

Network Inspection Using Heterogeneous Sensors for Detecting Strategic Attacks

Bobak McCann
 Georgia Institute of Technology
bmccann6@gatech.edu

Mathieu Dahan
 Georgia Institute of Technology
mathieu.dahan@isye.gatech.edu

Abstract

We consider a two-player network inspection game, in which a defender allocates sensors with potentially heterogeneous detection capabilities in order to detect multiple attacks caused by a strategic attacker. The objective of the defender (resp. attacker) is to minimize (resp. maximize) the expected number of undetected attacks by selecting a potentially randomized inspection (resp. attack) strategy. We analytically characterize Nash equilibria of this large-scale zero-sum game when every vulnerable network component can be monitored from a unique sensor location. We then leverage our equilibrium analysis to design a heuristic solution approach based on minimum set covers for computing inspection strategies in general. Our computational results on a benchmark cyber-physical distribution network illustrate the performance and computational tractability of our solution approach.

1. Introduction

Critical infrastructure networks such as electric, gas, and water distribution systems are paramount for the well-being of society. However, these networks regularly face random disruptions as well as attacks from strategic adversaries [1, 2]. In particular, recent incidents have demonstrated that adversarial attackers can disrupt or gain control of the cyber-physical systems deployed in these networks by exploiting cyber insecurities or physical faults. A most recent example is the cyberattack against a major US fuel pipeline, which caused disruptions in the fuel supply of the Eastern United States [3]. Additional examples can be found in [4, 5].

A key part of any defense strategy is to detect attacks using sensors positioned in various locations that continuously monitor the network. If a network is small, this can be done easily by placing a sensor at each location of interest. However, for medium or large networks, it can be infeasible to position a sensor at

every location. Thus the problem of how to strategically position a restricted number of sensors is crucial.

We employ a game-theoretic approach to study this problem. Game theory has successfully been used to study problems in the domain of cybersecurity (and network security more broadly) [6, 7, 8, 9, 10, 11, 12, 13]. In particular, it has proven successful for sensor allocation problems [13, 14, 15]. In our model, the defender allocates heterogeneous sensors in order to detect multiple attacks caused by a strategic attacker. The sensors may differ in their detection accuracies, which typically depend on the sensing technology utilized. The defender (resp. the attacker) aims to minimize (resp. maximize) the expected number of undetected attacks. Thus we model the interactions between both players using a zero-sum game, in which both players may potentially select randomized strategies. This feature is known to be desirable in security settings in which finite resources are allocated [14, 16].

Previous simultaneous security models, such as in [17, 18, 19, 20], assume that each detection device is homogeneous. In this work, we extend the model in [14] by accounting for the potential heterogeneity in detection accuracy of the sensors available to the defender. In particular, we study how the detection heterogeneity of the defender's resources affects the strategies of both players.

We study the mixed Nash Equilibria (NE) of this game. As this is a zero-sum game, NE can be computed by solving a linear program [21]. However, as the network's size increases, this linear program becomes too computationally expensive to solve because of the combinatorial nature of the players' action sets. Thus, we analyze equilibrium properties under certain conditions, and leverage our results to provide a computationally tractable heuristic solution approach that computes inspection strategies in the general case with good detection performance.

Our contributions are twofold: First, we analytically solve the game and provide equilibrium properties when

each component in the network is monitored from a unique sensor location. These results provide us with valuable insight regarding the impact of the detection accuracies, number of attacks, and network topology on the players' equilibrium strategies. Second, we leverage our equilibrium results to design a heuristic solution approach for computing inspection strategies in general. Our approach is based on solutions to a minimum set cover problem, which have been shown to be effective for different inspection games [10, 11, 12, 13]. We then conduct a computational study on a benchmark cyber-physical distribution network and empirically validate the performance and computational tractability of our solution approach.

The paper is structured as follows. In Section 2, we introduce the network inspection game. In Section 3, we derive equilibrium properties and solve the game when each component is monitored from a unique sensor location. We then present in Section 4 our heuristic approach for computing inspection strategies in the general case and provide computational results to validate our approach. Finally, we summarize our contributions and plans for future work in Section 5.

2. Problem Description

We consider a network containing a set of vulnerable components E that can be targeted by an attacker. A defender has access to $b_1 \in \mathbb{N}$ sensors that can be positioned among a set of locations (nodes) V for network monitoring. A sensor positioned at node $v \in V$ monitors a subset of components $E_v \subseteq E$, which we refer to as the *monitoring set* of v . For ease of exposition, we denote $n := |V|$ and $[k] := \{1, \dots, k\}$ for every $k \in \mathbb{N}$.

We consider that sensors can potentially differ in their detection capabilities. Specifically, for each sensor $k \in [b_1]$, we let $\lambda_k \in (0, 1]$ denote its accuracy, i.e., the probability that it detects an attack conducted against a given component within the monitoring set of the node at which it is positioned. We order the sensors so that $\lambda_1 \geq \dots \geq \lambda_{b_1}$. Without loss of generality, we assume that multiple sensors cannot be simultaneously positioned at the same node. Indeed, positioning additional sensors at a node $v \in V$ can be equivalently viewed as positioning them among $b_1 - 1$ different copies of node v , where each copy has an identical monitoring set E_v . A *sensor positioning* is then represented as a vector $s = (s_1, \dots, s_{b_1}) \in (V \cup \{0\})^{b_1}$ such that $s_i \neq s_j$ for every $(i, j) \in [b_1]^2$ with $i \neq j$ and $s_i, s_j \neq 0$. Here, $s_k \in V$ represents the node at which sensor $k \in [b_1]$ is positioned by the defender, and $s_k = 0$ corresponds to sensor k not being positioned within the

network. For consistency, we let $E_0 := \emptyset$. We denote the set of all sensor positionings as A_1 .

