A Game-Theoretic Framework for Securing Interdependent Assets in Networks

Ashish R. Hota, Abraham A. Clements, Saurabh Bagchi, Shreyas Sundaram

1 Introduction

The prevalence of networked engineered systems in the twenty-first century has made it increasingly challenging to ensure their security and resiliency. While the growing interdependency between cyber and physical entities has led to improved system performance, it has also led to new avenues for attackers to target a large number of entities by exploiting those interdependencies. The magnitude, sophistication, and scope of such cyber-attacks have seen rapid growth; examples include a cyber-attack on the power grid in Ukraine [14], and a Distributed Denial of Service (DDOS) attack launched via Internet-of-Things devices [36].

These large-scale attacks share several important characteristics; i) these attacks proceed in multiple stages, ii) exploit interdependencies between diverse entities (e.g., vulnerabilities in devices made by third-party vendors have been exploited in an attacks [36]), iii) are stealthy, and iv) often exploit unknown or *zero-day* vulnerabilities. The existing literature has investigated approaches to protect against individual attack characteristics mentioned above. Vulnerabilities and their interdependencies are often modeled as *attack graphs* [42]. Bayesian networks are often used to determine how to defend against attacks within the attack graph representation

The authors are with the School of Electrical and Computer Engineering, Purdue University, e-mail: {ahota,clemen19,sbagchi,sundara2}@purdue.edu. In a preliminary version of this work [20], we only investigated the security risk minimization game, and considered different sets of case studies.

Abraham Clements is supported by Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2017-10889 B

This material is based in part upon work supported by the National Science Foundation under Grant Number CNS-1548114. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

[33]. The authors in [43] quantify security risks due to zero-day vulnerabilities in multi-stage attacks. Mathematical models of stealthy attacks include the framework of FlipIt games [40]. Dynamic defense strategies, such as *moving target defense (MTD)* [23] are increasingly being deployed to prevent stealthy [3] and multi-stage attacks [18, 44]. Note that under MTD, the system being protected is reconfigured either periodically or based on some events, so that it is harder to penetrate by an external adversary.

Most of the existing literature has focused on network security aspects from the perspective of a single or centralized defender. However, large-scale cyber-physical systems are seldom managed by a single entity, and instead, they are operated by multiple self-interested stakeholders. For instance, different independent system operators (ISOs) are responsible for managing different portions of the power grid. Nonetheless, assets that belong to these different stakeholders remain interdependent, a fact which is exploited by attackers to increase the magnitude and spread of their attacks. There is a growing body of work that applies tools and ideas from game theory in network security in order to model decentralized decision-making by multiple stakeholders. We provide a short summary of the existing literature that is relevant for our setting. A more comprehensive discussion of the literature is beyond the scope of this chapter, and for this, we refer interested readers to [1, 26, 39].

Most of the existing work can be classified into two distinct paradigms. In the first class of games, referred to as interdependent security games [25, 26], each node in the network is treated as an independent decision maker responsible for protecting itself. In the context of interdependent security games, the existing literature has investigated inefficiency of equilibria [24], effectiveness of cyber insurance [37], impacts of behavioral decision-making [19], and applications in industrial control systems [2], among others. In the second class of games, there are typically two players, an attacker and a defender, who compete over attacking and defending multiple targets. Game-theoretic models in this second framework include Stackelberg Security Games [39], Colonel Blotto Games [34], and Network Interdiction Games [21]. Applications of these models include protecting physical assets [39], analyzing military conflict [34], and securing cyber [9] and cyber-physical systems [15].

However, these models do not adequately capture interactions between multiple defenders each protecting multiple nodes in the network, while simultaneously facing strategic adversaries. In order to bridge this gap, we present a game-theoretic framework that i) incorporates essential features of both the above paradigms, ii) systematically captures the characteristics of sophisticated attacks discussed above, and ii) allows us to quantify the security risk of interdependent assets under both centralized and decentralized (game-theoretic) allocation of defense resources. While there are a few recent papers on multi-defender security games [27, 28], these papers assume that the strategy space of a defender is discrete, leading to very different analysis compared to our work.

We model the interdependencies between the assets that belong to possibly different defenders as a directed graph referred to as an *interdependency graph*. We present two complementary game-theoretic formulations. In both settings, the defenders assign defense resources to reduce attack success probabilities on the edges

of the interdependency graph, but with different objectives. In the first class of games, referred to as the *security risk minimization game*, the defenders minimize their expected loss (formally defined in Section 2) due to cyber-attacks on a subset of assets that they own (or are valuable to them). In the second class of games, referred to as the defense cost minimization game, each defender minimizes the cost of defense allocation subject to a maximum security risk (referred to as its *risk tolerance*) it is willing to tolerate on each asset it values. In this second class of games, the set of feasible strategies for a defender is a function of the strategies of other defenders, which makes it an instance of a generalized Nash equilibrium problem (see the chapter appendix for a discussion on this class of problems).

We establish the existence of a pure Nash equilibrium (PNE) in the security risk minimization game, and a generalized Nash equilibrium (GNE) (Definition 1 in the appendix to the chapter) in the defense cost minimization game. For both settings, we show that a defender can compute its best response (i.e., its optimal defense allocation for a given allocation by other defenders) by solving appropriately defined convex optimization problems. We demonstrate how our framework can be used to identify certain important aspects of MTD deployment; specifically, how frequently the configurations should be updated to meet security requirements.

We illustrate the application of our framework in two case studies arising in diverse applications. First we consider the IEEE 300 bus power grid network topology with three ISOs who manage different subsets of the buses. We compare the Nash equilibrium outcomes in both games with the outcomes where a central authority minimizes the sum of expected losses (and defense cost) of all defenders. For the security risk minimization game, we show that as the total budget decreases, the total expected losses under centralized and Nash equilibrium defense allocations increase exponentially. For the defense cost minimization game, we show that as the risk tolerance decreases, the total defense cost increases much faster under a Nash equilibrium allocation compared to the centralized allocation. We also study the increase in total defense cost at an equilibrium when multiple assets are supplied by a common vendor that can be compromised directly by an attacker. The second case study is on an e-commerce system adapted from [29]. We compute optimal MTD deployment by applying our framework, and investigate how security risk varies as a function of the defense budget.

2 Model

In this section, we introduce the mathematical framework that captures the different network security scenarios discussed above. We introduce the notion of an **interdependency graph** to model network interactions at different levels of abstractions. For example, interdependency graphs capture essential features of *attack graphs* [17] where a node represents a single attack step or vulnerability that can be exploited. Similarly, the nodes can also correspond to cyber or physical entities, such as firewalls or Human Machine Interfaces (HMIs), present in enterprise networks

and industrial control systems. Edges capture whether two nodes communicate with each other. At a higher level of abstraction, the interdependency graph can model large-scale networks such as the electric power grid where nodes represent buses, and edges represent physical interconnections between them.

Formally, an **interdependency graph** is a directed graph $G = \{V, E\}$. We refer to each node $v \in V$ as an *asset*. The presence of a directed edge $(v_j, v_i) \in E$ (with index $j \neq i$) indicates that when the asset v_j is compromised, it can be used to launch an attack on asset v_i. In the absence of any defense action, this attack succeeds with a probability $p^0_{j,i} \in (0,1]$. The success of attack on an edge is independent of the success of attacks on other edges.

We consider strategic attackers who launch sophisticated cyber attacks, such as Advanced Persistent Threats (APTs), into the network. These tools exploit the interdependencies between the assets to move within the network and compromise valuable assets. Without loss of generality, let s be the source node from which an attacker launches the attack from outside the network. If there are multiple entry points to the network, we can effectively replace them by a single entry point s by adding edges from s to all neighbors of all entry points in the original graph (with attack probabilities on these edges same as the original graph).

For an asset $v_i \in V$, let $\mathscr{P}_i$ be the set of directed paths from the source s to v_i on the graph; a path $P \in \mathscr{P}_i$ is a collection of edges $\{(s,u_1),(u_1,u_2),\ldots,(u_l,v_i)\}$ where $u_1,\ldots,u_l \in V$. The probability that v_i is compromised due to an attacker exploiting a given path $P \in \mathscr{P}_i$ is $\prod_{(u_m,u_n)\in P} p^0_{m,n}$ which is the product of probabilities (due to our independence assumption) on the edges that belong to the path P.

Remark 1. There are systematic ways to assign initial attack probabilities depending on the application. For instance, in the attack graph representation, initial attack probabilities are typically defined based on Common Vulnerability Scoring System (CVSS) scores [33]. The CVSS score is a widely adopted metric for assessing the severity of computer system security vulnerabilities. It incorporates the factors of how a vulnerability may be exploited, how difficult it is to exploit a vulnerability, what level of authentication is needed by an adversary, and which of the security dimensions of confidentiality, integrity, and availability are affected by the exploit.

Strategic Defender(s): Let $\mathscr{D}$ be the set of defenders; we use the index k to represent a defender. Defender $D_k \in \mathscr{D}$ is responsible for the security of a set $V_k \subseteq V \setminus \{s\}$ of assets. For each asset $v_m \in V_k$, there is a financial loss $J_m \in \mathbb{R}_{\geq 0}$ that defender D_k incurs if v_m gets compromised. If an asset v_m is not considered valuable, we can set $J_m = 0$. A defender allocates its resources to reduce the attack probabilities on the edges interconnecting different assets on the interdependency graph.

Let the feasible (defense) strategy set of defender D_k be $\mathbb{R}^{n_k}_{\geq 0}$, where n_k represents the (finite) different dimensions of responses that the defender can deploy. The defense resources reduce the attack probabilities on the edges of the interdependency graph. Accordingly, we introduce a transformation matrix $\mathbf{T}_k \in \mathbb{R}^{|E|\times n_k}_{\geq 0}$ which maps a feasible defense strategy $\mathbf{x}_k \in \mathbb{R}^{n_k}_{\geq 0}$ to a defense allocation on edges.

By appropriately defining the matrix $\mathbf{T}_k$, we can capture very general classes of defense strategies. We discuss two such examples.

Edge-based Defense Strategy: In this case, a defender D_k allocates defense resources on a subset of edges $E_k \subseteq E$ of the graph G, and accordingly $n_k = |E_k|$. For example, E_k can represent the set of all the edges that are incoming to nodes in V_k, i.e., defender D_k can reduce the attack probabilities on all the edges that are incoming to the nodes under its ownership. In general, it is not required to assume E_k's to be mutually disjoint; certain edges can potentially be managed by multiple defenders.

