averse policy almost never results in a ratio (achieved total profit over maximum possible total profit) close to one. In fact, the probability that the risk-averse approach exceeds a ratio of* 0.5 *is only* 0.14*, while our proposed approach results in a ratio of* 0.5 *or greater with a probability of* 0.64*.*

6.4 Affect of the probability of experiencing a successful attack on the local servers (p_τ) and the normalized damage ($d^i_j/\min P_i$) on the performance

To investigate the effects of p_τ (probability of experiencing a successful attack on the local servers) and $d^i_j/\min P_i$ (damage per attack type normalized by the minimum possible payout from all the local servers) on the ratios (achieved total profit from the three polices over the maximum possible total profit), we perform two different simulations. In the first simulation, we still assume the probability of a successful attack follows the lognormal distribution ($p_\tau = 0.3827$), and investigate the affect of damage from attacks. In this experiment, we set all the attack types to have the same damage value ($d^i_j := D$) for a constant D. In Fig. 3a, we plot the average ratios from 100 iterations for each different values of D. We show the x-axis as D values normalized by the minimum possible achievable payouts from all the local servers (11000).

In the second simulation, we fix the damage as in Table (3), but we vary the probability of a successful attack (p_τ). We plot in Fig. 3b the average ratios from 100 iterations for each different values of p_τ.

Insight 3. *As can be seen from Fig. 3, when the probability of a successful attack* (p_τ) *or damage* (D) *have small values, our algorithm performance matches the risk-ignorant performance and even outperforms it. On the other hand, when* p_τ *or* D *are large, the performance of risk-ignorant policy drops drastically, while the performance of our algorithm matches the risk-averse policy performance.*

To conclude this subsection, we illustrate the average ratio of our proposed approach from 500 iterations for each different values of p_τ and D using a bar plot (Fig. 4). For $p_\tau = 0$ or $d^i_j = 0$, maximizing the expected total profit (our MILP-based approach) matches maximizing the total possible profit, and hence the ratio is 1. Furthermore, when $p_\tau = 1$, the ratio is also close to 1. This observation leads to insight 4:

Insight 4. *In the situations where there exist no* uncertainty *on the occurrence of attacks (i.e. the chance of a successful attack is either* 0 *or* 1*), our proposed approach reaches the maximum achievable total profit.*