To analyze the problem of strategically positioning sensors in the network, we introduce a zero-sum game $\Gamma := \langle \{1, 2\}, (\Delta(A_1), \Delta(A_2)), (-U, U) \rangle$. In this game, Player 1 (**P1**) is the defender who selects a sensor positioning $s \in A_1$. Simultaneously, Player 2 (**P2**) is an attacker who selects a subset of components $T \in 2^E$ to target, where $|T| \leq b_2$ and $b_2 \in [|E|]$ is the number of attack resources he has at his disposal. We refer to such a subset of components as an *attack plan*, and denote the set of all attack plans as A_2 .

In such security settings, it may be beneficial for one or both players to randomize their strategies. This feature is especially important for applications where sensing resources can be regularly moved throughout a network, which increases the strategic uncertainty faced by the attacker and hence generally achieves a higher protection level [10, 22]. Thus, we allow **P1** and **P2** to select mixed strategies. A *mixed strategy* for the defender (resp. attacker) is a probability distribution over the set of sensor positionings A_1 (resp. the set of attack plans A_2). Namely, we define the set of mixed inspection and attack strategies as $\Delta(A_1) := \{\sigma^1 \in [0, 1]^{|A_1|} \mid \sum_{s \in A_1} \sigma_s^1 = 1\}$ and $\Delta(A_2) := \{\sigma^2 \in [0, 1]^{|A_2|} \mid \sum_{T \in A_2} \sigma_T^2 = 1\}$ respectively, where σ_s^1 (resp. σ_T^2) represents the probability assigned to the sensor positioning $s \in A_1$ (resp. the attack plan $T \in A_2$) under the inspection strategy σ^1 (resp. the attack strategy σ^2). We assume that the players' strategies are independent randomizations.

In this model, we assume that the sensors are safe from possible damage during an attack; only the components in the network can be targeted. Additionally, we assume that detection is independent across attacks and sensors, and that if an attack against a component is detected, then the defender can nullify the damage. Hence, in our model we consider an attack on a component by **P2** to be successful if and only if it is not detected by **P1**. As such, **P1** (resp. **P2**) seeks to minimize (resp. maximize) the expected number of undetected attacks which, for any strategy profile $(\sigma^1, \sigma^2) \in \Delta(A_1) \times \Delta(A_2)$, is given by

$$U(\sigma^1, \sigma^2) := \mathbb{E}_{(\sigma^1, \sigma^2)} \left[\sum_{T \in A_2} \prod_{k=1}^{b_1} \left(1 - \lambda_k \mathbb{1}_{\{e \in E_{s_k}\}} \right) \right],$$

where the expectation is taken over all pairs of actions $(s, T) \in A_1 \times A_2$, which are selected with probability $\sigma_s^1 \cdot \sigma_T^2$ by the players' strategies.

Next, we show an instantiation of the zero-sum game Γ via an example.

Example 1. We consider an example of a network represented in Figure 1.

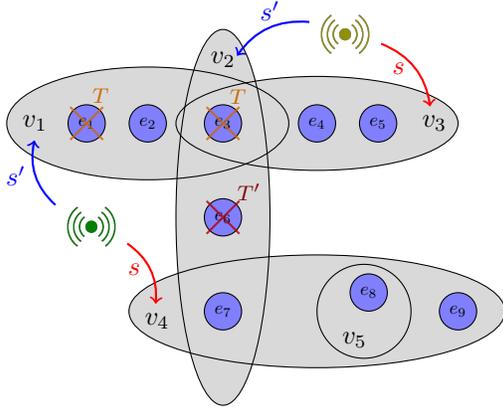


Figure 1. Game instance on a network containing 5 nodes and 9 components.

In this example, the set of nodes is $V = \{v_1, \dots, v_5\}$, the set of components is $E = \{e_1, \dots, e_9\}$, and the monitoring sets are $E_{v_1} = \{e_1, e_2, e_3\}$, $E_{v_2} = \{e_3, e_6, e_7\}$, $E_{v_3} = \{e_3, e_4, e_5\}$, $E_{v_4} = \{e_7, e_8, e_9\}$, and $E_{v_5} = \{e_8\}$. The defender has two sensors. Sensor 1 (in green) has accuracy $\lambda_1 = 0.9$, and sensor 2 (in yellow) has accuracy $\lambda_2 = 0.5$. In this example, the defender selects the randomized inspection strategy σ^1 defined by $\sigma_s^1 = 0.4$ and $\sigma_{s'}^1 = 0.6$, with $s = (v_4, v_3)$ and $s' = (v_1, v_2)$. Simultaneously, the attacker selects the randomized attack strategy σ^2 defined by $\sigma_T^2 = 0.2$ and $\sigma_{T'}^2 = 0.8$, with $T = \{e_1, e_3\}$ and $T' = \{e_4\}$. For this example, the expected number of undetected attacks is given by

$$\begin{aligned} U(\sigma^1, \sigma^2) &= \sigma_s^1 \sigma_T^2 (1 + (1 - \lambda_2)) \\ &\quad + \sigma_s^1 \sigma_{T'}^2 (1) \\ &\quad + \sigma_{s'}^1 \sigma_T^2 ((1 - \lambda_1) + (1 - \lambda_1)(1 - \lambda_2)) \\ &\quad + \sigma_{s'}^1 \sigma_{T'}^2 (1 - \lambda_2) \\ &= 0.698. \end{aligned}$$