Node-based Defense Strategy: In this case, a defender D_k allocates defense resources to the set of nodes in V_k, and accordingly, $n_k = |V_k|$. Specifically, the defense resource x_i^k being allocated to node v_i implies that all the incoming edges to v_i in the graph G have a defense allocation x_i^k. Here $\mathbf{T}_k$ maps the allocation on a node into the edges that are incoming to it. An example of node-based defense strategy is IP-address randomization, an MTD technique where x_i^k potentially captures how frequently the IP-address on v_i is updated.

We now illustrate the concepts defined above. Consider the interdependency graph shown in Figure 1 with a source node and three nodes or assets. Let there be two defenders; defender 1 is responsible for assets 1 and 3, while defender 2 is responsible for asset 2. Under edge-based defense, let E_1 be all edges incoming to nodes 1 and 3, and accordingly, $n_1 = 4$, while under node-based defense, $n_1 = 2$. The respective transformation matrices are

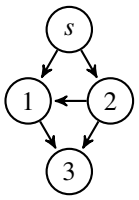

Fig. 1: Interdependency graph

$$\mathbf{T}_{e,1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \qquad \mathbf{T}_{n,1} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix}.$$

Both matrices have five rows corresponding to the edges in the graph in the order $(s,1),(2,1),(s,2),(1,3),(2,3)$. Note that under edge-based defense, defender 1's defense resources are not applied on the $(s,2)$ edge. Under node-based defense, both incoming edges to node 1 (as well as 3) receive identical defense resources.

We now introduce our assumptions behind defense effectiveness and cost.

Defense Effectiveness: Let $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{|\mathscr{D}|})$ be a joint defense strategy of the defenders. The attack success probability of an edge (v_j, v_i) under this joint defense strategy is denoted by $p_{j,i}(\mathbf{x})$. Note that, in our framework, it is possible to have multiple defenders simultaneously reducing the attack probability on a single edge. We make the following assumption on $p_{j,i}(\mathbf{x})$:

$$p_{j,i}(\mathbf{x}) := p_{j,i}^0 \exp(-\sum_{k=1}^{|\mathscr{D}|} \mathbf{t}_{j,i}^k \mathbf{x}_k), \tag{1}$$

where exp is the exponential function, and $\mathbf{t}^k_{j,i} \in \mathbb{R}^{1\times n_k}$ is the row vector in the transformation matrix $\mathbf{T}_k$ that maps the defense allocation $\mathbf{x}_k$ to the edge (v_j, v_i). In the example pertaining to Figure 1, $\mathbf{t}^1_{1,3} = \begin{bmatrix}0 & 0 & 1 & 0\end{bmatrix}$ under edge-based defense.

Under our assumption, the marginal reduction in attack probability decreases with increasing security investment.

Defense Cost: For a defender D_k and feasible defense strategy $\mathbf{x}_k$, we define the cost of defense allocation

$$c_k(\mathbf{x}_k) := \sum_{i=1}^{n_k} g^k_i(x^k_i). \tag{2}$$

We assume that g^k_i is strictly increasing and convex for every defender D_k and every $i \in \{1,2,\ldots,n_k\}$, and $g^k_i(0) = 0$. The convexity assumption captures increasing marginal cost of deploying more effective mitigation strategies.

Remark 2. Our assumptions on the defense cost and effectiveness are motivated by a recurring assumption in the security economics literature that probability of successful attack is decreasing and convex (similar to (1)) as a function of security investment [12, 24]. Under this interpretation, x represents investments in monetary or dollar amount, and it suffices to assume that $g(x) = x$. Our motivation behind choosing $g(x)$ to be increasing and convex is two fold.

1. It enables us to indirectly capture a broader class of defense effectiveness functions than (1). For example, suppose every edge is defended by at most one defender, and security investment reduces attack probability as $p(x) = p^0\exp(-\sqrt{x})$. We can capture such a scenario indirectly by defining $w = \sqrt{x}$ as the defense resource and cost function $g(w) = w^2$.
2. On the other hand, x could represent the unit of defense resource deployed, and $g(x)$ is the (possibly non-monetary) cost to the system. For example, in the context of IP-address randomization, x might represent the rate at which the IP-addresses are updated, while $g(x)$ could capture certain types of implementation overhead that are often nonlinear in x; examples of convex overhead costs include probability of genuine connection loss [6] and decrease in bandwidth [41].

Security Risk of an Asset: For an asset v_m, we define its security risk as

$$r_m(\mathbf{x}) := \max_{P\in\mathscr{P}_m} \prod_{(v_j,v_i)\in P} p_{j,i}(\mathbf{x}). \tag{3}$$

In other words, the security risk of an asset is given by the highest probability of attack on any path from the source to that asset on the interdependency graph. This is motivated by practical cyber-physical systems that face sophisticated adversaries and APTs and the security maxim that any interdependent system is only as secure as its weakest link. Our choice of defining security in terms of the worst case attack probabilities on an asset in (3) implicitly captures strategic attackers who aim to compromise valuable assets and choose a plan of attack that has the highest probability of success for each asset.

We consider two complementary problems that the defenders face.

Security Risk Minimization: In this problem, a defender minimizes its expected loss, where security risk on every asset is defined in equation (3), subject to a budget constraint on defense allocation. Let $\mathbf{x}_{-k}$ denote the defense allocation profile of all defenders except D_k. Then, the objective of D_k is to

$$\underset{\mathbf{x}_k \in \mathbb{R}^{n_k}_{\geq 0}}{\text{minimize}} \qquad \phi_k(\mathbf{x}_k, \mathbf{x}_{-k}) := \sum_{v_m \in V_k} J_m \cdot \left(\max_{P \in \mathscr{P}_m} \prod_{(v_j, v_i) \in P} p_{j,i}(\mathbf{x}_k, \mathbf{x}_{-k}) \right) \tag{4}$$

$$\text{subject to} \qquad \sum_{i=1}^{n_k} g_i^k(x_i^k) \leq b_k, \tag{5}$$

where $b_k > 0$ is the security budget of D_k. Note that the feasible strategy set $X_k := \{\mathbf{x}_k \in \mathbb{R}^{n_k}_{\geq 0} | \sum_{i=1}^{n_k} g_i^k(x_i^k) \leq b_k\}$ is non-empty, compact and convex. Furthermore, the cost function $\phi_k(\mathbf{x}_k, \mathbf{x}_{-k})$ depends on the strategy profile of all defenders. In Section 3.2, we analyze the existence of pure Nash equilibria (PNE) of the game between multiple defenders, and show to compute the best response of a player.

Defense Cost Minimization: In this problem, a defender minimizes its cost of defense allocation subject to constraints on the security risk on each asset it values. Let $\theta_m \in (0,1]$ be the *risk tolerance* of asset $v_m \in V_k$; it captures the maximum security risk (3) defender D_k is willing to tolerate on v_m. A smaller value of θ_m indicates that the defender prefers v_m to have a smaller security risk, and must choose its defense allocation accordingly. When $\theta_m = 1$ for an asset, the defender is essentially indifferent to whether the asset is attacked or remains secure. Thus, the defender can choose to not defend a subset of assets, e.g., by defining $\theta_m = 1$ for an asset v_m with $J_m = 0$. Note that $\theta_m \neq 0$ since the probability of successful attack is always nonzero under our assumptions.

Let $\mathbf{x}_{-k}$ denote the vector of defense allocation by all defenders other than D_k. The objective of D_k is to

$$\underset{\mathbf{x}_k \in \mathbb{R}^{n_k}_{\geq 0}}{\text{minimize}} \qquad f_k(\mathbf{x}_k) := \sum_{i=1}^{n_k} g_i^k(x_i^k) \tag{6}$$

$$\text{subject to} \qquad r_m(\mathbf{x}_k, \mathbf{x}_{-k}) \leq \theta_m, \qquad v_m \in V_k. \tag{7}$$

In other words, D_k has a risk tolerance for every asset it owns (denoted by the vector θ), and it wants to allocate the defense resources with minimum cost to achieve the desired risk tolerance.

Note that the cost function for defender D_k is independent of the strategies of other players, but the set of constraints in (7) is a function of the strategies of all other players. This class of problems is referred to as *generalized Nash equilibrium problems* (GNEPs). A brief overview is presented in the appendix to the chapter. In Section 4, we establish the existence of a generalized Nash equilibrium (GNE) in the game between defenders, and discuss how to compute the best response of a defender.

3 Security Risk Minimization Game

The analysis in this section relies on establishing the convexity of the optimization problem defined in eqs. (4) and (5). We start by introducing certain auxiliary variables. We define the *length* or distance of an edge (v_j, v_i) in terms of the attack probability under the given joint defense allocation $\mathbf{x}$ as,

$$l_{j,i}(\mathbf{x}) := -\log(p_{j,i}(\mathbf{x})) \geq 0, \tag{8}$$

where $p_{j,i}(\mathbf{x})$ is given by (1). A higher probability of an attack on an edge leads to smaller length for the edge. It follows from (1) that the modified length of the edge under a joint strategy profile $\mathbf{x}$ is given by

$$l_{j,i}(\mathbf{x}) := l^0_{j,i} + \sum_{D_k \in \mathscr{D}} \mathbf{t}^k_{j,i}\mathbf{x}_k := l^0_{j,i} + x_{j,i}, \tag{9}$$

where $l^0_{j,i} := -\log(p^0_{j,i})$, and $x_{j,i} = \sum_{D_k \in \mathscr{D}} \mathbf{t}^k_{j,i}\mathbf{x}_k$ captures the total defense allocation on the edge (v_j, v_i). Recall that $\mathbf{t}^k_{j,i}$ is the row vector corresponding to edge (v_j, v_i) in the transformation matrix $\mathbf{T}_k$. We denote the vector of modified lengths of the graph under joint defense strategy $\mathbf{x}$ as $\mathbf{L}(\mathbf{x}) = \mathbf{L}^0 + \sum_{D_k \in \mathscr{D}} \mathbf{T}_k\mathbf{x}_k$, where $\mathbf{L}^0$ is the vector of lengths in the absence of any defense allocation.