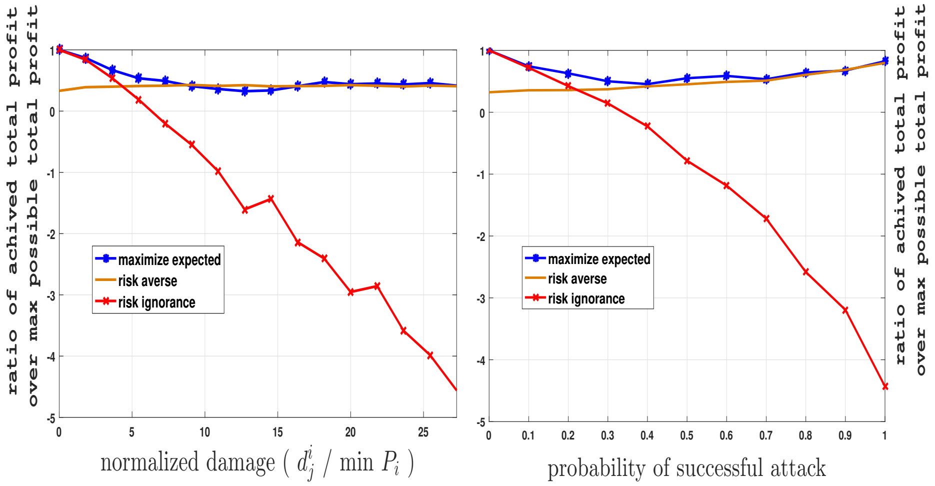

(a) dependency to risk damage

(b) dependency to p_τ

Fig. 3: Effect of p_τ and $d_j^i/\min\, P_i$ on the ratios

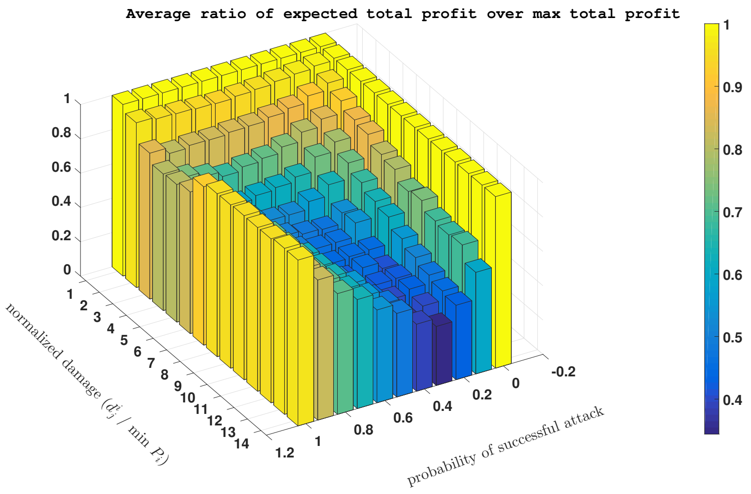

Fig. 4: Affect of different p_τ and d_j^i on the average ratio

Insight 5. *Generally, the loss in the total profit achieved by our proposed policy is caused by uncertainty on the occurrence of successful attacks. Particularly, when the probability of a successful attack is roughly about* 0.5*, the lowest performance of our approach is caused. More precisely, based on different levels of damage due to the attack, Table (4) shows the probabilities of a successful attack that leads to the lowest performance of our proposed method.*

Table 4: Successful attack probabilities that cause the lowest performance of our proposed approach for the different level of damage due to the attack.

$\boldsymbol{D}$	**13**	**12**	**11**	**10**	**9**	**8**	**7**	**6**	**5**	**4**	**3**	**2**	**1**
$\boldsymbol{p_{\mathcal{T}}}$	0.3	0.4	0.4	0.4,0.5	0.3,0.5	0.5	0.5,0.4	0.5,0.4	0.5,0.6	0.7	0.7	0.9	0.6,0.7

Table 5: Data and parameters of the synthetic dataset.

	Missions		
$\boldsymbol{\mathscr{W}_0 = 200}$	$\boldsymbol{M}_{\text{high}}$	$\boldsymbol{M}_{\text{med}}$	$\boldsymbol{M}_{\text{low}}$
Number of missions	5	10	10
$\boldsymbol{G_i}$	U([20-40])	U([20-40])	U([20-40])
$\boldsymbol{b_i}$	1	1	1
$\boldsymbol{c_i}$	2	2	2
Number of different types of attack	10	10	10
$\boldsymbol{\alpha_j^i}$	0.25	0.25	0.25
Estimated $\boldsymbol{\alpha_j^i}$	$\hat{\alpha}$	$\hat{\alpha}$	$\hat{\alpha}$
$\boldsymbol{P_i}$	U([120-150])	U([100-120])	U([80-100])
$\boldsymbol{Q_i}$	U([100-120])	U([80-100])	U([60-80])

6.5 Sensitivity Analysis

In this scenario, we study the affect of incorrect estimation of the probability of occurrence of different attack types (α_j^i) on the total profit. Over- or underestimating this probability can result in non-optimal decisions:

(a) Decide to continue a mission while the optimal decision is to stop the mission, and vice versa.

(b) Decide to allocate the secure resources to a mission while the optimal decision is to only assign the local resources to that mission.

Consider the following scenario where there are three different types of mission: 10 missions have a low-level payout, 10 missions have a medium-level payout, and 5 missions have a high-level payout. The capacity of the secure special resources/server $\mathscr{W}_0 = 200$, and each mission requires $U([20-40])$ units of resource. The cost for each mission for using their local and the secure servers are $b_i = 1$ and $c_i = 2$, respectively. The payout of the high-level missions from their local server is in the range of $U([120-150])$ while their payout from the secure server is in $U([100-120])$. Further, the payout of the medium-level missions from their local server is in $U([100-120])$ while their payout from the secure server is in $U([80-100])$. The payout of the low-level missions from their local server is in $U([80-100])$ while their payout from the secure server is in $U([60-80])$. In this scenario, each mission is exposed to 10 different types of attack. The true probability of each type of attack for all missions is $\alpha_j^i = 0.25$. Table (5) summarizes the parameters for this scenario.