△

In simultaneous games, a solution concept is given by Nash Equilibrium. Specifically, a strategy profile $(\sigma^{1*}, \sigma^{2*}) \in \Delta(A_1) \times \Delta(A_2)$ is a *Nash Equilibrium* (NE) of Γ if for all $(\sigma^1, \sigma^2) \in \Delta(A_1) \times \Delta(A_2)$, we have

$$U(\sigma^{1*}, \sigma^2) \leq U(\sigma^{1*}, \sigma^{2*}) \leq U(\sigma^1, \sigma^{2*}).$$

Equivalently, at a NE, σ^{1*} (resp. σ^{2*}) is a best response to σ^{2*} (resp. σ^{1*}). We refer to σ^{1*} (resp. σ^{2*})

as an equilibrium inspection strategy (resp. equilibrium attack strategy). Additionally, we refer to $U(\sigma^{1*}, \sigma^{2*})$ as the *value of the game*. Since Γ is a finite zero-sum game, the value $U(\sigma^{1*}, \sigma^{2*})$ exists and is identical for every strategy profile $(\sigma^{1*}, \sigma^{2*}) \in \Delta(A_1) \times \Delta(A_2)$ that is a NE. In other words, the value of the game is unique and well-defined.

Furthermore, the zero-sum game Γ can be solved using the following linear programming problem [21]:

$$(\mathcal{P}) \quad \min_{\sigma^1 \in \Delta(A_1)} \max_{T \in A_2} U(\sigma^1, T).$$

Specifically, the equilibrium inspection strategies, equilibrium attack strategies, and value of the game Γ are given by the optimal primal solutions, optimal dual solutions, and optimal value of (\mathcal{P}) , respectively.

However, solving (\mathcal{P}) becomes intractable even for medium-sized networks due to the combinatorial nature of the players' sets of actions: the number of variables and constraints in (\mathcal{P}) are given by $1 + |A_1| = 1 + \sum_{i=0}^{b_1} i! \binom{n}{i}$ and $1 + |A_2| = 1 + \sum_{j=0}^{b_2} \binom{|E|}{j}$, respectively. Thus, in this paper, we present an approach to provide approximate solutions to the game Γ . We first derive an analytical characterization of a class of NE when the monitoring sets are mutually disjoint. We then leverage this result in Section 4 to derive a heuristic method for computing an approximate solution in general.

Henceforth, we assume without loss of generality that (i) $b_1 \leq n$, (ii) $b_2 \leq |E|$, (iii) each monitoring set E_v ($v \in V$) is nonempty, and (iv) every component $e \in E$ belongs to at least one monitoring set. Indeed, if some components do not belong to any monitoring set, then **P2** will always target these components and allocate his remaining resources among the components that belong to at least one monitoring set.

Our game models scenarios where, for instance, each component represents an asset that can be hacked, and each node represents a computer on which software protocols can be installed to detect cyber attacks. In this scenario, our sensors are the software security protocols, which each have a certain probability of detecting a cyber attack. Stronger protocols are harder to be bypassed, and will detect an intrusion with a higher probability than a weaker protocol, which a hacker can more easily bypass.

Finally, we note that in a zero-sum game, no player has a first-mover advantage. This implies that if the players were to play sequentially, the equilibrium solutions would remain valid. Thus, the game Γ can be used to model scenarios where the attacker selects his attack strategy after observing the defender's inspection strategy. This type of situation is frequently encountered

in cybersecurity applications and various other security problems more broadly.

3. Game-Theoretic Analysis for Mutually Disjoint Monitoring Sets

In this section, we study the game Γ when all the monitoring sets are mutually disjoint. That is, when $E_v \cap E_w = \emptyset$ for all $(v, w) \in V^2$ such that $v \neq w$. Without loss of generality, we rewrite the set of nodes as $V = \{v_1, \dots, v_n\}$ so that $|E_{v_1}| \geq \dots \geq |E_{v_n}|$. Furthermore, to simplify the equilibrium analysis, we define for every $(\sigma^1, v) \in \Delta(A_1) \times V$ the *detection probability* of node v under σ^1 as:

$$p_{\sigma^1}(v) := \sum_{j=1}^{b_1} \lambda_j \sum_{\{s \in A_1 | s_j = v\}} \sigma_s^1.$$

That is, $p_{\sigma^1}(v)$ represents the probability that an attack in the monitoring set E_v is detected under the inspection strategy σ^1 .

Similarly, we define for every $(\sigma^2, e) \in \Delta(A_2) \times E$ the *attack probability* of component e under σ^2 as

$$p_{\sigma^2}(e) := \sum_{\{T \in A_2 | e \in T\}} \sigma_T^2.$$

That is, $p_{\sigma^2}(e)$ represents the probability that e is targeted under the attack strategy σ^2 .

In order to maximize the expected number of undetected attacks, **P2**'s incentive is to spread his attacks across the monitoring sets, thus making it more challenging for **P1** to detect the attacks. However, **P2** is constrained by the topology of the network, and more particularly by the sizes of the different monitoring sets. This in turn will impact **P1**'s best-response inspection strategy.

More formally, we consider the following quantity:

$$k^* = \min \left\{ k \in [n] \left| \frac{b_2 - \sum_{j=k+1}^n |E_{v_j}|}{k} \geq |E_{v_{k+1}}| \right. \right\},$$

where we let $|E_{v_{n+1}}| := 0$. Essentially, $\{E_{v_1}, \dots, E_{v_{k^*}}\}$ represents the monitoring sets that are not fully targeted by **P2** when he spreads his attacks.