With this additional notation, we can express the probability that a node v_m is compromised via a given $P \in \mathscr{P}_m$ by

$$\prod_{(v_j,v_i)\in P} p_{j,i}(x_{j,i}) = \exp\left(-\sum_{(v_j,v_i)\in P} l_{j,i}(x_{j,i})\right). \tag{10}$$

Accordingly, the security risk on asset v_m is given by

$$r_m(\mathbf{x}) = \max_{P\in\mathscr{P}_m} \prod_{(v_j,v_i)\in P} p_{j,i}(x_{j,i}) = \exp\left(-\min_{P\in\mathscr{P}_m} \sum_{(v_j,v_i)\in P} l_{j,i}(x_{j,i})\right). \tag{11}$$

In other words, the path with the largest probability of successful attack is the path that has the smallest length under the transformation stated in equation (8). This observation enables us to utilize concepts from shortest path problems on graphs, discussed subsequently.

3.1 Existence of a Pure Nash Equilibrium

We are now ready to show the existence of a PNE in the game between multiple defenders.

Proposition 1. *The strategic game with multiple defenders where a defender minimizes its cost* $\phi_k(\mathbf{x}_k, \mathbf{x}_{-k})$ *defined in* (4)*, subject to* $\mathbf{x}_k \in X_k$ *defined in* (5)*, possesses a pure Nash equilibrium.*

Proof. From our transformation of attack probabilities into lengths on edges given in (8) and (9), the probability of successful attack on a node $v_m \in V_k$ due to a path $P \in \mathscr{P}_m$ and joint defense strategy $\mathbf{x}$ is equal to

$$\prod_{(u_j,u_i)\in P} p_{j,i}(\mathbf{x}) = \exp\left(-\sum_{(v_j,v_i)\in P}\left[l^0_{j,i} + \sum_{D_k\in\mathscr{D}} \mathbf{t}^k_{j,i}\mathbf{x}_k\right]\right).$$

Following (11), we can express the cost function of a defender D_k, defined in (4), as a function of its strategy $\mathbf{x}_k$ and the joint strategy of other defenders $\mathbf{x}_{-k}$ as

$$\phi_k(\mathbf{x}_k, \mathbf{x}_{-k}) = \sum_{v_m\in V_k} J_m \exp\left(-\min_{P\in\mathscr{P}_m}\sum_{(v_j,v_i)\in P}(l_{j,i}(\mathbf{x}_{-k}) + \mathbf{t}^k_{j,i}\mathbf{x}_k)\right), \tag{12}$$

where $l_{j,i}(\mathbf{x}_{-k}) = l^0_{j,i} + \sum_{D_l\in\mathscr{D}, l\neq k} \mathbf{t}^l_{j,i}\mathbf{x}_l$ for an edge (v_j, v_i).

Note that $\sum_{(v_j,v_i)\in P}\left[l_{j,i}(\mathbf{x}_{-k}) + \mathbf{t}^k_{j,i}\mathbf{x}_k\right]$ is an affine and, therefore, concave function of $\mathbf{x}_k$. The minimum of a finite number of concave functions is concave [5]. Finally, $\exp(-z)$ is a convex and decreasing function of z. Since the composition of a convex decreasing function and a concave function is convex, $\phi_k(\mathbf{x}_k, \mathbf{x}_{-k})$ is convex in $\mathbf{x}_k$ for any given $\mathbf{x}_{-k}$. Furthermore, the feasible strategy set $X_k = \{\mathbf{x}_k \in \mathbb{R}^{n_k}_{\geq 0} | \sum_{i=1}^{n_k} g^k_i(x^k_i) \leq b_k\}$ is non-empty, compact and convex for every defender D_k. As a result, the game is an instance of a *concave game* and has a PNE following Theorem 1 of [35]. □

3.2 Computing the Best Response of a Defender

The best response of D_k at a given strategy profile $\mathbf{x}_{-k}$ of others is defined as $\mathbf{x}^*_k := \operatorname{argmin}_{\mathbf{x}_k\in X_k} \phi_k(\mathbf{x}_k, \mathbf{x}_{-k})$. While the previous proposition shows that $\phi_k(\mathbf{x}_k, \mathbf{x}_{-k})$ is convex in $\mathbf{x}_k$, the cost function in (4) is non-differentiable. We now present an equivalent formulation of the problem below with a smooth cost function. Let $\mathbf{L}(\mathbf{x}_{-k}) = \mathbf{L}^0 + \sum_{D_r\in\mathscr{D}, r\neq k} \mathbf{T}_r\mathbf{x}_r$ be the vector of edge lengths under defense allocation $\mathbf{x}_{-k}$. For a given $\mathbf{x}_{-k}$, consider the following convex optimization problem:

$$\underset{\mathbf{y}\in\mathbb{R}^{|V|}_{\geq 0},\mathbf{x}_k\in\mathbb{R}^{n_k}_{\geq 0}}{\text{minimize}} \quad \sum_{v_m\in V_k} J_m e^{-y_m} \tag{13}$$

$$\text{subject to} \quad \mathbf{B}\mathbf{y} - \mathbf{T}_k\mathbf{x}_k \leq \mathbf{L}(\mathbf{x}_{-k}), \tag{14}$$

$$y_s = 0, \tag{15}$$

$$\sum_{i=1}^{n_k} g_i^k(x_i^k) \leq b_k, \tag{16}$$

where $\mathbf{B}$ is the node-edge incidence matrix of the graph G. Note that the constraint in (14) is affine. This formulation is motivated by similar ideas explored in the shortest path interdiction games literature [21, 38].

Remark 3. For a directed graph G, its incidence matrix is $\mathbf{B} \in \mathbb{R}^{|E|\times|V|}$, where the row corresponding to the edge (v_j, v_i) has entry -1 in the j^{th} column and 1 in the i^{th} column.

Remark 4. We refer to the vector $\{y_u\}_{u\in V}$ as a *feasible potential* if it satisfies (14) for every edge in the graph. We make the following observations.

1. The inequality in (14) for an edge is precisely the inequality that the Bellman-Ford algorithm tries to satisfy in every iteration. As shown in (8), the length of every edge is nonnegative in our setting. Therefore, the Bellman-Ford algorithm terminates with a feasible potential [8]. Note that we don't actually use the Bellman-Ford (or Dijkstra's) algorithm in solving the above problem.
2. Consider a path P from s to a node $v \in V$. Then, $y_v - y_s \leq \sum_{(v_j,v_i)\in P} l_{j,i}(\mathbf{x}_k, \mathbf{x}_{-k})$. In other words, when $y_s = 0$, y_v is a lower bound on the length of every path (and consequently the shortest path) from s to v.
3. In the absence of negative cycles, there always exists a feasible potential where y_v is *equal* to the length of the shortest path from s to v [8, Theorem 2.14] for every $v \in V$ (the solution of the Bellman-Ford algorithm).

We now prove the following result.

Proposition 2. *A defense strategy $\mathbf{x}_k^* \in \mathbb{R}^{n_k}_{\geq 0}$ is the optimal solution of the problem defined in eqs.* (13) *to* (16) *if and only if it is the minimizer of $\phi_k(\mathbf{x}_k, \mathbf{x}_{-k})$ defined in* (4) *subject to constraint* (5)*.*

Proof. Consider a strategy profile $\mathbf{x}_{-k}$ of all defenders other than D_k. Consider a feasible defense allocation vector $\mathbf{x}_k$ satisfying the constraint in (16). The joint strategy profile $\mathbf{x} = (\mathbf{x}_k, \mathbf{x}_{-k})$ defines a modified length vector $\mathbf{L}(\mathbf{x}_k, \mathbf{x}_{-k}) = \mathbf{L}(\mathbf{x}_{-k}) + \mathbf{T}_k\mathbf{x}_k$ on the edges of G. Let $\{y_u^{\mathbf{x}}\}_{u\in V}$ be the feasible potential where $y_u^{\mathbf{x}}$ is equal to the length of the shortest path from s to u under the joint defense allocation $\mathbf{x}$ for every $u \in V$. Now consider a path P from s to $v_m \in V_k$, and let P^* be a path of shortest length s to v_m. From Remark 4, we have

$$y^{\mathbf{x}}_{v_m} \leq \sum_{(u_j,u_i)\in P} l_{j,i}(\mathbf{x}_k,\mathbf{x}_{-k}) = -\sum_{(u_j,u_i)\in P} \log(p_{j,i}(\mathbf{x}))$$
$$\implies e^{-y^{\mathbf{x}}_{v_m}} \geq \prod_{(u_j,u_i)\in P} p_{j,i}(\mathbf{x}), \tag{17}$$

with equality for the path P^*. Accordingly, if $\mathbf{x}^*_k$ is optimal for the problem in eqs. (4) and (5), $\{\mathbf{x}^*_k, \{y_u^{(\mathbf{x}^*_k,\mathbf{x}_{-k})}\}_{u\in V}\}$ is feasible for the problem in eqs. (13) to (16), and both have identical cost. Therefore, the optimal cost for the problem in eqs. (13) to (16) is at most the optimal cost of eqs. (4) and (5).

Now let $\{\mathbf{x}^*_k, \{y^*_u\}_{u\in V}\}$ be the optimal solution of the problem defined in eqs. (13) to (16) for a given $\mathbf{x}_{-k}$. We claim that $y^*_{v_m}$ is equal to the length of the shortest path from s to v_m for every v_m with $J_m > 0$.

Assume on the contrary that $y^*_{v_m}$ is strictly less than the length of the shortest path from s to v_m, under the defense allocation $\mathbf{x}^*_k$. From Remark 4 we know that there exists a feasible potential $\{\hat{y}_u\}_{u\in V}$ such that $\hat{y}_{v_m}$ is equal to the length of the shortest path from s to v_m for every node $v_m \in V_k$ with length of every edge (u_j,u_i) given by $l_{j,i}(\mathbf{x}^*_k,\mathbf{x}_{-k})$. As a result, we have $y^*_{v_m} < \hat{y}_{v_m}$, and the objective is strictly smaller at $\hat{y}_{v_m}$, contradicting the optimality of $\{\mathbf{x}^*, \{y^*_u\}_{u\in V}\}$.

Therefore, at the optimal $\{\mathbf{x}^*_k, \{y^*_u\}_{u\in V}\}$, the cost in (13) is equal to the cost in (4) with defense allocation $\mathbf{x}^*_k$ (following similar arguments as the above paragraph). Furthermore, $\mathbf{x}^*_k$ is feasible for the problem in eqs. (4) and (5). Accordingly, the optimal cost for the problem in eqs. (4) and (5) is at most the optimal cost of eqs. (13) to (16). Combining both observations, we have the required result. □

We now discuss the security risk minimization problem from the perspective of a central authority.

