Optimal and Game-Theoretic Deployment of Security Investments in Interdependent Assets

Ashish R. Hota$^{(\boxtimes)}$, Abraham A. Clements,
Shreyas Sundaram, and Saurabh Bagchi

School of Electrical and Computer Engineering,
Purdue University, West Lafayette, USA
{ahota,clemen19,sundara2,sbagchi}@purdue.edu

Abstract. We introduce a game-theoretic framework to compute optimal and strategic security investments by multiple defenders. Each defender is responsible for the security of multiple assets, with the interdependencies between the assets captured by an *interdependency graph*. We formulate the problem of computing the optimal defense allocation by a single defender as a convex optimization problem, and establish the existence of a pure Nash equilibrium of the game between multiple defenders. We apply our proposed framework in two case studies on interdependent SCADA networks and distributed energy resources, respectively. In particular, we investigate the efficiency loss due to decentralized defense allocations.

AQ1
AQ2

1   Introduction

Modern critical infrastructures have a large number of interdependent assets, operated by multiple stakeholders each working independently to maximize their own economic benefits. In these cyber-physical systems, interdependencies between the assets owned by different stakeholders have significant implications on the reliability and security of the overall system. For instance, in the electricity grid, industrial control systems at the power generator are managed by a different entity (the generator) than the smart meters deployed by the distribution utility companies. If certain components of these assets are from a common vendor, then a sophisticated attacker can exploit potential shared vulnerabilities and compromise the assets managed by these different entities [21].

Security interdependencies are often modeled in varying degrees of abstractions. While the *attack graph* formalism [7] captures detailed models of how an attacker might exploit vulnerabilities within an enterprise network, representations of interdependencies in large-scale cyber-physical networks, such as the

electric grid, are often captured in terms of coupled dynamical systems [12]. In addition to the interdependencies, individual stakeholders are often myopic and resource constrained, which makes identification and mitigation of vulnerabilities in a large number of cyber and physical assets prohibitively expensive. Furthermore, decentralized deployment of defense strategies by these self-interested defenders often leads to increased security risks for the entire system.

In this paper we present a systematic framework that can be used to efficiently compute optimal defense allocations under interdependencies. We model the network security problem as a game between multiple defenders, each of whom manages a set of assets. The interdependencies between these assets are captured by an *interdependency graph*. Each defender minimizes her own expected loss, where the attack probabilities of her assets are a function of her own defense strategies, strategies of other defenders, and the interdependency graph. In particular, attacker(s) are assumed to exploit the interdependencies to target valuable assets in the network. We first establish the existence of a pure Nash equilibrium in the game between self-interested defenders. For a general class of defense strategies, we show that the problem of computing an optimal defense allocation for a defender (i.e., her *best response*) is equivalent to solving a convex optimization problem.

We evaluate the inefficiency of decentralized decision-making in two case studies; the first is a SCADA system with multiple control networks managed by independent entities, and the second is a distributed energy resource failure scenario identified by the US National Electric Sector Cybersecurity Organization Resource (NESCOR). In both settings, we find that when entities have similar risks but disparate budgets, the total expected loss at a Nash equilibrium can be much larger than the total expected loss under the socially optimal solution. Furthermore, we show that it can be in the interest of a selfish actor to defend assets that belong to another entity due to mutual interdependencies.

1.1 Related Work

Security games on networks with multiple defenders have recently been considered within the broad framework of *Stackelberg security games* [9,22]. A Stackelberg security game is defined as an extensive form leader-follower game where a defender randomizes her defense allocations across multiple targets and an attacker observes the randomized strategies and chooses the target with highest successful attack probability. Several papers have considered multiple defenders and network interdependencies within this framework [14–16]. A recurring assumption in these papers is that the strategy space of a defender is discrete, e.g., a node is either fully protected or is vulnerable. In contrast, we consider defense strategies that are continuous variables. In addition, our work is related to recent explorations of attack graph games [3], though the defense strategies considered in that paper are very different from the ones explored here.

Our work is also related to the substantial body of literature on *interdependent security games*; [13] contains a comprehensive review. A common feature in this line of work is that each node is an independent decision maker, i.e.,

a player is responsible for the defense of a single node in the graph. We relax this assumption in this paper. In our formulation, a player is responsible for the defense of multiple assets (nodes) in the (interdependency) graph.