The next theorem then characterizes a class of NE of the game Γ when the monitoring sets are mutually disjoint:

Theorem 1. *If $E_v \cap E_w = \emptyset$ for all $(v, w) \in V^2$ such that $v \neq w$, then a strategy profile $(\sigma^{1*}, \sigma^{2*}) \in$*

$\Delta(A_1) \times \Delta(A_2)$ is a NE if it satisfies the following conditions:

$$p_{\sigma^{1*}}(v_i) = \begin{cases} \frac{1}{k^*} \sum_{j=1}^{\min\{b_1, k^*\}} \lambda_j & \text{if } 1 \leq i \leq k^* \\ \lambda_i & \text{if } k^* < i \leq b_1 \\ 0 & \text{if } \max\{b_1, k^*\} < i \leq n, \end{cases} \quad (1)$$

$$\sum_{e \in E_{v_i}} p_{\sigma^{2*}}(e) = \begin{cases} \frac{1}{k^*} \left(b_2 - \sum_{j=k^*+1}^n |E_{v_j}| \right) & \text{if } 1 \leq i \leq k^* \\ |E_{v_i}| & \text{if } k^* < i \leq n. \end{cases} \quad (2)$$

Furthermore, the value of the game is given by the following expression:

$$b_2 - \sum_{i=1}^{b_1} \lambda_i \min \left\{ |E_{v_i}|, \frac{1}{k^*} \left(b_2 - \sum_{j=k^*+1}^n |E_{v_j}| \right) \right\}.$$

From Theorem 1, we obtain that when the monitoring sets are mutually disjoint, a class of NE can be described analytically using the players' resources and the sizes of the monitoring sets. In particular, we find that in equilibrium, **P2** targets all the components in $E_{v_{k^*+1}}, \dots, E_{v_n}$, and allocates his remaining resources uniformly among the first k^* monitoring sets $E_{v_1}, \dots, E_{v_{k^*}}$. By definition of k^* , we have that:

$$\frac{b_2 - \sum_{j=k^*+1}^n |E_{v_j}|}{k^*} \geq |E_{v_{k^*+1}}|.$$

Therefore, since **P1** aims to minimize the number of undetected attacks, her incentive, given **P2**'s equilibrium attack strategy, is to position her best sensors (i.e., those with the highest accuracy) among the nodes $\{v_1, \dots, v_{k^*}\}$. Moreover, since **P2** targets all components in $E_{v_{k^*+1}}, \dots, E_{v_n}$, **P1**'s incentive is to position her next best sensor (if available) to the remaining node with the largest monitoring set, namely v_{k^*+1} . **P1** then repeats this process until all her sensors are positioned.

Since the monitoring sets $E_{v_1}, \dots, E_{v_{k^*}}$ are not fully targeted under **P2**'s equilibrium attack strategy, **P1** must randomize the positioning of her best sensors among the nodes $\{v_1, \dots, v_{k^*}\}$ to ensure that **P2** does not have an incentive to deviate from his strategy. Thus, **P1**'s equilibrium inspection strategy is such that the detection probability of each node in $\{v_1, \dots, v_{k^*}\}$ is identical, and given by $\frac{1}{k^*} \sum_{j=1}^{\min\{b_1, k^*\}} \lambda_j$. In the next

lemma, we construct a strategy profile that satisfies the detection and attack probability conditions (1)-(2) of Theorem 1:

Lemma 1.

1. – If $b_1 \leq k^*$, consider for every $l \in [k^*]$ the following sensor positioning:

$$s^l := \begin{cases} (v_l, \dots, v_{l+b_1-1}) & \text{if } 1 \leq l \leq k^* - b_1 + 1 \\ (v_l, \dots, v_{k^*}, v_1, \dots, v_{l+b_1-k^*-1}) & \text{if } k^* - b_1 + 1 < l \leq k^*. \end{cases}$$

– If $b_1 > k^*$, consider for every $l \in [k^*]$ the following sensor positioning:

$$s^l := \begin{cases} (v_1, \dots, v_{k^*}, v_{k^*+1}, \dots, v_{b_1}) & \text{if } l = 1 \\ (v_l, \dots, v_{k^*}, v_1, \dots, v_{l-1}, v_{k^*+1}, \dots, v_{b_1}) & \text{if } 1 < l \leq k^*. \end{cases}$$

Then, $\sigma^{1*} \in \Delta(A_1)$ defined by

$$\sigma_{s^l}^{1*} = \frac{1}{k^*} \forall l \in [k^*], \text{ and } \sigma_s^{1*} = 0 \text{ otherwise,}$$

satisfies condition (1) in Theorem 1.

2. Let $b'_2 := k^* \left\lceil \frac{1}{k^*} \left(b_2 - \sum_{j=k^*+1}^n |E_{v_j}| \right) \right\rceil$, and for $l \in [k^*]$ let

$$C^l := \{1, \dots, l + b_2 - b'_2 - k^* - 1\} \cup \{l, \dots, \min\{l + b_2 - b'_2 - 1, k^*\}\}.$$

Consider attack plans T^l ($l \in [k^*]$) defined as follows:

$$|T^l \cap E_{v_j}| := \begin{cases} \frac{b'_2}{k^*} + 1 & \text{if } j \in C^l \\ \frac{b'_2}{k^*} & \text{if } j \in [k^*] \setminus C^l \\ |E_{v_j}| & \text{if } k^* < j \leq n. \end{cases}$$

Then, $\sigma^{2*} \in \Delta(A_2)$ defined by

$$\sigma_{T^l}^{2*} = \frac{1}{k^*} \forall l \in [k^*], \text{ and } \sigma_T^{2*} = 0 \text{ otherwise,}$$

satisfies condition (2) in Theorem 1.