Centralized Defense Allocation to Minimize Security Risk: The security risk minimization problem for a central authority is to find a defense allocation $\mathbf{x}^{\texttt{OPT}} \in \{\mathbb{R}^{\sum_{D_k\in\mathscr{D}} n_k}_{\geq 0} | \sum_{D_k\in\mathscr{D}} \sum_{i=1}^{n_k} g^k_i(x^k_i) \leq \sum_{D_k\in\mathscr{D}} b_k\}$ which minimizes $\sum_{D_k\in\mathscr{D}} \phi_k(\mathbf{x})$. This problem can also be solved via an analogous reformulation as eqs. (13) to (16), and the equivalence result from Proposition 2 applies for this case as well. In our case study in Section 6, we compare the security risks under both centralized and game-theoretic defense allocations.

Nash Equilibrium Computation: We compute the PNE strategy profile by iteratively computing the best responses for the defenders. This family of algorithms is referred to as *best response dynamics* [11]. Specifically, we apply the *sequential best response dynamics* in our case studies, and this scheme converges in all considered instances. However, proving theoretical guarantees on the convergence of best response-based update schemes is challenging for the following reasons. First, the expected loss of a defender represented in (12) is non-differentiable. Second, in the equivalent formulation eqs. (13) to (16), the players' cost minimization problems are coupled through their constraints which makes it an instance of a GNEP. Analysis of best response schemes for GNEPs is challenging with few algorithms that provide convergence guarantees. Therefore, a theoretical investigation of convergence of best response dynamics is beyond the scope of this chapter.

4 Defense Cost Minimization Game

In this section, we analyze the defense cost minimization game between multiple defenders. We start by showing that the risk tolerance constraints (7) are equivalent to a set of affine constraints in the defense allocation vector $\mathbf{x}$, and this fact will be useful in our proofs. Consider a node $v_m \in V_k$. Let $P_m \in \mathscr{P}_m$ be a path from the source node s to v_m. Let $r^0_{P_m} := (\prod_{(v_j,v_i)\in P_m} p^0_{j,i})^{-1}$.

Now consider the transformation matrix $\mathbf{T}_k$ for a defender D_k. Let $\mathbf{t}^k_{j,i}$ be the row vector that corresponds to the edge (v_j, v_i) as before. Furthermore, let $\mathbf{t}^k_{P_m} := \sum_{(v_j,v_i)\in P_m} \mathbf{t}^k_{j,i}$. We assume that for every node $v_m \in V_k$, and every path $P_m \in \mathscr{P}_m$, $\mathbf{t}^k_{P_m}$ has at least one nonzero entry, i.e., for every path from s to v_m, there exists at least one edge that D_k can defend. We compute

$$
\begin{aligned}
& r_m(\mathbf{x}) = \max_{P_m\in\mathscr{P}_m} \prod_{(v_j,v_i)\in P_m} p_{j,i}(\mathbf{x}) \leq \theta_m \\
\iff & \prod_{(v_j,v_i)\in P_m} p_{j,i}(\mathbf{x}) \leq \theta_m, \qquad \forall P_m \in \mathscr{P}_m \\
\iff & \left(\prod_{(v_j,v_i)\in P_m} p^0_{j,i}\right) \exp\left(-\sum_{(v_j,v_i)\in P_m}\sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{j,i}\mathbf{x}_l\right) \leq \theta_m, \qquad \forall P_m \in \mathscr{P}_m \\
\iff & \exp\left(-\sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{P_m}\mathbf{x}_l\right) \leq \theta_m r^0_{P_m}, \qquad \forall P_m \in \mathscr{P}_m \\
\iff & \sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{P_m}\mathbf{x}_l \geq -\log(\theta_m r^0_{P_m}), \qquad \forall P_m \in \mathscr{P}_m. \qquad (18)
\end{aligned}
$$

Therefore, each constraint in (7) can be expressed as a set of affine constraints.

4.1 Existence of a Generalized Nash Equilibrium

We now prove the existence of a GNE. Note that in the chapter appendix, we have formally defined the notion of a GNE, and provided a general result on the existence of a GNE (Theorem 1). First observe that Theorem 1 requires each X_k (for defender D_k) to be compact, while $\mathbb{R}^{n_k}_{\geq 0}$ is unbounded. In the proof, we define an appropriate compact subset of $\mathbb{R}^{n_k}_{\geq 0}$ for every player that contains the optimal defense allocation irrespective of the strategies of others.

Proposition 3. *The defense cost minimization problems contains a GNE.*

Proof. Let $\mathbf{x}^0_k$ be the optimal defense allocation of defender D_k when the allocation by every other player is 0. Let $\beta_k \in \mathbb{R}^{n_k}_{>0}$. Then $\hat{\mathbf{x}}_k := \mathbf{x}^0_k + \beta_k$ satisfies

$$\mathbf{t}^k_{P_m}\hat{\mathbf{x}}_k \geq -\log(\theta_m r^0_{P_m}) + \mathbf{t}^k_{P_m}\beta_k, \qquad \forall P_m \in \mathscr{P}_m, \forall v_m \in V_k, \tag{19}$$

and $\mathbf{t}^k_{P_m}\beta_k > 0$ following our assumption that $\mathbf{t}^k_{P_m}$ has at least one nonzero entry. We now define

$$X_k := \{\mathbf{x}_k \in \mathbb{R}^{n_k}_{\geq 0} | f_k(\mathbf{x}_k) \leq f_k(\hat{\mathbf{x}}_k)\}, \tag{20}$$

where $f_k(\mathbf{x}_k)$ is the cost of defense allocation $\mathbf{x}_k$ defined in (6). From the definition, it is easy to see that X_k is nonempty, convex (f_k is convex and X_k is its sublevel set) and compact (f_k is strictly increasing). In particular, $\mathbf{x}^0_k$ and $\hat{\mathbf{x}}_k$ belong to the set X_k because i) $f_k(\mathbf{x}^0_k) < f_k(\hat{\mathbf{x}}_k)$ by the optimality of $\mathbf{x}^0_k$, and ii) f_k is strictly increasing.

Now consider the set of constraints (7) for D_k that depend on the defense allocation of others. Formally, these constraints can be represented as a correspondence

$$C_k(\mathbf{x}_{-k}) := \{\mathbf{x}_k \in \mathbb{R}^{n_k}_{\geq 0} | \quad \mathbf{t}^k_{P_m}\mathbf{x}_k \geq -\log(\theta_m r^0_{P_m}) - \sum_{D_l \in \mathscr{D}, l \neq k} \mathbf{t}^l_{P_m}\mathbf{x}_l, \forall P_m \in \mathscr{P}_m, v_m \in V_k\}. \tag{21}$$

First observe that for any $\mathbf{x}_{-k} \in \mathbb{R}^{\sum_{l\neq k} n_l}_{\geq 0}$, $\mathbf{x}^0_k \in C_k(\mathbf{x}_{-k})$ since entries in $\mathbf{T}_k$ are nonnegative for every $D_k \in \mathscr{D}$. Therefore, the optimal solution of the problem eqs. (6) and (7), denoted by $\mathbf{x}^*_k(\mathbf{x}_{-k})$, has cost $f_k(\mathbf{x}^*_k(\mathbf{x}_{-k})) \leq f_k(\mathbf{x}^0_k)$, and accordingly $\mathbf{x}^*_k(\mathbf{x}_{-k}) \in X_k$. Therefore, without loss of generality, we can consider X_k to be the set of feasible defense allocation, and redefine the constraint correspondence as $\hat{C}_k(\mathbf{x}_{-k}) := C_k(\mathbf{x}_{-k}) \cap X_k \subseteq X_k$.

Now, suppose $\mathbf{x}_{-k} \in \mathbb{R}^{\sum_{l\neq k} n_l}_{\geq 0}$. Then $\hat{C}_k(\mathbf{x}_{-k})$ is nonempty (contains $\hat{\mathbf{x}}_k$ following (19) and (21)), closed and convex (intersection of closed and convex sets X_k and $C_k(\mathbf{x}_{-k})$). In addition, the constraint correspondence $\hat{C}_k$ is stated in terms of a set of inequalities where the associated functions (21) are continuous and affine (thereby, convex). Furthermore, from the definition of $\hat{\mathbf{x}}_k$ in (19), it satisfies all of the affine inequalities in (21) with strict inequality. Therefore, from Theorem 2 (in the chapter appendix), $\hat{C}_k$ is both upper and lower semi-continuous in X_{-k}.

Finally, the cost function f_k is independent of $\mathbf{x}_{-k}$, and is continuous and convex in X_k. Therefore, a straightforward application of Theorem 1 establishes the existence of a GNE. □

4.2 Computing the Best Response of a Defender

Recall that the cost function $f_k(\mathbf{x}_k)$ in (6) is independent of the strategies of other defenders, and is convex. The set of constraints in (18) are affine. Note that (18) defines one constraint for every path P_m from s to a given node v_m. Thus, the number of such constraints can be exponentially large in the worst case. We therefore propose the following equivalent problem where the number of constraints is equal to the sum of the number of nodes and edges in the interdependency graph.

Consider the following problem:

$$\underset{\mathbf{y}\in\mathbb{R}^{|V|},\mathbf{x}_k\in\mathbb{R}^{n_k}_{\geq 0}}{\text{minimize}} \quad \sum_{i=1}^{n_k} g_i^k(x_i^k) \tag{22}$$

$$\text{subject to} \quad \mathbf{B}\mathbf{y} - \mathbf{T}_k\mathbf{x}_k \leq \mathbf{L}(\mathbf{x}_{-k}), \tag{23}$$

$$y_s = 0, \tag{24}$$

$$y_m \geq -\log(\theta_m), \qquad \forall m \in V_k, \tag{25}$$

where $\mathbf{B}$ is the incidence matrix and $\mathbf{L}(\mathbf{x}_{-k})$ is the vector of edge lengths under the defense allocation $\mathbf{x}_{-k}$ by defenders other than D_k. We now prove the following equivalence result.