Our game-theoretic formulation and analysis borrows ideas and techniques from the literature on *network interdiction games* [8]. In the classical shortest path interdiction game [8], there is an underlying network; an attacker aims to find a path of shortest length from a given source to a target, while the defender can *interdict* the attack by increasing the lengths of the paths. Extensions to cases where multiple attackers and/or defenders operate on a given network are few, with the exception of our recent work [19]. The model we propose in this paper generalizes the formulation in [19] as we consider defenders who defend multiple nodes and with possibly nonlinear cost functions. Finally, our paper has a similar perspective as [5] as we develop a systems-theoretic framework that is readily applicable in a broad class of interdependent network security settings.

2 Security Game Framework

Interdependency Graph: We represent the assets in a networked (cyber-physical) system as nodes of a directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, i.e., each node $v_i \in \mathcal{V}$ represents an asset. The presence of a directed edge $(v_j, v_i) \in \mathcal{E}$ indicates that when the asset v_j is compromised, it can be used to launch an attack on asset v_i. This attack succeeds with a probability $p_{j,i} \in (0, 1]$, independent of analogous attack probabilities defined on the other edges. Without loss of generality, let s be the source node from which an attacker launches the attack from outside the network.[1] We refer to such a graph as an *interdependency graph*.[2]

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

Strategic Defenders: Let $\mathcal{D}$ be the set of defenders. A defender $D_k \in \mathcal{D}$ is responsible for the security of a set $\mathcal{V}_k \subseteq \mathcal{V} \setminus \{s\}$ of assets. For each asset $v_m \in \mathcal{V}_k$, there is a financial loss $L_m \in \mathbb{R}_{\geq 0}$ that defender D_k incurs if v_m gets compromised. The defender can allocate its resources to reduce the attack probabilities on the edges interconnecting different assets on the interdependency graph, subject to certain constraints. We denote the feasible (defense) strategy set of defender D_k as $\mathcal{X}_k \subset \mathbb{R}_{\geq 0}^{n_k}$, where $n_k < \infty$. We require that $\mathcal{X}_k$ is non-empty, compact and convex. The defense resources reduce the attack success probabilities on the edges. We will discuss the exact transformation of defense allocation into the reduction of attack probabilities in the next subsection.

---

[1] If there are multiple entry points to the network, we can add a source node s and add edges from s to all entry points with attack probabilities equal to 1.

[2] Interdependency graphs also capture essential features of *attack graphs* [3,7] where the nodes represent intermediate steps in multi-stage attacks.

Now, let $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_{|\mathcal{D}|}]$ be a joint defense strategy of the defenders, with $\mathbf{x}_k \in \mathcal{X}_k$ for every defender D_k. The attack success probability of an edge (v_j, v_i) under this joint defense strategy is denoted as $\hat{p}_{j,i}(\mathbf{x})$. The goal of each defender D_k is to minimize the cost function given by

$$C_k(\mathbf{x}) \triangleq \sum_{v_m \in \mathcal{V}_k} L_m \left(\max_{P \in \mathcal{P}_m} \prod_{(v_j, v_i) \in P} \hat{p}_{j,i}(\mathbf{x}) \right), \tag{1}$$

subject to $\mathbf{x}_k \in \mathcal{X}_k$. In other words, a defender minimizes her expected loss, where the probability of loss of an asset is given by the highest probability of attack on any path from the source to that asset on the interdependency graph.

Strategic Attacker(s): Cyber-physical systems in the field face multiple strategic adversaries with different objectives, capabilities and knowledge about the system. As a result, detailed modeling of strategic attackers is challenging. Nonetheless, a defender must be able to assess her security risks and allocate defense resources under inadequate information about the attackers. This motivates our choice of minimizing the worst case attack probabilities on an asset in (1), which 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. The defender can assess her risk profile against attackers of different capabilities by appropriately varying the probabilities on each edge.

As an example of a setting that can be modeled within our framework, consider the SCADA based control system shown in Fig. 1. There are two control subsystems, with interdependencies due to a shared corporate network and a common vendor for the remote terminal units (RTUs). Figure 2 shows the resulting interdependency graph. We further discuss this setting in Sect. 3.
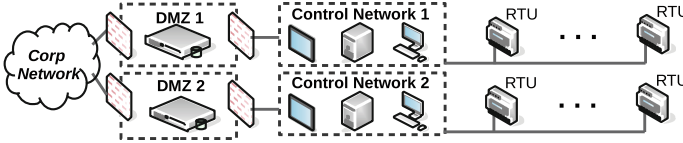

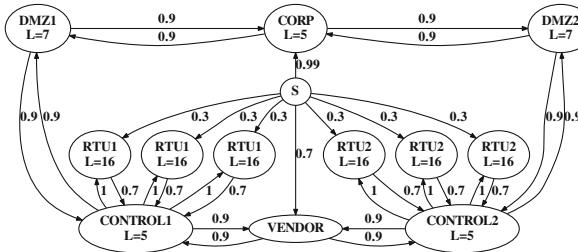**Fig. 1.** A SCADA system diagram of two interacting control systems.