From Lemma 1, we find that an equilibrium inspection strategy can be constructed by “cycling” the positioning of sensors $1, \dots, \min\{k^*, b_1\}$ among the nodes v_1, \dots, v_{k^*} : s^1 positions sensor 1 at node v_1 , sensor 2 at node v_2 , and so on. Then, s^2 positions sensor 1 at node v_2 , sensor 2 at node v_3 and so on. Furthermore, if $b_1 > k^*$, then **P1** deterministically positions sensors $k^* + 1, \dots, b_1$ at the remaining nodes, in decreasing order of their monitoring sets’ size: she positions sensor $k^* + 1$ at node v_{k^*+1} , sensor $k^* + 2$ at node v_{k^*+2} , and so on.

Similarly, an equilibrium attack strategy can be constructed by first deterministically targeting all the components in $E_{v_{k^*+1}}, \dots, E_{v_n}$. Then,

$\left\lfloor \frac{1}{k^*} \left(b_2 - \sum_{i=k^*+1}^n |E_{v_i}| \right) \right\rfloor$ components are deterministically targeted within each monitoring set in $E_{v_1}, \dots, E_{v_{k^*}}$. Finally, **P2** “cycles” his remaining attack resources (if any are remaining) over the remaining components in $E_{v_1}, \dots, E_{v_{k^*}}$.

Next, we illustrate Theorem 1 and Lemma 1 with an example.

Example 2. Consider the network shown in Figure 2.

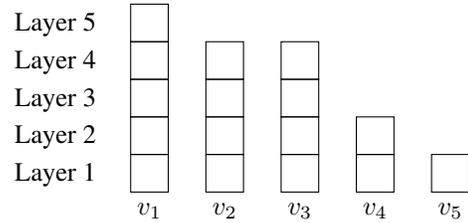


Figure 2. Example of a network with 5 nodes, 16 components, and mutually disjoint monitoring sets.

In this illustration, each square represents a component that can only be monitored from the node indicated below it. Thus, in this example, $|E_{v_1}| = 5$, $|E_{v_2}| = 4$, $|E_{v_3}| = 4$, $|E_{v_4}| = 2$, and $|E_{v_5}| = 1$. To simplify our equilibrium description, let $e_{i,j}$ ($i \in [n]$, $j \in [|E_{v_i}|]$) represent the component in layer j of monitoring set E_{v_i} . This example can be used to represent a computer network in which each computer lies within a closed section of the network, and such that each computer in a given closed section can detect cyberattacks conducted against only the components in its section.

Suppose that **P1** has 4 sensors. Furthermore, we consider that **P2** has $b_2 = 10$ attack resources. To spread his attacks in equilibrium, **P2** can first allocate 5 attack resources to target one component in each monitoring set (in layer 1). Then **P2** can allocate 4 attack resources to target one more component in each monitoring set that is not fully targeted (in layer 2). Finally, **P2**

can uniformly randomize his remaining attack resource among the remaining 3 monitoring sets that still have untargeted components.

In particular, $k^* = 3$ in this example, and an attack strategy σ^{2*} constructed from Lemma 1 is given as follows:

$$\sigma_T^{2*} = \begin{cases} \frac{1}{3} & \text{if } T = T_0 \cup \{e_{1,3}\} \\ \frac{1}{3} & \text{if } T = T_0 \cup \{e_{2,3}\} \\ \frac{1}{3} & \text{if } T = T_0 \cup \{e_{3,3}\} \\ 0 & \text{otherwise,} \end{cases}$$

where $T_0 = \{e_{1,1}, e_{2,1}, e_{3,1}, e_{4,1}, e_{5,1}, e_{1,2}, e_{2,2}, e_{3,2}, e_{4,2}\}$. We note that σ^{2*} satisfies conditions (2).

Since **P1** has $4 > k^*$ sensors, she cycles the positioning of her 3 most accurate sensors among the nodes v_1, v_2, v_3 , and deterministically positions her remaining sensor at v_4 . The construction of such an equilibrium inspection strategy σ^{1*} from Lemma 1 is given as follows:

$$\sigma_s^{1*} = \begin{cases} \frac{1}{3} & \text{if } s = (v_1, v_2, v_3, v_4) \\ \frac{1}{3} & \text{if } s = (v_2, v_3, v_1, v_4) \\ \frac{1}{3} & \text{if } s = (v_3, v_1, v_2, v_4) \\ 0 & \text{otherwise.} \end{cases}$$

The NE $(\sigma^{1*}, \sigma^{2*})$ is illustrated in Figure 3. In this example, sensor 1 (in green) has accuracy $\lambda_1 = 0.9$, sensor 2 (in yellow) has accuracy $\lambda_2 = 0.5$, sensor 3 (in orange) has accuracy $\lambda_3 = 0.4$, and sensor 4 (in maroon) has accuracy $\lambda_4 = 0.2$.

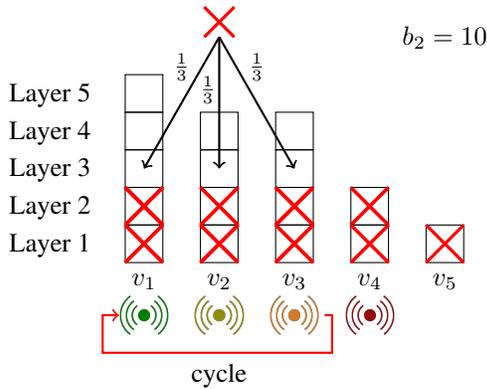


Figure 3. Example of a NE.