Proposition 4. *A defense strategy $\mathbf{x}_k^* \in \mathbb{R}^{n_k}_{\geq 0}$ is the optimal solution of the problem defined in eqs.* (22) *to* (25) *if and only if it is the minimizer of the problem defined in eqs.* (6) *and* (7)*.*

Proof. Let $\mathbf{x}_{-k}$ be the defense allocation by other defenders. Let $(\mathbf{x}_k, \mathbf{y})$ be feasible for the problem defined in eqs. (22) to (25). We show that $\mathbf{x}_k$ is feasible for the problem defined in eqs. (6) and (7). In particular, consider a path $P_m \in \mathscr{P}_m$ from s to $v_m \in V_k$. For every $(v_j, v_i) \in P_m$, (23) is equivalent to

$$\begin{aligned}
& y_j - y_i - \sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{j,i}\mathbf{x}_l \leq -\log(p^0_{j,i}) \\
\iff & y_m - \sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{P_m}\mathbf{x}_l \leq -\log\Big(\prod_{(v_j,v_i)\in P_m} p^0_{j,i}\Big) \qquad \text{(adding over all } (v_j,v_i)\in P_m\text{)} \\
\iff & \sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{P_m}\mathbf{x}_l \geq -\log(\theta_m) + \log\Big(\prod_{(v_j,v_i)\in P_m} p^0_{j,i}\Big) \\
\iff & \sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{P_m}\mathbf{x}_l \geq -\log(\theta_m r^0_{P_m}),
\end{aligned}$$

which satisfies (18). Therefore, $\mathbf{x}_k$ is feasible for the problem defined in eqs. (6) and (7), and the optimal cost of the problem eqs. (6) and (7) is at most that of the problem eqs. (22) to (25).

Now, let $\mathbf{x}_k$ be feasible for the problem eqs. (6) and (7). Define

$$y_m := \min_{P_m\in\mathscr{P}_m} \Big[\sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{P_m}\mathbf{x}_l - \log\Big(\prod_{(v_j,v_i)\in P_m} p^0_{j,i}\Big)\Big],$$

and $y_s = 0$. In other words, y_m is the length of the shortest path from s to v_m under the joint strategy profile $(\mathbf{x}_k, \mathbf{x}_{-k})$. Thus, it satisfies (23). In addition, it follows from (18) that for every $P_m \in \mathscr{P}_m$,

$$\begin{aligned}
& \sum_{D_l\in\mathscr{D}} \mathbf{t}^l_{P_m}\mathbf{x}_l - \log\Big(\prod_{(v_j,v_i)\in P_m} p^0_{j,i}\Big) \geq -\log(\theta_m) \\
\implies & y_m \geq -\log(\theta_m).
\end{aligned}$$

Thus, $(\mathbf{x}_k, \mathbf{y})$ is feasible for the problem eqs. (22) to (25). Therefore, the optimal cost of the problem eqs. (22) to (25) is at most that of the problem eqs. (6) and (7).

Combining both observations, we have the desired result. □

Centralized Defense Allocation to Minimize Defense Cost: The defense cost minimization problem for a central authority is to find a defense allocation $\mathbf{x}^{\texttt{OPT}} \in \mathbb{R}_{\geq 0}^{\sum_{D_k \in \mathscr{D}} n_k}$ which minimizes $\sum_{D_k \in \mathscr{D}} f_k(\mathbf{x})$ subject to risk tolerance constraints for every $v_m \in V_k$, $D_k \in \mathscr{D}$. This problem can be solved via an analogous reformulation as eqs. (22) to (25); the equivalence result from Proposition 4 applies in this case.

In our case studies, we compute GNE defense allocations by employing the sequential best response algorithm as discussed earlier, and compare the security risks under both centralized and game-theoretic defense allocations.

In the following section, we show how to compute optimal deployment of MTD by applying the framework developed thus far.

5 Moving Target Defense

As discussed in the introduction, one of our goals is to consider MTD techniques that eliminate the advantage that strategic adversaries have against a static defended system. This advantage arises from the fact that the adversary can seek to breach such a static system repeatedly, with different (and likely continually learning) attack techniques. In order to capture this mathematically, we consider the notion of *Time-to-Compromise* of an asset [30]. Specifically, the time to successfully compromise an asset v_i, via an attack launched from v_j, is a random variable denoted $Q_{j,i}$ with an associated distribution function $F_{j,i}$. We assume that the support of $Q_{j,i}$ is $[0,\infty)$ for every $(v_j, v_i) \in E$. As before, we denote the *baseline* attack probability on v_i, launched from v_j, by $p^0_{j,i} \in (0,1]$; this represents the probability of successful attack when i) there is no defense allocation on the edge (v_j, v_i), and ii) the attacker has an infinite amount of time to compromise v_i.

While deploying MTD, a key variable that determines its effectiveness as well as the deployment cost is *how fast the configuration is changed dynamically*. For instance, consider the class of Dynamic Network defense techniques that relies on randomizing network IP addresses that have been shown to be effective against many types of attacks [3, 22]. If the network addresses are changed more slowly (for instance, once every few months), it gives the attacker sufficient time to learn about system vulnerabilities and execute its attack. On the other hand, if the addresses are changed more frequently, then it deters certain types of attacks more effectively. However, this also increases the overhead cost, such as the number of IP addresses that the defender must own, as well as the cost to legitimate clients, e.g., due to disconnections of network sessions. We now formalize this idea.

For ease of exposition, we only discuss an edge-based defense strategy where each edge receives an independent MTD deployment. We denote $\tau_{j,i} \in [0,\infty)$ as the time period between two successive changes of configuration of the edge (v_j, v_i) under a certain MTD deployment. A smaller $\tau_{j,i}$ represents a higher frequency of configuration changes. While evaluating the effectiveness of MTD, we only con-

sider attacks that succeed within a given configuration in this section. A change of configuration while the attack is in progress i) prevents the attack from succeeding, and ii) enables the defender to detect the attack and take corrective measures. In other words, for the attack on v_i to succeed, we must have $Q_{j,i} \leq \tau_{j,i}$. Accordingly, the probability of a successful attack on v_i is given by

$$p_{j,i}(\tau_{j,i}) = p^0_{j,i} F_{j,i}(\tau_{j,i}). \tag{26}$$

More generally, we refer to $\tau_{j,i}$ as the defense allocation on the edge (v_j, v_i) and τ_E as the vector of defense allocations on all edges. As before, we assume that the success of this attack is independent of the success of attacks propagating through other edges in the graph.

The defender incurs a cost $g^m_{j,i}(\tau_{j,i})$ for its choice of MTD allocation $\tau_{j,i}$ on the edge (v_j, v_i). We make the following assumptions on the function $g^m_{j,i}$.

Assumption 1. *The functions $g^m_{j,i}$ have the following properties.*

1. *$g^m_{j,i}$ is strictly decreasing and convex.*
2. *$g^m_{j,i}(0) = \infty$ and $g^m_{j,i}(\tau) > 0$ for any finite $\tau \in [0, \infty)$.*

In other words, the defender incurs a higher cost for more frequent configuration updates and this cost is infinite for updating continuously. For finite choice of period τ, the defender incurs a nonzero cost. As an example, the functions $g^m_{j,i}(\tau) = e^{-\alpha\tau}, \alpha > 0$ and $g^m_{j,i}(\tau) = \frac{1}{\tau}$ satisfy the above assumption.

In the context of MTD deployment, we will consider both security risk minimization, and defense cost minimization problems stated in Section 2. Formally, the security risk minimization problem for a single defender is to

$$\text{minimize} \qquad \sum_{v_m \in V} J_m \cdot \left(\max_{P \in \mathscr{P}_m} \prod_{(v_j, v_i) \in P} p_{j,i}(\tau_{j,i}) \right) \tag{27}$$

$$\text{subject to} \qquad \tau_{j,i} \geq \gamma_{j,i}, \qquad (v_j, v_i) \in E, \tag{28}$$

$$\sum_{i=1}^{n_k} g^m_{j,i}(\tau_{j,i}) \leq b, \tag{29}$$

where $\gamma_{j,i}$ is a bound on how fast the configuration can be updated, possibly due to physical constraints, and b is the budget. The game-theoretic setting and the defense cost minimization problem can be defined in an analogous manner, and is omitted.

Convexity under Exponential Distributions

Many probability distribution functions (for instance, Exponential and Laplace) are log-concave [4]. Log-concavity does not necessarily imply that the function is convex. Nonetheless, for Exponentially distributed $Q_{j,i}$'s, we obtain sufficient conditions under which the problem defined in eqs. (27) to (29) is in fact convex.

Let $F_{j,i}$ be any continuous strictly monotone distribution function, such as the distribution function of an exponential random variable. Similar to Section 3.2, we define the *length* of an edge (v_i, v_j) under defense allocation $\tau_{j,i}$ as

$$\begin{aligned} l_{j,i}(\tau_{j,i}) &:= -\log(p_{j,i}(\tau_{j,i})) = -\log(p^0_{j,i}) - \log(F_{j,i}(\tau_{j,i})) \\ &:= l^0_{j,i} + x_{j,i}(\tau_{j,i})). \end{aligned} \tag{30}$$

In other words,

$$x_{j,i}(\tau_{j,i})) := -\log(F_{j,i}(\tau_{j,i})) \tag{31}$$

$$\iff e^{-x_{j,i}} = F_{j,i}(\tau_{j,i}) \iff \tau_{j,i} = F^{-1}_{j,i}(e^{-x_{j,i}}). \tag{32}$$

Note that $l^0_{j,i}$ is the length of the edge without any defense allocation and the quantity $x_{j,i}$ (a function of $\tau_{j,i}$) increases the length linearly. Let $\mathbf{L}^0$ be the vector of lengths without any defense allocation, and let $\mathbf{x}$ be the vector of $x_{j,i}$ variables. We now state the following assumptions on cost functions $g_{j,i}$ and exponentially distributed time-to-compromise random variable $Q_{j,i}$.