Fig. 2. Interdependency graph for the SCADA system in Fig. 1.

2.1 Defense Strategies

As noted above, the defense resources reduce the attack probabilities on the edges of the interdependency graph. Accordingly, we introduce a transformation matrix $T_k : \mathbb{R}^{n_k} \to \mathbb{R}^{|\mathcal{E}|}$ which maps a feasible defense strategy $\mathbf{x}_k$ to a defense allocation on edges. By appropriately defining the matrix T_k, we can capture very general class 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 $\mathcal{E}_k \subseteq \mathcal{E}$ of the graph $\mathcal{G}$, and accordingly $n_k = |\mathcal{E}_k|$. For example, $\mathcal{E}_k$ can represent all the edges that are incoming to a node in $\mathcal{V}_k$, i.e., defender D_k can reduce the attack probabilities of all the edges that are incoming to the nodes under its ownership. Furthermore, an edge can potentially be managed by multiple defenders. Under edge-based defense scenarios, we will define the feasible strategy space of a defender D_k as $\mathcal{X}_k := \{x_{j,i}^k \in \mathbb{R}_{\geq 0}, (v_j, v_i) \in \mathcal{E}_k | \sum_{(v_j, v_i) \in \mathcal{E}_k} x_{j,i}^k \leq B_k\}$, where B_k is the total defense budget for defender D_k. In this case, T_k has a sub-matrix which is an identity matrix of dimension $|\mathcal{E}_k|$ and the other entities are equal to 0. An example of edge-based defense is when a device inspects the incoming traffic depending on the traffic source.

Node-based defense strategy: In this case, a defender D_k allocates defense resources to the set of nodes in $\mathcal{V}_k$, and accordingly, $n_k = |\mathcal{V}_k|$. 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 $\mathcal{G}$ have a defense allocation x_i^k, i.e., $x_{j,i}^k = x_i^k$ for every $(v_j, v_i) \in \mathcal{E}$. Node-based defense strategy is motivated by *moving target defense* techniques [10]. Here x_i potentially represents the rate at which system configurations (such as the IP-address) of a node v_i are being updated. Here T_k maps the allocation on a node into the edges that are incoming to it.

We now define the *length* or distance of an edge (v_j, v_i) in terms of the attack probability as,

$$l_{j,i} \triangleq -\log(p_{j,i}) \geq 0. \tag{2}$$

A higher probability of attack on an edge leads to smaller length for the edge.

In this paper, we assume that the defense allocations on an edge linearly increase the length of the edge. Mathematically, let $x_{j,i} = \sum_{(v_j, v_i) \in \mathcal{E}_k} x_{j,i}^k = \sum_{D_k \in \mathcal{D}} [T_k \mathbf{x}_k]_{(j,i)}$ denote the total defense allocation by all the defenders on the edge (v_j, v_i). Then, the modified length of the edge under a joint strategy profile $\mathbf{x}$ is given by

$$\hat{l}_{j,i}(\mathbf{x}) \triangleq l_{j,i} + x_{j,i} \implies -\log(\hat{p}_{j,i}(\mathbf{x})) = l_{j,i} + x_{j,i} \tag{3}$$

$$\implies \hat{p}_{j,i}(\mathbf{x}) \triangleq p_{j,i} e^{-x_{j,i}}, \tag{4}$$

i.e., the total defense allocation on an edge $x_{j,i}$ leads to a relative reduction of the corresponding attack success probability given by $e^{-x_{j,i}}$. This captures diminishing effectiveness of defense allocations and leads to a tractable formulation of the cost minimization problem (1). We denote the vector of modified lengths of the graph under joint defense strategy $\mathbf{x}$ as $\hat{\mathbf{L}}(\mathbf{x}) = \mathbf{L} + \sum_{D_k \in \mathcal{D}} T_k \mathbf{x}_k$, where $\mathbf{L}$ is the vector of lengths in the absence of any defense allocation, given by (2).

2.2   Existence of a Pure Nash Equilibrium (PNE)

We first show the existence of a PNE in the game between multiple defenders, each with a defender-specific transformation matrix T_k.