In this NE, we observe that the 3 most accurate sensors are randomized so that the detection probability of each node in $\{v_1, v_2, v_3\}$ has an identical detection probability given by $\frac{1}{3}(0.9 + 0.5 + 0.4) = 0.6$. We also note that node v_5 is never monitored in this NE. The

expected number of attacks in each of the monitoring sets $E_{v_1}, E_{v_2}, E_{v_3}$ is given by $\frac{7}{3}$. Every component in the remaining monitoring sets is deterministically targeted. Thus, the value of the game Γ , i.e., the expected number of undetected attacks in equilibrium, for this example is given by $10 - 3 \times 0.6 \times \frac{7}{3} - 0.2 \times 2 = 5.4$. \triangle

Theorem 1 demonstrates that there are scenarios where it is beneficial for **P1** to leave some components completely unmonitored and instead allocate her resources on parts of the network where there will be a larger number of attacks. Such scenarios occur when $k^* < n$, i.e., when the number of attack resources b_2 is large enough and the monitoring sets are of heterogeneous sizes.

Conversely, when $k^* = n$, which occurs if and only if $b_2 < n|E_{v_n}|$, **P1** randomizes her sensors over all the nodes in the network and monitors every component with identical probability. In fact, in such cases, we have the following result:

Corollary 1. *The set of equilibrium inspection strategies is identical for any number of attack resources satisfying $b_2 < n|E_{v_n}|$.*

Hence, if **P1** does not know the exact number of attack resources **P2** has at his disposal, but knows that $b_2 < n|E_{v_n}|$, then she can compute an equilibrium inspection strategy by simply assuming that $b_2 = 1$.

Next, we investigate conditions under which **P2** needs to use all of his b_2 resources in equilibrium when the monitoring sets are mutually disjoint:

Proposition 1. *If $b_1 \geq k^*$ and $\lambda_j = 1$ for every $j \in [k^*]$, then for any $b_2 > k^*|E_{v_{k^*+1}}| + \sum_{j=k^*+1}^n |E_{v_j}|$, an attack plan T^* of size $k^*|E_{v_{k^*+1}}| + \sum_{j=k^*+1}^n |E_{v_j}|$ that satisfies*

$$\forall j \in [n], |T^* \cap E_{v_j}| = \min\{|E_{v_j}|, |E_{v_{k^*+1}}|\}$$

is an equilibrium attack strategy.

Otherwise, for any $b_2 \leq |E|$, any equilibrium attack strategy σ^{2} necessarily randomizes over attack plans T of size exactly b_2 .*

This proposition shows that if **P1** has at least k^* sensors with perfect detection accuracy, then **P2** does not need to utilize more than $k^*|E_{v_{k^*+1}}| + \sum_{j=k^*+1}^n |E_{v_j}|$ attack resources in equilibrium. Indeed, any additional attack resource would be necessarily allocated to components monitored by perfect sensors, and hence will be detected with probability 1. Therefore, simply targeting $\min\{|E_{v_j}|, |E_{v_{k^*+1}}|\}$ components within each monitoring set E_{v_j} ensures a maximum expected number of undetected attacks in equilibrium.

Finally, the following proposition shows that **P1** must always use all her sensors in equilibrium:

Proposition 2. *For any $b_1 \leq n$, any equilibrium inspection strategy σ^{1*} necessarily randomizes over sensor positionings $s \in A_1$ such that $s_k \neq 0$ for all $k \in [b_1]$.*

From this proposition, we conclude that in any NE, **P1**'s inspection strategy must randomize over sensor positionings that utilize all her resources when $b_1 \leq n$.

4. General Case Approximation

4.1. Solution Approach

In this section, we leverage our equilibrium results in the case of disjoint monitoring sets to design a heuristic approach for computing an approximate equilibrium inspection strategy in general. In the general case when monitoring sets are not necessarily disjoint, the main challenge lies in determining the subset of nodes that should receive sensors in equilibrium. As observed in Section 3, **P2** aims to spread his attacks to maximize the number of undetected attacks. Therefore, **P1**'s incentive is to position her sensors on nodes that collectively monitor a large number of network components.

One natural candidate set of nodes to receive sensors is given by a *minimum set cover*, i.e., a set of nodes $S \in 2^V$ of minimum size that collectively monitors all network components. Minimum set covers can be obtained by solving the following optimization problem, which can be formulated as an integer program:

$$\min_{S \in 2^V} |S| \text{ subject to } \cup_{v \in S} E_v = E.$$

Although the minimum set cover problem is NP-hard, modern mixed-integer optimization solvers can be used to optimally solve large-scale problem instances [14].

To utilize our results in Section 3, we must recreate an instantiation where the monitoring sets are mutually disjoint. To this end, we partition the set of network components by utilizing the monitoring sets of the nodes in a minimum set cover $S = \{v'_1, \dots, v'_m\} \in 2^V$. In Theorem 1, we observed that **P2** cannot spread his attacks as much in the disjoint case when the monitoring sets are of heterogeneous sizes, thus leading to a lower expected number of undetected attacks. Hence, we partition the set of network components into m subsets by greedily assigning each component to the largest monitoring set containing that component. Specifically, we first determine the monitoring set E_v , $v \in S$ of maximum size, suppose it is $E_{v'_1}$, and then remove every component that belongs to $E_v \cap E_{v'_1}$ (for all $v \in S \setminus \{v'_1\}$)

from E_v . We then repeat this process with the second largest monitoring set, and so on until each network component belongs to exactly one set in the partitioning.