Assumption 2. *For every edge (v_j, v_i), i) $g_{j,i}(\tau) = e^{-\alpha_{j,i}\tau}, \alpha_{j,i} > 0$, ii) $F_{j,i}(\tau) = 1 - e^{-\beta_{j,i}\tau}, \beta_{j,i} > 0$, and iii) $\beta_{j,i} < \alpha_{j,i}$.*

Now consider the following optimization problem.

$$\underset{\mathbf{y}\in\mathbb{R}^{|V|}_{\geq 0}, \mathbf{x}\in\mathbb{R}^{|E|}_{\geq 0}}{\text{minimize}} \quad \sum_{v_m \in V} J_m e^{-y_m} \tag{33}$$

$$\text{subject to} \quad \mathbf{B}\mathbf{y} - \mathbf{x} \leq \mathbf{L}^0, \tag{34}$$

$$y_s = 0, \tag{35}$$

$$\sum_{(v_i,v_j)\in E} (1 - e^{-x_{j,i}})^{\frac{\alpha_{j,i}}{\beta_{j,i}}} \leq b, \tag{36}$$

$$0 \leq x_{j,i} \leq \log(\frac{\alpha_{j,i}}{\beta_{j,i}}), \tag{37}$$

where $\mathbf{B}$ is the incidence matrix of the interdependency graph. Our main result in this subsection shows that the above problem is convex and solves the optimization problem stated in eqs. (27) to (29) when $\gamma_{j,i} = \frac{-1}{\beta_{j,i}} \log\left(1 - \frac{\beta_{j,i}}{\alpha_{j,i}}\right)$. We start with the following lemmas. We drop the subscript i, j in the following analysis.

Lemma 1. *Under Assumption 2, the function $g(F^{-1}(e^{-x})) := (1 - e^{-x})^{\frac{\alpha}{\beta}}$ is convex in x over the domain $x \in [0, \log(\frac{\alpha}{\beta})]$.*

Proof. For the exponential distribution function, we have

$$F(\tau) = 1 - e^{-\beta\tau} \implies e^{-\beta\tau} = 1 - F(\tau) \implies \tau = \frac{-1}{\beta}\log(1 - F(\tau))$$
$$\implies F^{-1}(w) := \frac{-1}{\beta}\log(1-w),$$

where $w := F(\tau)$. Then, the cost function can be expressed as

$$g(F^{-1}(e^{-x})) = g\left(\frac{-1}{\beta}\log(1-e^{-x})\right) = \exp\left(\frac{\alpha}{\beta}\log(1-e^{-x})\right) = (1-e^{-x})^{\frac{\alpha}{\beta}}.$$

We now verify that the function $h(x) := g(F^{-1}(e^{-x})) = (1-e^{-x})^{\frac{\alpha}{\beta}}$ is increasing and convex for $x \in [0, \log(\frac{\alpha}{\beta})]$. We denote $\frac{\alpha}{\beta} = z$ and compute

$$\begin{aligned}
h'(x) &= z(1-e^{-x})^{(z-1)}(e^{-x}) \\
h''(x) &= z(z-1)(1-e^{-x})^{(z-2)}e^{-2x} + z(1-e^{-x})^{(z-1)}(-e^{-x}) \\
&= z(1-e^{-x})^{(z-2)}e^{-x}[(z-1)e^{-x} - (1-e^{-x})] \\
&= z(1-e^{-x})^{(z-2)}e^{-x}[ze^{-x} - 1].
\end{aligned}$$

We need $z > e^x$ for $h''(x) > 0$, or equivalently, $x < \log(z) = \log(\frac{\alpha}{\beta})$. □

Lemma 2. *Let $\beta < \alpha$, Then, $x \leq \log(\frac{\alpha}{\beta}) \iff \tau \geq \frac{-1}{\beta}\log(1-\frac{\beta}{\alpha}) = \gamma$.*

Proof. Recall from (31) that $x = -\log(F(\tau)) = -\log(1-e^{-\beta\tau})$. Then,

$$x = -\log(F(\tau)) \leq \log(\frac{\alpha}{\beta}) \iff F(\tau) = 1 - e^{-\beta\tau} \geq \frac{\beta}{\alpha},$$
$$\iff e^{-\beta\tau} \leq 1 - \frac{\beta}{\alpha} \iff -\beta\tau \leq \log(1-\frac{\beta}{\alpha}) \iff \tau \geq \frac{-1}{\beta}\log\left(1-\frac{\beta}{\alpha}\right).$$

This concludes the proof. □

We now prove the following result.

Proposition 5. *Suppose Assumption 2 holds, and let $\gamma_{j,i} = \frac{-1}{\beta_{j,i}}\log(1-\frac{\beta_{j,i}}{\alpha_{j,i}})$. Then eqs.* (33) *to* (37) *represent a convex optimization problem that is equivalent to the security risk minimization problem stated in eqs.* (27) *to* (29)*.*

Proof. From Lemma 1, we observe that the constraints (29) and (36) are equivalent, and are convex. Similarly, from Lemma 2, we observe that the constraints (28) and (37) are equivalent. We reach the desired result following identical arguments as the proof of Proposition 2. □

In the following section, we compare centralized and PNE defense allocation in a case study on the IEEE 300 bus power grid network. In Section 7, we compute optimal MTD deployment for an e-commerce system for both security risk and defense cost minimization problems.

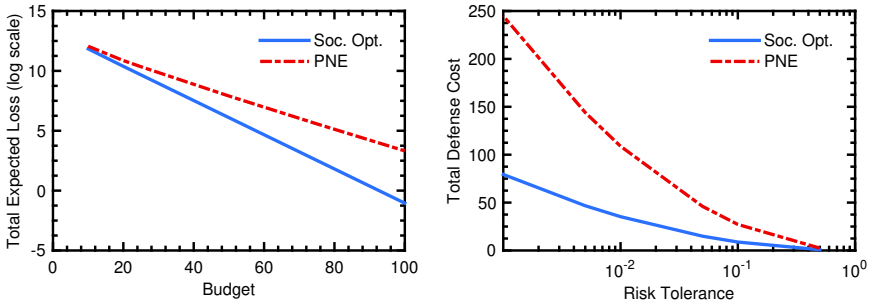

(a) Total expected loss

(b) Optimal and equilibrium cost

Fig. 2: Comparison of centralized and PNE defense allocations for the IEEE 300 bus power grid network. Figures 2a and 2b correspond to the security risk minimization and defense cost minimization games respectively. The costs are in an abstract unit.

6 Case Study 1 - IEEE 300 Bus Power Network

A large-scale network, such as the power grid, contains thousands of cyber and physical entities. Therefore, many different types of attacks are possible against such a system. Our first case study illustrates how our framework is applicable in this context via the following stylized example. Note that the following choice of the interdependency graph, cost functions, and attack probabilities are only made for illustrative purposes. Depending on the setting, a practitioner must instantiate the model appropriately.

We consider the widely used benchmark IEEE 300 bus power grid network [7]. We define the network itself as the interdependency graph where each node represents a bus (i.e., the network has 300 nodes), and the physical interconnection between the buses represent the edges. Each bus has generators and/or load centers associated with it. The 300 bus network data divides the buses or nodes into 3 different regions containing 159, 78 and 63 nodes respectively [7]. We assume that each region is managed by an independent entity or defender. The defenders want to protect the buses within their region that contain the generators; each generator is valued at its maximum generation capacity. The attacker can directly access three nodes (specifically, bus 39, 245 and 272).

All computations in this section are carried out in MATLAB using the convex optimization solver CVX [13].

We first consider the security risk minimization problem. We assume that the cost function is $g(x) = x$. Here x potentially represents the monetary amount spent on securing an asset, while our assumption in (1) ($p(x) = p^0 \exp(-x)$) captures how effective the monetary investment is in reducing the attack probability. We further assume that every edge has an initial probability of successful attack of magnitude 1. For a given total budget, we compute the centralized defense allocation that

minimizes the total expected loss. We divide the total budget among the players proportional to the number of nodes they control, and compute the PNE defense allocation by iteratively computing their respective best responses. We observe that both simultaneous and sequential best response dynamics converge to PNE within 25 iterations starting from random initial defense allocations. Figure 2a shows the total expected loss (in the logarithmic scale with base e) experienced by all three players at the PNE and under the centralized defense allocation for different total budgets. The total expected loss is larger at the PNE, and the relative change in the total expected loss at the PNE grows from 1.8% to over 7500% as the budget increases from 1 to 100. When the total budget is 100, the total expected loss at the social optimum is 0.35 while it is 27.56 (or 3.3164 in the log scale as shown in the plot) at the PNE.

We then consider the defense cost minimization problem. We assume that the cost function is given by $g(x) = x^2$ for every defender and for every edge in the network. Our motivation behind this choice is the crash overdrive malware attack on the Ukraine power grid [14]. MTD techniques such as IP-address randomization are effective against reconnaissance scans which the above malware relies on; here x potentially represents the rate at which IP-addresses are updated. Following the discussion in Remark 2, we choose $g(x) = x^2$ which better captures nonlinear growth of certain types of overhead costs [6, 41]. Since $g(x)$ could be interpreted as both monetary as well as overhead costs, we assume that it is in an abstract unit. In Figure 2b, we compare the total defense cost required to enforce a given tolerance level (shown in the x-axis) at each generator node under centralized and PNE defense allocations. As the risk tolerance decreases, the defense cost at the PNE increases faster than the defense cost under the centralized defense allocation.

Interdependency Through Common Vendor

As we discussed earlier, strategic attackers have exploited vulnerabilities in assets prepared by a common vendor to increase the spread of their attacker in recent years. In this subsection, we show how our framework can be used by practitioners to quantify the (potentially higher) security risk they face when multiple assets are from a common vendor. This is a common occurrence in practice where the same hardware or software (or both) is in use at multiple sub-systems owned by different stakeholders and any vulnerability in it can affect multiple assets. We again consider the IEEE 300 bus network with $g(x) = x^2$ (with an abstract unit). We represent the vendor by a new node, and connect the vendor to eight different generator nodes (belonging to different players), i.e., if an attacker successfully compromises the vendor, it can launch attacks on the generators directly. The attacker can directly attack the vendor node.