Proposition 1. *The strategic game with multiple defenders where a defender minimizes her cost defined in* (1) *possesses a pure Nash equilibrium.*

Proof. From our transformation of attack probabilities into lengths on edges given in (3) and (4), the probability of successful attack on a node $v_m \in \mathcal{V}_k$ due to a path $P \in \mathcal{P}_m$ and joint defense strategy $\mathbf{x}$ is equal to

$$\prod_{(u_j,u_i)\in P} \hat{p}_{j,i}(\mathbf{x}) = \exp\left(-\sum_{(v_j,v_i)\in P}\left[l_{j,i} + \sum_{D_r\in\mathcal{D}}[T_r\mathbf{x}_r]_{(j,i)}\right]\right),$$

where $\exp(\cdot)$ is the exponential function, i.e., $\exp(z) = e^z$. Accordingly, we can express the cost function of a defender D_k, defined in (1), as a function of her strategy $\mathbf{x}_k$ and the joint strategy of other defenders $\mathbf{x}_{-k}$ as

$$C_k(\mathbf{x}_k,\mathbf{x}_{-k}) = \sum_{v_m\in\mathcal{V}_k} L_m\exp\left(-\min_{P\in\mathcal{P}_m}\sum_{(v_j,v_i)\in P}\left[\hat{l}_{j,i}(\mathbf{x}_{-k}) + [T_k\mathbf{x}_k]_{(j,i)}\right]\right),\quad (5)$$

where $\hat{l}_{j,i}(\mathbf{x}_{-k}) = l_{j,i} + \sum_{D_r\in\mathcal{D},D_r\neq D_k}[T_r\mathbf{x}_r]_{(j,i)}$ for an edge (v_j,v_i).

Note that $\sum_{(v_j,v_i)\in P}\left[\hat{l}_{j,i}(\mathbf{x}_{-k}) + [T_k\mathbf{x}_k]_{(j,i)}\right]$ is an affine and, therefore, a concave function of $\mathbf{x}_k$. The minimum of a finite number of concave functions is concave [1]. 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, $C_k(\mathbf{x}_k,\mathbf{x}_{-k})$ is convex in $\mathbf{x}_k$ for any given $\mathbf{x}_{-k}$. Furthermore, the feasible strategy set $\mathcal{X}_k$ 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 [18].   □

2.3   Computing the Best Response of a Defender

Let $\mathbf{x}_{-k}$ be the joint defense strategy of all defenders other than D_k. Then the *best response* of D_k is a strategy $\mathbf{x}_k^* \in \mathcal{X}_k$ which minimizes her cost $C_k(\mathbf{x}_k,\mathbf{x}_{-k})$ defined in (1). Let $\hat{\mathbf{L}}(\mathbf{x}_{-k}) = \mathbf{L} + \sum_{D_r\in\mathcal{D},r\neq k}T_r\mathbf{x}_r$ be the vector of edge lengths under defense allocation $\mathbf{x}_{-k}$. We show that $\mathbf{x}_k^*$ can be computed by solving the following convex optimization problem:

$$\underset{y\in\mathbb{R}^{|\mathcal{V}|},x_k\in\mathbb{R}^{n_k}}{\text{minimize}} \qquad \sum_{v_m\in\mathcal{V}_k} L_m e^{-y_m} \qquad (6)$$

$$\text{subject to} \qquad \mathcal{I}y - T_k\mathbf{x}_k \leq \hat{\mathbf{L}}(\mathbf{x}_{-k}), \qquad (7)$$

$$y_s = 0, \qquad (8)$$

$$\mathbf{x}_k \in \mathcal{X}_k, \qquad (9)$$

where $\mathcal{I}$ is the node-edge incidence matrix of the graph $\mathcal{G}$. Note that the constraint in (7) is affine. This formulation is motivated by similar ideas explored in the shortest path interdiction games literature [8,19].

We refer to the vector $\{y_u\}_{u \in \mathcal{V}}$ as a *feasible potential* if it satisfies (7) for every edge in the graph. In graphs without a negative cycle, the well known Bellman-Ford algorithm for shortest paths corrects the inequality in (7) for an edge in every iteration and terminates with a feasible potential [2]. In our setting, the length of every edge is nonnegative. We now prove the following result.

Proposition 2. *A defense strategy* $\mathbf{x}_k^* \in \mathcal{X}_k$ *is the optimal solution of the problem defined in Eqs.* (6)–(9) *if and only if it is the minimizer of* $C_k(\mathbf{x}_k, \mathbf{x}_{-k})$ *defined in* (1).