Once this partitioning is obtained, we have an instance with m disjoint monitoring sets. From this, we construct an inspection strategy $\sigma^{1'}$ according to Lemma 1 that satisfies (1) in Theorem 1. Since equilibrium inspection strategies are optimal solutions of (\mathcal{P}) (see Section 2), we evaluate the performance of our approximate inspection strategy $\sigma^{1'}$ by computing its objective value in (\mathcal{P}) , i.e., $\max_{T \in A_2} U(\sigma^{1'}, T)$. This determines the worst-case expected number of undetected attacks if **P1** selects $\sigma^{1'}$ as her inspection strategy. Since for every attack plan $T \in A_2$, $U(\sigma^{1'}, T) = \sum_{e \in T} U(\sigma^{1'}, e)$, the largest number of undetected attacks can be efficiently computed by greedily selecting the b_2 components with highest probability of undetection under $\sigma^{1'}$.

Our heuristic approach can be summarized as follows:

Algorithm 1: Heuristic Approach

Input: – Set of nodes V
– Set of components E
– Monitoring sets E_v , $v \in V$
– Number of sensors $b_1 \in \mathbb{N}$
– Number of attack resources $b_2 \in \mathbb{N}$
– Sensors' accuracies $\lambda_k \in (0, 1]$, $k \in [b_1]$

Result: – Inspection strategy $\sigma^{1'} \in \Delta(A_1)$

- 1 Compute a minimum set cover $S = \{v'_1, \dots, v'_m\}$
- 2 Set $E'_v \leftarrow E_v, \forall v \in S$
- 3 Set $V' \leftarrow S$
- 4 **while** $V' \neq \emptyset$ **do**
- 5 Select $v' \in \arg \max\{|E'_v|, v \in V'\}$
- 6 $E'_v \leftarrow E'_v \setminus (E'_{v'} \cap E'_v), \forall v \in V' \setminus \{v'\}$
- 7 $V' \leftarrow V' \setminus \{v'\}$
- 8 **end**
- 9 Order the nodes in S so that $|E'_{v'_1}| \geq \dots \geq |E'_{v'_m}|$
- 10 $k^* \leftarrow \min \left\{ k \in [m] \mid \frac{b_2 - \sum_{j=k+1}^m |E'_{v'_j}|}{k} \geq |E'_{v'_k}| \right\}$
- 11 Construct $\sigma^{1'}$ according to Lemma 1.

Next, we implement our heuristic approach on an example network and evaluate the performance of the resulting inspection strategies.

4.2. Computational Study

We consider the benchmark cyber-physical distribution network given in Figure 4.



Figure 4. Benchmark Kentucky distribution network.

This real-world network from Kentucky is composed of 420 nodes that can receive sensors, and 492 components that are vulnerable to cyber-physical attacks, which induce disruptions. To detect these attacks, we consider that the defender has access to flow and pressure sensors that can be deployed at access points and shifted from one to another. These sensors can measure signals which can be used to detect the sudden rate of change of pressure or mass flow at different locations of the network. In our study, we compute the monitoring set of each node through simulations using a threshold-based detection model, as proposed in [23, 24]. All network simulations were implemented in Matlab, and all optimization problems were solved using Gurobi on a computer with a 2.3 GHz 8-Core Intel Core i9 processor and 32 GB of RAM.

To evaluate the performance of our heuristic approach we consider 10 game instances where **P2** has $b_2 = 1$ attack resource and **P1** has $b_1 \in [10]$ sensors, with sensor $k \in [b_1]$ having accuracy $\lambda_k = 1 - 0.05(k - 1)$. For such instances, (\mathcal{P}) only has 494 constraints since $b_2 = 1$. Therefore, equilibrium inspection strategies of Γ can be obtained by solving (\mathcal{P}) using the column generation algorithm.

We now implement our heuristic approach: We solve the minimum set cover problem, and obtain a set of 19 nodes. Next, we greedily partition the set of network components into 19 sets. Finally, we construct an inspection strategy $\sigma^{1'}$ according to Lemma 1. The worst-case expected number of undetected attacks under the inspection strategy $\sigma^{1'}$ is then computed by selecting the b_2 components with the highest probability of not being detected under $\sigma^{1'}$. In Figure 5, we illustrate for

$b_1 \in [10]$ the optimality gap achieved by $\sigma^{1'}$, i.e., the relative difference between the worst-case performance of $\sigma^{1'}$ and the value of the game (given by the optimal value of (\mathcal{P})).

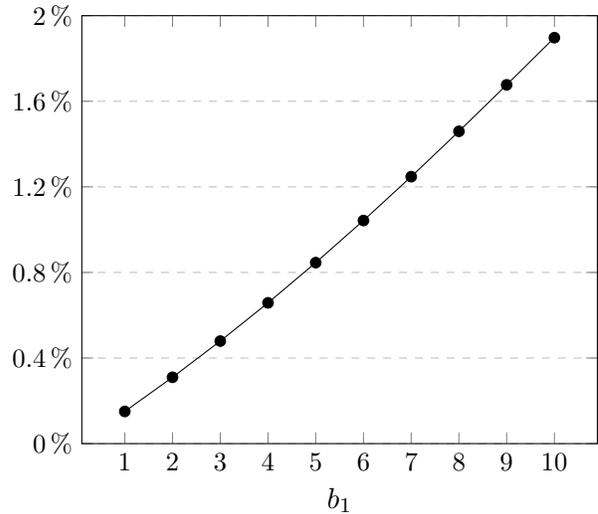


Figure 5. Optimality gap of the heuristic solution when $b_2 = 1$.

From Figure 5, we observe that our heuristic solution achieves a detection performance that is close to the detection performance in equilibrium. However, we note that as the number of sensors increases, the optimality gap associated with our heuristic solution increases. This is due to the fact that when **P1** has more sensors, she can strategically coordinate their positioning so as to maximize the detection probabilities of the components that are monitored from multiple locations. In contrast, our heuristic approach assigns such components to a single monitoring set to construct an inspection strategy using a disjoint instance.