For each value of $\hat{\alpha}$ (estimated probability of occurrence of different attack types) and d_j^i (damage due to those undesirable events), we run 500 iterations,

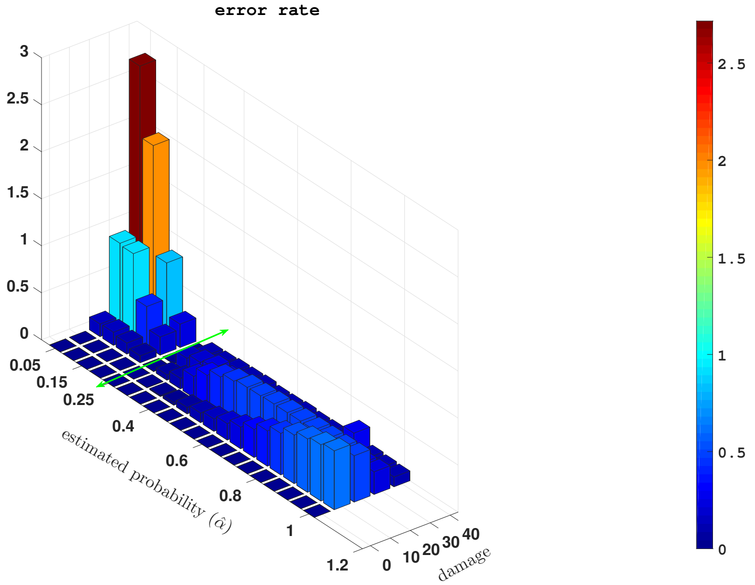

Fig. 5: Sensitivity to under/over-estimating the probability of occurrence of different types of attack. The green arrow shows the borderline of under- and over-estimation of the probability.

and find the average error rate where the error rate is calculated by first finding the difference between the total profit when there is no error on the estimated probability ($\hat{\alpha}$) and the total profit when under/over-estimating $\hat{\alpha}$. Then, the absolute value of the difference is divided by the total profit when there is no error on the estimated probability. An error rate of greater than 1 indicates that the obtained total profit by miscalculating the probability of occurrence of attacks has a negative value. As expected the more we deviate from the true probability of occurrence of attacks the higher the error rate is. Fig. 5 depicts the results when the true probability of occurrence of all types of attack is $\alpha_j^i = 0.25$.

Insight 6. *We observe that under-estimating the occurrence of different types of attack probabilities has more severe consequences than over-estimating them.*

Insight 7. *A counter-intuitive observation is that when over-estimating the probability of occurrence of different attack types, as the damage increases the error rate increases also. However, after the normalized damage reaches a certain level, the error rate starts to decrease.*

7 Limitations and Discussion

In this work, for mathematical simplicity, we assumed payouts (P_i's and Q_i's) and potential damage are proportional (linear) to the amount of allocated resources.

In general, partial payouts and damage might have a nonlinear relation to the amount of allocated resources. However, these mathematical assumptions are prevalent, and usually, do not have a significant impact on the overall performance.

Another limitation of our framework is that we assumed the secure resources are immune from any successful attack, and hence, a mission whose request is fully satisfied from the secure server will not suffer from any attack. This might not be the case in real-world scenarios, and there could be a chance that the secure server is vulnerable to risk too. However, our model can be easily modified to address this scenario.

8 Conclusions

In this paper, we presented a mission-oriented security model—an optimization-based framework to allocating resources that integrates security risk, cost and payout metrics to optimally allocate constrained secure resources to discrete actions called missions. We modeled this deployment or adaptive decision problem as a Mixed Integer Linear Program (MILP) which can be solved efficiently by different optimization solvers such as MATLAB MILP solver. Additionally, we proposed a novel quantitative risk assessment technique to learn the attack rates and risk characteristic from historical data. We used this technique to maximize the expected total profit gained from all the missions. We evaluated our model using the existing risk surveys and validated the model robustness and uncovered a number of insights on the importance of risk valuation in resource allocation.