Proof. Consider a feasible defense allocation vector $\mathbf{x}_k \in \mathcal{X}_k$. The joint strategy profile $(\mathbf{x}_k, \mathbf{x}_{-k})$ defines a modified length vector $\hat{\mathbf{L}}(\mathbf{x}_k, \mathbf{x}_{-k}) = \hat{\mathbf{L}}(\mathbf{x}_{-k}) + T_k \mathbf{x}_k$ on the edges of $\mathcal{G}$. Now consider a feasible potential $\{y_u\}_{u \in \mathcal{V}}$ which satisfies (7). A feasible potential exists, since the vector $y_u = 0$ for every $u \in \mathcal{V}$ satisfies (7).

Now consider a P from s to a node $v_m \in \mathcal{V}_k$. Then, the feasible potential at node v_m satisfies $y_{v_m} - y_s = y_{v_m} \leq \sum_{(u_j, u_i) \in P} \hat{l}_{j,i}(\mathbf{x}_k, \mathbf{x}_{-k})$. In other words, y_{v_m} is a lower bound on the length of every path (and consequently the shortest path) from s to v_m. Furthermore, in the absence of negative cycles, there always exists a feasible potential where y_{v_m} is *equal* to the length of the shortest path from s to v_m [2, Theorem 2.14] (the solution of the Bellman-Ford algorithm).

Now let $\{\mathbf{x}_k^*, \{y_u^*\}_{u \in \mathcal{V}}\}$ be the optimal solution of the problem defined in Eqs. (6)–(9) for a given $\mathbf{x}_{-k}$. The length of every edge (u_j, u_i) at the optimal defense allocation $\mathbf{x}_k^*$ is given by $\hat{l}_{j,i}(\mathbf{x}_k^*, \mathbf{x}_{-k})$. 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 $L_m > 0$. 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 [2, Theorem 2.14]) we know that there exists a feasible potential $\{\hat{y}_u\}_{u \in \mathcal{V}}$ such that $\hat{y}_{v_m}$ is equal to the length of the shortest path from s to v_m for every node $v_m \in \mathcal{V}_k$ with length of every edge (u_j, u_i) given by $\hat{l}_{j,i}(\mathbf{x}_k^*, \mathbf{x}_{-k})$. 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 $\{\mathbf{x}^*, \{y_u^*\}_{u \in \mathcal{V}}\}$.

Let P be a path from s to v_m in the optimal solution, and let P^* be a path of shortest length. The length of this path is given by

$$y_{v_m}^* \leq \sum_{(u_j, u_i) \in P} \hat{l}_{j,i}(\mathbf{x}_k^*, \mathbf{x}_{-k}) = - \sum_{(u_j, u_i) \in P} \log(\hat{p}_{j,i}(\mathbf{x}^*))$$
$$\implies e^{-y_{v_m}^*} \geq \prod_{(u_j, u_i) \in P} \hat{p}_{j,i}(\mathbf{x}^*),$$

with equality for the path P^*. Therefore the optimal cost of the problem defined in Eqs. (6)–(9) is equal to the cost in (1). ☐

As a result, a defender can efficiently (up to any desired accuracy) compute her optimal defense allocation given the strategies of other defenders. Furthermore, the problem of social cost minimization, where a central planner minimizes

the sum of expected losses of all defenders, can be represented in a form that is analogous to Eqs. (6)–(9) and can be solved efficiently.

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 (5) is non-differentiable and we cannot apply gradient-based update schemes. Second, in the equivalent formulation Eqs. (6)–(9), the players' cost minimization problems are coupled through their constraints. As a result, the problem belongs to the class of *generalized Nash equilibrium problems* [4], which has very few general convergence results. We leave further theoretical investigations of convergence of different dynamics to PNE for future work.

3   Numerical Case Studies

We apply our proposed framework in two case studies. Our goal is to understand the loss of welfare due to decentralized decision making by the defenders with asymmetric defense budgets compared to the socially optimal defense allocation. The social optimum corresponds to the solution computed by a central authority as it minimizes the total expected loss of all the players. The ratio of the highest total expected loss at any PNE and the total expected loss at the social optimum is often used as a metric (*Price of Anarchy*) to quantify the inefficiency of Nash equilibrium. We consider PNE strategy profiles obtained by iteratively computing best responses of the players; the sequence of best response strategies converged to a PNE in all of the experiments in this section. We use the MATLAB tool CVX [6] for computing the best response of a defender and the social optimum. In both the experiments, we randomly initialize the attack success probabilities on the edges of the respective interdependency graphs.