We first assume that $p^0_{j,i} = 1$ on every edge in the network, except for the edge from the attacker node to the vendor node. In Figure 3a, we show the total defense cost at the PNE to meet a given risk tolerance; the quantity p_v represents the prob-

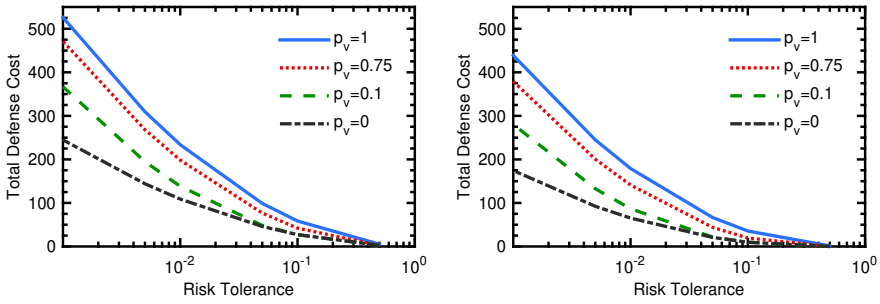

(a) Initial probability on edges: 1 (b) Initial probability on edges: 0.75

Fig. 3: Total defense cost at a PNE for the IEEE bus power grid network with common vendor. The probability of successful attack on the vendor directly from the attacker is represented as p_v. Initial probabilities of successful attack on all edges are 1 and 0.75 respectively. The defense cost is in an abstract unit.

ability of successful direct attack on the vendor by the attacker. The case where the vendor is not present is denoted by $p_v = 0$ (which is the case from Figure 2b). As p_v increases, it becomes easier for the attacker to attack the generators via the vendor, and accordingly the budget required to meet a given tolerance increases. We find identical trends when the $p^0_{j,i} = 0.75$ on every edge in the network (except the edge from the attacker node to the vendor node) the results for which are shown in Figure 3b. When the risk tolerance is 0.5, the figure shows that the total defense costs are equal when $p_v = 0$, $p_v = 0.1$ and $p_v = 0.5$, which is expected because the attack probability via the vendor is smaller than 0.5.

The practical implication of this result is that quantifying the security risks due to assets from third-party vendors could lead to designing adequate countermeasures and financial incentives (such as adding appropriate security requirements in procurement and support contracts with the vendors), which will then potentially reduce the likelihood and spread of such attacks in the future. Our treatment enables any stakeholder to quantitatively calculate the risk of compromise of its asset due to shared vulnerability at a vendor.

7 Case Study 2 - Moving Target Defense of E-Commerce System

We consider an e-commerce distributed system studied in [29] to illustrate how our framework can be used to compute optimal MTD deployment. Figure 4b shows the devices, and Figure 4c shows the corresponding attack graph for the e-commerce system. The attacker aims to obtain the customer information, such as credit card numbers, from a database. The attacker needs to find a suitable subset of the twelve

Nodes	Devices	Attack	Defense					
			DN	DP	DE	DS	DD	CD
1,2	Web Server	Network Reconnaissance[a]	x					
3	Web Server	Vulnerability Exploit[b]		x	x	x		
4	Web Server	Credential attack[c]						x
5	Internal Network	Network Reconnaissance	x					
6	Web Server	Credential attack						x
7,8	DB or App. Server	Exploit a service		x	x	x		
9,11	DB or App. Server	Credential attack						x
10,12	Database Server	Read from Database					x	
T	Database Server	Read credit card data	No Defense					

(a) Description of attack steps and respective MTDs

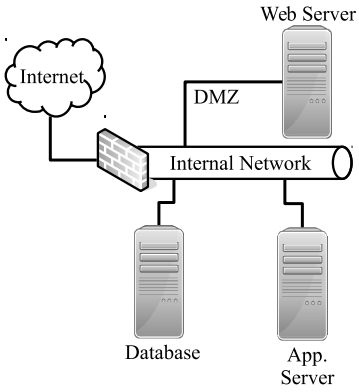

(b) E-commerce distributed system [29]

[a] Network reconnaissance actively or passively probes the network configuration to identify vulnerable systems.

[b] Exploitation relies on vulnerabilities in software which enable attackers to perform otherwise prohibited operations.

[c] Credential attacks involve obtaining valid credentials (e.g. passwords) to a system (for example, by brute force).

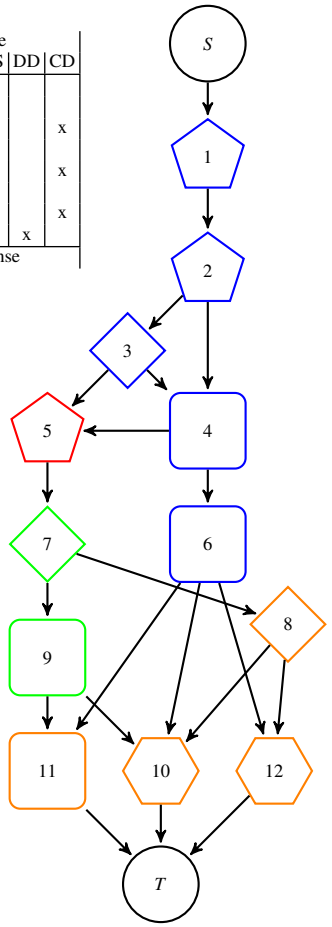

(c) Attack graph representation

Fig. 4: Representation of e-commerce network. In the attack graph (c) colors indicate the targeted device and shapes indicate type of attack. For mapping see (a); Defense used in (a) are: Dynamic Networks (DN), Dynamic Platforms (DP), Dynamic Environments (DE), Dynamic Software (DS), Dynamic Data (DD), and Credential Defense (CD).

attack steps on four devices to achieve this goal. The devices are a web server (located in a DMZ), the internal network, a database server, and an application server; both servers are located on the internal network. For each node in the attack graph, we describe the type of attack that can compromise it, and the type of MTD from [32] that can be deployed in Figure 4a.

We treat the attack graph in Figure 4c as the interdependency graph. The attacker has a single entry point into the network at node S, and it targets node T. We as-

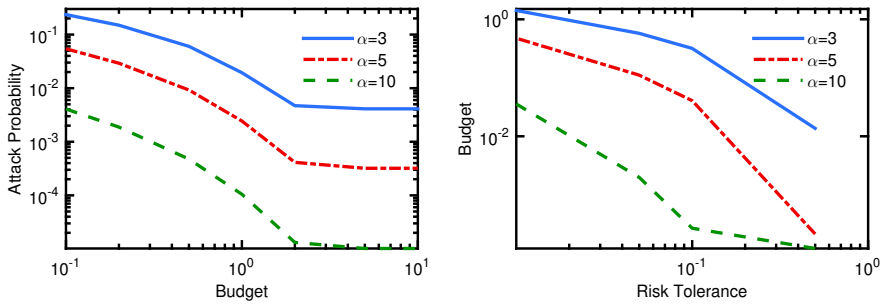

(a) Attack probability as a function of budget

(b) Optimal defense cost (budget) as a function of risk tolerance

Fig. 5: Optimal MTD deployment on e-commerce network. Figure 5a corresponds to the security risk minimization problem and Figure 5b corresponds to the defense cost minimization problem. The budget is in an abstract unit.

sume that the initial probability of successful attack on every edge is 1. In practice, these initial probabilities can be defined in terms of their CVSS scores [33]. We consider a node-based defense strategy. At every node, the frequency at which the corresponding MTD is updated represents a decision variable.

We consider the setting in Section 5 where a higher frequency of updating the MTD reduces the attacker's advantage. We assume that the cost function and distribution of time required for successful compromise satisfy Assumption 2. Specifically, for every edge (v_j, v_i), we assume that the random variable $Q_{j,i}$ is exponentially distributed with distribution function $F_{j,i}(\tau) = 1 - e^{-\tau}$, i.e., the parameter $\beta_{j,i} = 1$. We also consider an identical cost function $g_{j,i}(\tau) = e^{-\alpha\tau}$ for every edge in an abstract unit, and consider three different values of $\alpha \in \{3, 5, 10\}$ in our simulations. We consider both security risk minimization and defense cost minimization problems from the perspective of a single (centralized) defender. We used MATLAB's *fmincon* routine with *active-set* and *sqp* solvers to compute optimal defense allocation for both problems. Numerical results show that nodes $1, 2, 4, 6$ and 10 receive higher defense allocation than other nodes. This is expected because the initial attack probabilities are identical, and these nodes lie on a path with the smallest number of edges.

Figure 5a shows how the attack probability on the target node decreases under the optimal defense allocation with a given budget. First observe that at a given budget, when α is larger, the attack probability is smaller. For a given β, higher values of α implies that we can assign a larger defense allocation on the edges (i.e., $x_{j,i}$'s) without violating constraints (36) and (37). As a result, we obtain a smaller attack probability at the target node. Note further that as the budget increases, the attack probability initially decreases, but it gets saturated beyond a certain budget. The reason for this is the constraints on the defense allocation (37) limits how fast the MTD configuration can be updated. While the constraint in (37) is imposed to

preserve the convexity of the optimization problems, qualitatively similar behavior will emerge when the constraints are due to physical limitations on the frequency of configuration updates.

Figure 5b shows the cost of defense allocation to enforce that the probability of attack on the target is smaller than the risk tolerance. We observe that the relationship between the two is approximately piecewise linear in the logarithmic scale. As before, a higher value of α implies a smaller budget requirement for a given level of risk tolerance.

8 Conclusion

In this chapter, we presented two complementary game-theoretic models to study the security of networked systems. We considered multiple self-interested defenders, each of whom manages a set of assets represented by nodes in a directed graph. Attacks spread through the network via the interconnecting links in the graph. In the first class of games, each defender minimizes its expected loss subject to budget constraints on the defense investments, while in the second class of games, each defender minimizes its cost of defense investment subject to upper bounds on the probability of successful attack on its assets (or its risk tolerance). Under suitable assumptions on the effectiveness of defense investments in reducing attack probabilities, we showed the existence of (generalized) Nash equilibria in both settings, and showed that each defender can compute its optimal defense allocation for a given allocation by other defenders by solving a convex optimization problem.

We demonstrated how our framework can be applied in diverse settings, including large-scale cyber-physical systems such as the power grid as well as enterprise networks. Motivated by recent cyber attacks that exploit vulnerabilities in assets supplied by third-party vendors, we specifically studied the impact of such vendors on the Nash equilibrium defense allocation in a case study on the IEEE 300 bus power grid network. As the probability of successful attack on the vendor increases, the defenders need to invest more to meet a given risk tolerance constraints. In a second case study on an e-commerce network, we computed optimal deployment of moving target defense using our framework.

Our framework leaves several interesting avenues for future research. The impact of incentive mechanisms, such as imposing fines on defenders or vendors who do not take adequate security measures, can be studied within our framework. Another important future direction is to consider real-time interaction between attacker(s) and defender(s) in a dynamic game framework. Such interaction can proceed in multiple stages and can consider various levels of misinformation about the strategies of different parties.

Acknowledgements We thank Dr. Shaunak Bopardikar (United Technologies Research Center) and Dr. Pratyusha Manadhata (HP Labs) for fruitful discussions.