This work is an effort in developing techniques on how to allocate resources in the face of risk. The capability afforded by this model will allow us to explore different policies by evaluating their effectiveness when dealing with varying characteristics of risk. In the future, we will explore a wide range of environments and assess its ability to promote effectiveness to different adversarial assumptions.

9 Acknowledgments

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

1. Anderson, R., Moore, T.: The economics of information security. Science (2006)
2. Anderson, R., Moore, T.: Information security economics–and beyond. Advances in cryptology (2007)

3. Celik, Z.B., Hu, N., Li, Y., Papernot, N., McDaniel, P., Walls, R., Rowe, J., Levitt, K., Bartolini, N., La Porta, T.F., et al.: Mapping sample scenarios to operational models. In: Military Communications Conference (MILCOM) (2016)
4. Cherdantseva, Y., Burnap, P., Blyth, A., Eden, P., Jones, K., Soulsby, H., Stoddart, K.: A review of cyber security risk assessment methods for scada systems. computers & security (2016)
5. Dekker, M., Liveri, D.: Cloud security guide for smes–cloud computing security risks and opportunities for smes. European Union Agency for Network and Information Security (ENISA) (2015)
6. Floudas, C.A.: Nonlinear and mixed-integer optimization: fundamentals and applications. Oxford University Press (1995)
7. Gordon, L.A., Loeb, M.P.: The economics of information security investment. ACM Transactions on Information and System Security (TISSEC) (2002)
8. Gordon, L.A., Loeb, M.P., Lucyshyn, W., Richardson, R.: 2006 csi/fbi computer crime and security survey. www.lfca.net/Reference Documents/2006 CSI-FBI Survey.pdf (2006), [Online; accessed 11-January, 2018]
9. Holm, H.: A large-scale study of the time required to compromise a computer system. IEEE Transactions on Dependable and Secure Computing (2014)
10. Hoo, K.J.S.: How much is enough? A risk management approach to computer security. Consortium for Research on Information Security and Policy ,Stanford University (2000)
11. Hu, N., La Porta, T., Bartolini, N.: Self-adaptive resource allocation for event monitoring with uncertainty in sensor networks. In: IEEE Mobile Ad Hoc and Sensor Systems (MASS) (2015)
12. Information technology – Security techniques – Information security risk management: `http://www.iso27001security.com/html/27005.html` (2017), [Online; accessed 11-January, 2018]
13. Jajodia, S., Ghosh, A.K., Swarup, V., Wang, C., Wang, X.S.: Moving target defense: creating asymmetric uncertainty for cyber threats, vol. 54. Springer Science & Business Media (2011)
14. Kaplan, S., Garrick, B.J.: On the quantitative definition of risk. Risk analysis (1981)
15. McDaniel, P., Jaeger, T., La Porta, T.F., Papernot, N., Walls, R.J., Kott, A., Marvel, L., Swami, A., Mohapatra, P., Krishnamurthy, S.V., et al.: Security and science of agility. In: ACM Workshop on Moving Target Defense (2014)
16. Nicol, D.M., Sanders, W.H., Trivedi, K.S.: Model-based evaluation: from dependability to security. IEEE Transactions on dependable and secure computing (2004)
17. Papoulis, A., Pillai, S.U.: Probability, random variables, and stochastic processes. Tata McGraw-Hill Education (2002)
18. Ponemon Institute: Cost of cyber crime study & the risk of business innovation. http://www8.hp.com/us/en/software-solutions/ponemon-cyber-security-report/ (2016), [Online; accessed 11-January, 2018]
19. Richardson, R.: Issues and trends: 2003 csi/fbi computer crime and security survey. www.lfca.net/Reference Documents/2003-CSI-FBI-Survey.pdf (2003), [Online; accessed 11-January, 2018]
20. Richardson, R.: 2010/2011 computer crime and security survey (2010)
21. Schneidewind, N.F.: Cyber security prediction models. Systems and Software Engineering with Applications (2005)