3.1   An Interdependent SCADA Network with Two Utilities

We first consider the SCADA network shown in Fig. 1, based on NIST's guidelines for industrial control systems [20]. As discussed earlier, there are two control subsystems with interdependencies due to a shared corporate network and vendors for RTUs. Each subsystem is owned by a different defender. The resulting interdependency graph is shown in Fig. 2. The number in the name of a node indicates the defender who owns it and the amount of loss to its owner, if it is compromised. The corporate network is owned by both defenders. The compromise of the control network causes loss of the RTUs, and as a result, the corresponding edges have an attack success probability 1 and are indefensible.
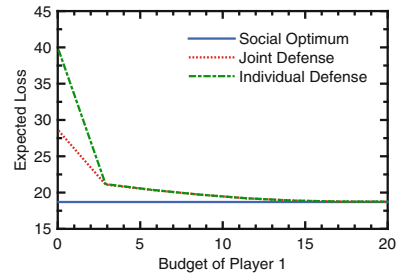
In our experiments, we keep the total defense budget fixed for the overall system, and compare the resulting total expected loss (under this total budget) at the social optimum with the expected loss that arises at a PNE when each subsystem is defended independently. We consider an edge-based defense strategy for all our results. In the decentralized defense case, we consider two scenarios. First, the defenders can only defend edges that are incoming to a node

under their ownership. We refer to this scenario as *individual defense*. Second, the defenders can *jointly defend* all the edges in the interdependency network, i.e., a defender can defend an edge within the subsystem of the other defender.

We plot our results in Fig. 3a for a SCADA network where each utility has 3 RTUs. The total budget is 20, and we vary the budget of defender 1 as shown in the x-axis of the plot. Defender 2 receives the remaining amount (20 minus the budget of defender 1). We observe that the joint defense case leads to a smaller total expected loss compared to the individual defense case at the respective PNEs. The difference between the two cases is most significant when the budgets of the two defenders are largely asymmetric. Our results show that it is beneficial for a selfish decision maker with a large budget to defend parts of the network owned by another defender with a smaller budget in presence of interdependencies. As the asymmetry in budgets decreases, the expected losses under joint defense and the individual defense approach the social optimum. This is because the network considered is symmetric for the defenders. In Fig. 3b, we plot analogous results when each utility has 30 RTUs with a total budget 40, and observe similar trends in the respective expected losses.
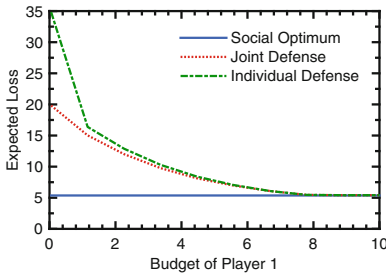
(a) 3 RTUs per subsystem   (b) 30 RTUs per subsystem

Fig. 3. Comparison of total expected loss under the social optimal defense allocation with the PNE strategies under joint and individual defense scenarios. The total budgets across both defenders are 20 and 40, respectively.

3.2   Evaluation of a Distributed Energy Resource Failure Scenario

In our second case study, we consider a distributed energy resource failure scenario, DER.1, identified by the US National Electric Sector Cybersecurity Organization Resource (NESCOR) [17]. We build upon the recent work by [11], where the authors developed a tool *CyberSAGE*, which represents NESCOR failure scenarios as a security argument graph to identify the interdependencies between different attack steps. We reproduce the security argument graph for the DER.1 scenario in Fig. 4. The authors of [11] note that applying all mitigations for the DER.1 failure scenario can be expensive. Our framework enables computing the optimal (and PNE) defense strategy under budget constraints.