References

[1] Alpcan T, Başar T (2010) Network security: A decision and game-theoretic approach. Cambridge University Press

[2] Amin S, Schwartz GA, Sastry SS (2013) Security of interdependent and identical networked control systems. Automatica 49(1):186–192

[3] Antonatos S, Akritidis P, Markatos EP, Anagnostakis KG (2007) Defending against hitlist worms using network address space randomization. Computer Networks 51(12):3471–3490

[4] Bagnoli M, Bergstrom T (2005) Log-concave probability and its applications. Economic Theory 26(2):445–469

[5] Boyd S, Vandenberghe L (2004) Convex optimization. Cambridge university Press

[6] Carroll TE, Crouse M, Fulp EW, Berenhaut KS (2014) Analysis of network address shuffling as a moving target defense. In: Communications (ICC), 2014 IEEE International Conference on, IEEE, pp 701–706

[7] Christie R (1993) Power systems test case archives. URL `https://goo.gl/1AOSXj`, retrieved: 2017-06-07

[8] Cook WJ, Cunningham WH, Pulleyblank WR, Schrijver A (1998) Combinatorial optimization, vol 605. Springer

[9] Durkota K, Lisỳ V, Bošanskỳ B, Kiekintveld C (2015) Approximate solutions for attack graph games with imperfect information. In: Decision and Game Theory for Security, Springer, pp 228–249

[10] Dutang C (2013) Existence theorems for generalized Nash equilibrium problems. Journal of Nonlinear Analysis and Optimization: Theory & Applications 4(2):115–126

[11] Fudenberg D, Levine DK (1998) The theory of learning in games, vol 2. MIT Press

[12] Gordon LA, Loeb MP (2002) The economics of information security investment. ACM Transactions on Information and System Security (TISSEC) 5(4):438–457

[13] Grant M, Boyd S, Ye Y (2008) CVX: Matlab software for disciplined convex programming

[14] Greenberg A (2017) 'Crash Overdrive': The malware that took down a power grid. URL `http://bit.ly/2raojOf`, Wired Magazine, retrieved: 2017-09-20

[15] Gupta A, Schwartz G, Langbort C, Sastry SS, Basar T (2014) A three-stage Colonel Blotto game with applications to cyberphysical security. In: American Control Conference (ACC), 2014, IEEE, pp 3820–3825

[16] Hogan WW (1973) Point-to-set maps in mathematical programming. SIAM Review 15(3):591–603

[17] Homer J, Zhang S, Ou X, Schmidt D, Du Y, Rajagopalan SR, Singhal A (2013) Aggregating vulnerability metrics in enterprise networks using attack graphs. Journal of Computer Security 21(4):561–597

[18] Hong JB, Kim DS (2016) Assessing the effectiveness of moving target defenses using security models. IEEE Transactions on Dependable and Secure Computing 13(2):163–177
[19] Hota A, Sundaram S (2016) Interdependent security games on networks under behavioral probability weighting. IEEE Transactions on Control of Network Systems (To Appear)
[20] Hota AR, Clements AA, Sundaram S, Bagchi S (2016) Optimal and game-theoretic deployment of security investments in interdependent assets. In: International Conference on Decision and Game Theory for Security, Springer, pp 101–113
[21] Israeli E, Wood RK (2002) Shortest-path network interdiction. Networks 40(2):97–111
[22] Jafarian JH, Al-Shaer E, Duan Q (2012) Openflow random host mutation: Transparent moving target defense using software defined networking. In: Proceedings of the first workshop on Hot topics in software defined networks, ACM, pp 127–132
[23] Jajodia S, Ghosh AK, Subrahmanian V, Swarup V, Wang C, Wang XS (2013) Moving target defense II. Application of Game Theory and Adversarial Modeling Series: Advances in Information Security 100:203
[24] Jiang L, Anantharam V, Walrand J (2011) How bad are selfish investments in network security? Networking, IEEE/ACM Transactions on 19(2):549–560
[25] Kunreuther H, Heal G (2003) Interdependent security. Journal of risk and uncertainty 26(2-3):231–249
[26] Laszka A, Felegyhazi M, Buttyan L (2014) A survey of interdependent information security games. ACM Computing Surveys (CSUR) 47(2):23:1–23:38
[27] Letchford J, Vorobeychik Y (2013) Optimal interdiction of attack plans. In: AAMAS, pp 199–206
[28] Lou J, Smith AM, Vorobeychik Y (2017) Multidefender security games. IEEE Intelligent Systems 32(1):50–60
[29] Modelo-Howard G, Bagchi S, Lebanon G (2008) Determining placement of intrusion detectors for a distributed application through Bayesian network modeling. In: International Workshop on Recent Advances in Intrusion Detection, Springer, pp 271–290
[30] Nzoukou W, Wang L, Jajodia S, Singhal A (2013) A unified framework for measuring a network's mean time-to-compromise. In: Reliable Distributed Systems (SRDS), 2013 IEEE 32nd International Symposium on, IEEE, pp 215–224
[31] Ok EA (2007) Real analysis with economic applications, vol 10. Princeton University Press
[32] Okhravi H, Hobson T, Bigelow D, Streilein W (2014) Finding focus in the blur of moving-target techniques. IEEE Security & Privacy 12(2):16–26
[33] Poolsappasit N, Dewri R, Ray I (2012) Dynamic security risk management using Bayesian attack graphs. IEEE Transactions on Dependable and Secure Computing 9(1):61–74
[34] Roberson B (2006) The Colonel Blotto game. Economic Theory 29(1):1–24

[35] Rosen JB (1965) Existence and uniqueness of equilibrium points for concave n-person games. Econometrica: Journal of the Econometric Society 33(3):520–534

[36] Sanger DE, Perlroth N (2016) A new era of internet attacks powered by everyday devices. URL `https://nyti.ms/2nsqrlT`, The New York Times, retrieved: 2017-05-14

[37] Schwartz G, Shetty N, Walrand J (2013) Why cyber-insurance contracts fail to reflect cyber-risks. In: Communication, Control, and Computing (Allerton), 2013 51st Annual Allerton Conference on, IEEE, pp 781–787

[38] Sreekumaran H, Hota AR, Liu AL, Uhan NA, Sundaram S (2015) Multi-agent decentralized network interdiction games. arXiv preprint arxiv:150301100

[39] Tambe M (2011) Security and game theory: Algorithms, deployed systems, lessons learned. Cambridge University Press

[40] Van Dijk M, Juels A, Oprea A, Rivest RL (2013) Flipit: The game of stealthy takeover. Journal of Cryptology 26(4):655–713

[41] Van Leeuwen B, Stout WM, Urias V (2015) Operational cost of deploying moving target defenses defensive work factors. In: Military Communications Conference, MILCOM 2015-2015 IEEE, IEEE, pp 966–971

[42] Wang L, Noel S, Jajodia S (2006) Minimum-cost network hardening using attack graphs. Computer Communications 29(18):3812–3824

[43] Wang L, Jajodia S, Singhal A, Cheng P, Noel S (2014) K-zero day safety: A network security metric for measuring the risk of unknown vulnerabilities. IEEE Transactions on Dependable and Secure Computing 11(1):30–44

[44] Zhang M, Wang L, Jajodia S, Singhal A, Albanese M (2016) Network diversity: A security metric for evaluating the resilience of networks against zero-day attacks. IEEE Transactions on Information Forensics and Security 11(5):1071–1086

Chapter Appendix: Generalized Nash Equilibrium

In this section, we give a formal definition of a *Generalized Nash Equilibrium* (GNE) and state the required existence result that will be useful in our analysis.

Let there be N players. The strategy set of player i is denoted as $X_i \subseteq \mathbb{R}^{n_i}$. Let $X := \Pi_{i=1}^{N} X_i$, and $X_{-i} := \Pi_{j=1, j\neq i}^{N} X_j$. Let $C_i : X_{-i} \to 2^{X_i}$ be the set-valued map or correspondence that defines the feasible strategy set for player i at a given strategy profile of all other players. Let $f_i : X \to \mathbb{R}$ denote the cost function for player i. We denote this game as $\Gamma(N, \{X_i\}, \{C_i\}, \{f_i\})$.

Definition 1. A strategy profile $\mathbf{x}^* \in X$ is a GNE of $\Gamma(N, \{X_i\}, \{C_i\}, \{f_i\})$ if for every player i,

$$x_i^* \in \underset{x_i \in C_i(x_{-i}^*)}{\operatorname{argmin}} f_i(x_i, x_{-i}^*). \tag{38}$$

Our proof of GNE existence in this chapter is based on the following general result.

Theorem 1. *Consider the game $\Gamma(N, \{X_i\}, \{C_i\}, \{f_i\})$. Assume for all players, we have*

1. *X_i is a nonempty, convex and compact subset of an Euclidean space,*
2. *C_i is both upper and lower semi-continuous,*
3. *$C_i(x_{-i})$ is nonempty, closed and convex for every $x_{-i} \in X_{-i}$,*
4. *f_i is continuous on the graph of C_i, and*
5. *$f_i(x_i, x_{-i})$ is quasiconvex on $C_i(x_{-i})$ for every $x \in X$.*

Then there exists a GNE.

The proof of the above theorem relies on Kakutani Fixed Point theorem and Berge's Maximum theorem, and is presented in [10, Theorem 3.1].

In many application, including for the defense cost minimization game studied in this chapter, we encounter a parametrized constraint set, i.e., $C_i(x_{-i}) = \{x_i \in X_i | g_{ij}(x_i, x_{-i}) \leq 0, j = \{1, 2, \ldots, m_i\}\}$. For this class of constraints, we have the following sufficient conditions for the upper and lower semi-continuity of C_i [16, Theorem 10,12].

Theorem 2. *Let $C_i : X_{-i} \to 2^{X_i}$ be given by $C_i(x_{-i}) = \{x_i \in X_i | g_{ij}(x_i, x_{-i}) \leq 0, j = \{1, 2, \ldots, m_i\}\}$.*

1. *Let $X_i \subseteq \mathbb{R}^{n_i}$ be closed, and all components g_{ij}'s be continuous on X. Then, C_i is upper semi-continuous on X_{-i}.*
2. *Let g_{ij}'s be continuous and convex in x_i for each x_{-i}. If there exists $\bar{x}$ such that $g_{ij}(\bar{x}_i, \bar{x}_{-i}) < 0$ for all j, then C_i is lower semi-continuous at $\bar{x}_{-i}$, and in some neighborhood of $\bar{x}_{-i}$.*

Remark 5. Some authors use the term hemicontinuity instead of semi-continuity [31]. The definitions coincide for closed and compact-valued correspondences, which is the case here.