Note that the nature of interdependency in Fig. 4 is qualitatively different from the setting in the previous subsection. In Fig. 4, the nodes in the interdependency graph capture individual attack steps (similar to the representation in attack graphs). In contrast, the nodes in Fig. 2 correspond to disparate devices in the SCADA network. Furthermore, multiple attack steps can occur within a single device; all the intermediate nodes that belong to a common device are shown to be within a box in Fig. 4. For example, nodes $w3, w4, w5, w6, w7$ belong to the Human-Machine Interface (HMI) of the photovoltaic (PV) system. The node S represents the entry point of an attack, the nodes G_0 and G_1 represent the final target nodes that compromise the PV and electric vehicle (EV) components of the DER. A more detailed description is available in [11].
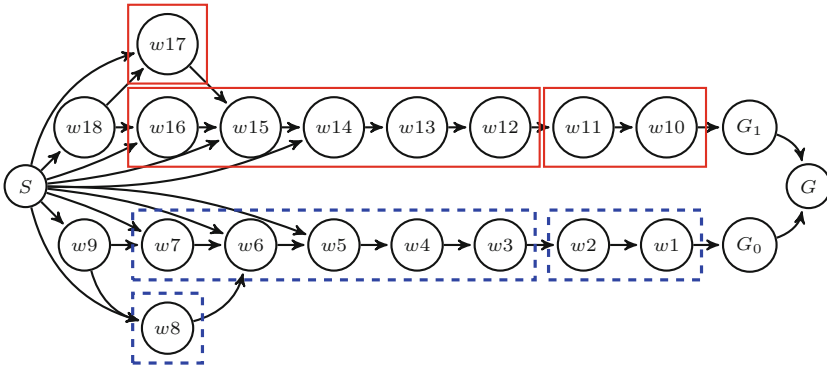

Fig. 4. Interdependency Graph of NESCOR DER.1 failure scenario [11]

We treat the security argument graph (Fig. 4) as the interdependency graph, and compute the globally optimal and Nash equilibrium strategies for two classes of defense strategies, (i) edge-based defense, where a player defends every edge independently, and (ii) device-based defense (such as IP-address randomization), where each device receives a defense allocation. In the second case, all the incoming edges to the nodes that belong to a specific device receive identical defense allocations. In the decentralized case, there are two players, who aim to protect nodes G_0 and G_1, respectively. In addition, each player experiences an additional loss if the other player is attacked successfully. This is captured by adding the extra node G which has edges from both G_0 and G_1 with attack probabilities equal to 1. Both players experience a loss if node G is compromised.

We plot the ratio of total expected losses under the socially optimal and PNE strategy profiles, for both edge-based and device-based defense strategies, in Figs. 5a and b, respectively. As the figures show, at a given total budget, the ratio of the expected losses at the social optimum and at a PNE is smaller when there is a larger asymmetry in the budgets of the individual players. In other words, when the individual players have similar defense budgets, the total
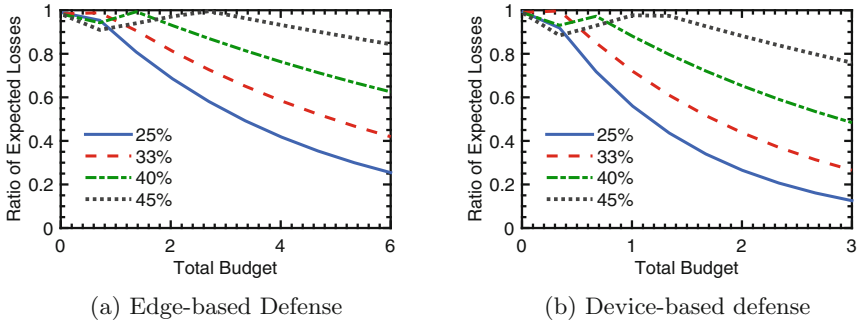
(a) Edge-based Defense   (b) Device-based defense

Fig. 5. The ratio of total expected losses of both defenders at the social optimum and a PNE in the DER.1 failure scenario under two different defense strategies. The total budget is divided among the two defenders, and defender 1 receives a percentage of the total budget as indicated in the legend.

expected loss at Nash equilibrium is not much larger than the expected loss under a globally optimal defense strategy.

4   Discussion and Conclusion

We presented a game-theoretic framework that enables systematic analysis of security trade-offs in interdependent networked systems. For a general class of defense strategies, the computation of optimal defense allocation for a defender is equivalent to solving a convex minimization problem. We also proved the existence of a pure Nash equilibrium for the game between multiple defenders. The SCADA network and DER.1 case studies illustrate how our framework can be used to study the security of interdependent systems at different levels of abstraction, from individual attack steps in the DER.1 scenario to an entire organization (vendor in the SCADA example) being abstracted to a single node. Our framework can be readily applied in practice by individual stakeholders to evaluate the effectiveness of different defense strategies and share information with other defenders to decide when and to what degree cooperative defense should be applied. The different levels of abstractions enable the creation of models with the available information a defender has. For example, the SCADA use case could be used to identify the degree to which the compromise of the vendor will affect the security of a system. This could translate into adding security requirements in procurement contracts with the vendors. In future, we will apply our framework in large-scale cyber-physical systems. Establishing convergence guarantees for best response dynamics and theoretical characterizations of inefficiencies at Nash equilibria remain as challenging open questions.

References

1. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press, New York (2004)
2. Cook, W.J., Cunningham, W.H., Pulleyblank, W.R., Schrijver, A.: Combinatorial Optimization, vol. 605. Springer, Heidelberg (1998)
3. Durkota, K., Lisý, V., Bošanský, B., Kiekintveld, C.: Approximate solutions for attack graph games with imperfect information. In: Khouzani, M.H.R., Panaousis, E., Theodorakopoulos, G. (eds.) GameSec 2015. LNCS, vol. 9406, pp. 228–249. Springer, Heidelberg (2015). doi:10.1007/978-3-319-25594-1_13
4. Facchinei, F., Kanzow, C.: Generalized nash equilibrium problems. Ann. Oper. Res. **175**(1), 177–211 (2010)
5. Fielder, A., Panaousis, E., Malacaria, P., Hankin, C., Smeraldi, F.: Game theory meets information security management. In: Cuppens-Boulahia, N., Cuppens, F., Jajodia, S., Abou El Kalam, A., Sans, T. (eds.) SEC 2014. IAICT, vol. 428, pp. 15–29. Springer, Heidelberg (2014). doi:10.1007/978-3-642-55415-5_2
6. Grant, M., Boyd, S., Ye, Y.: CVX: Matlab software for disciplined convex programming (2008)
7. Homer, J., Zhang, S., Ou, X., Schmidt, D., Du, Y., Rajagopalan, S.R., Singhal, A.: Aggregating vulnerability metrics in enterprise networks using attack graphs. J. Comput. Secur. **21**(4), 561–597 (2013)
8. Israeli, E., Wood, R.K.: Shortest-path network interdiction. Networks **40**(2), 97–111 (2002)
9. Jain, M., Conitzer, V., Tambe, M.: Security scheduling for real-world networks. In: AAMAS, pp. 215–222 (2013)
10. Jajodia, S., Ghosh, A.K., Subrahmanian, V., Swarup, V., Wang, C., Wang, X.S.: Moving Target Defense II. Application of Game Theory and Adversarial Modeling. Advances in Information Security, vol. 100, p. 203. Springer, New York (2013)
11. Jauhar, S., Chen, B., Temple, W.G., Dong, X., Kalbarczyk, Z., Sanders, W.H., Nicol, D.M.: Model-based cybersecurity assessment with nescor smart grid failure scenarios. In: 21st Pacific Rim International Symposium on Dependable Computing, pp. 319–324. IEEE (2015)
12. Kundur, D., Feng, X., Liu, S., Zourntos, T., Butler-Purry, K.L.: Towards a framework for cyber attack impact analysis of the electric smart grid. In: IEEE SmartGridComm, pp. 244–249 (2010)
13. Laszka, A., Felegyhazi, M., Buttyan, L.: A survey of interdependent information security games. ACM Comput. Surv. (CSUR) **47**(2), 23:1–23:38 (2014)
14. Letchford, J., Vorobeychik, Y.: Computing randomized security strategies in networked domains. In: Applied Adversarial Reasoning and Risk Modeling 2011, vol. 06 (2011)
15. Letchford, J., Vorobeychik, Y.: Optimal interdiction of attack plans. In: AAMAS, pp. 199–206 (2013)
16. Lou, J., Smith, A.M., Vorobeychik, Y.: Multidefender security games. arXiv preprint arXiv:1505.07548 (2015)
17. Electric sector failure scenarios and impact analyses, National Electric Sector Cybersecurity Organization Resource, EPRI (2014)
18. Rosen, J.B.: Existence and uniqueness of equilibrium points for concave n-person games. Econometrica: J. Econometric Soc. **33**(3), 520–534 (1965)
19. Sreekumaran, H., Hota, A.R., Liu, A.L., Uhan, N.A., Sundaram, S.: Multi-agent decentralized network interdiction games. arXiv preprint arXiv:1503.01100 (2015)

20. Stouffer, K.: Guide to industrial control systems (ICS) security. NIST special publication 800-82, 16-16 (2011)
21. Emerging threat: Dragonfly/Energetic Bear - APT Group (2014). http://www.symantec.com/connect/blogs/emerging-threat-dragonfly-energetic-bear-apt-group, Symantec Official Blog. Accessed 15 Aug 2016
22. Tambe, M.: Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned. Cambridge University Press, New York (2011)