Next, we compare in Figure 6 the running times of our heuristic method with the running times of the column generation algorithm for computing equilibrium inspection strategies.

Interestingly, we observe that our heuristic solution is obtained in 0.11 seconds, and this running time is almost identical for any number of sensors. The reason is that most of the running time is spent computing a minimum set cover. As previously mentioned, although this problem is NP-hard, it can be efficiently solved by modern mixed-integer optimization solvers. In contrast, the time required to compute an equilibrium inspection strategy using column generation increases exponentially with the number of sensors b_1 . This is due to the fact that the number of variables in (\mathcal{P}) grows combinatorially with respect to b_1 . For instance,

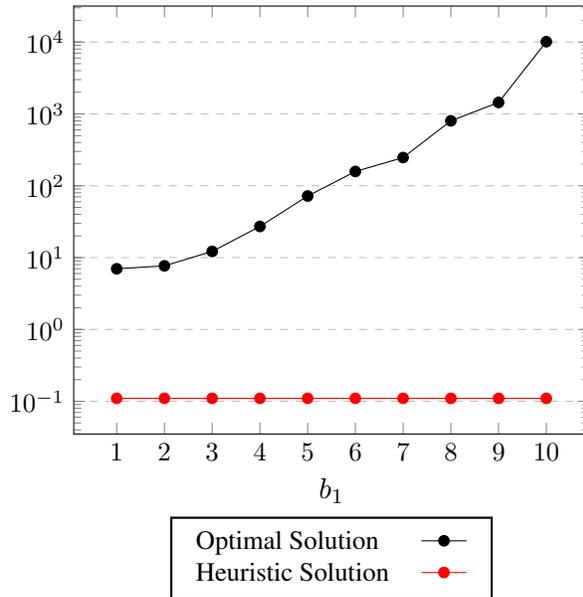


Figure 6. Running times (in seconds) of the column generation algorithm and the heuristic method.

when $b_1 = 10$, the number of variables in (\mathcal{P}) is approximately $1.54 \cdot 10^{26}$ for this network.

Finally, we note that the column generation algorithm for computing equilibrium inspection strategies cannot be used in practice when $b_2 > 1$, as the number of constraints in (\mathcal{P}) grows combinatorially with respect to b_2 . By leveraging the analytical characterization derived in Section 3, our heuristic approach remains scalable for any value of b_1 and b_2 , and can be implemented for large-scale networks, as minimum set covers have been shown to be efficiently solvable for networks containing more than 100,000 nodes and components [14].

5. Conclusion

In this paper, we studied a network inspection game in which a defender allocates sensors with potentially heterogeneous detection capabilities in order to detect multiple attacks caused by a strategic attacker. In this two-person zero-sum game, the defender (resp. attacker) seeks to minimize (resp. maximize) the expected number of undetected attacks by selecting a potentially randomized inspection (resp. attack) strategy. When the monitoring sets are mutually disjoint, we derived an analytical characterization of a class of NE for this game. Additionally, we studied the dependence of these NE on the network topology, sensor accuracies, and the number of resources the attacker has at his disposal. We then leveraged our equilibrium analysis to

design a heuristic solution approach for the general case based on minimum set covers. Our computational study on a benchmark cyber-physical distribution network showed that our heuristic approach is computationally tractable and provides inspection strategies with good detection performance. In future work, we aim to refine our heuristic solution approach and provide theoretical performance guarantees.

References

- [1] H. Sandberg, S. Amin, and K. Johansson, "Cyberphysical security in networked control systems: An introduction to the issue," *IEEE Control Systems Magazine*, vol. 35, no. 1, pp. 20–23, 2015.
- [2] S. Weerakkody and B. Sinopoli, "Challenges and opportunities: Cyber-physical security in the smart grid," in *Smart Grid Control: Overview and Research Opportunities* (J. Stoustrup, A. Annaswamy, A. Chakraborty, and Z. Qu, eds.), pp. 257–273, Springer International Publishing, 2019.
- [3] M. Russon, "US fuel pipeline hackers 'didn't mean to create problems'," *BBC News*, May 10 2021. <https://www.bbc.com/news/business-57050690>.
- [4] J. Slay and M. Miller, "Lessons learned from the Maroochy water breach," in *Critical Infrastructure Protection* (E. Goetz and S. Sheno, eds.), pp. 73–82, Springer US, 2008.
- [5] R. M. Lee, M. J. Assante, and T. Conway, "Analysis of the cyber attack on the Ukrainian power grid," tech. rep., Electricity Information Sharing and Analysis Center, 2016.
- [6] Q. Zhu and T. Basar, "Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: Games-in-games principle for optimal cross-layer resilient control systems," *IEEE Control Systems Magazine*, vol. 35, no. 1, pp. 46–65, 2015.
- [7] F. Miao, Q. Zhu, M. Pajic, and G. J. Pappas, "A hybrid stochastic game for secure control of cyber-physical systems," *Automatica*, vol. 93, pp. 55–63, 2018.
- [8] A. Gupta, C. Langbort, and T. Başar, "Dynamic games with asymmetric information and resource constrained players with applications to security of cyberphysical systems," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 71–81, 2017.
- [9] A. R. Hota, A. A. Clements, S. Sundaram, and S. Bagchi, "Optimal and game-theoretic deployment of security investments in interdependent assets," in *Decision and Game Theory for Security* (Q. Zhu, T. Alpcan, E. Panaousis, M. Tambe, and W. Casey, eds.), pp. 101–113, Springer International Publishing, 2016.
- [10] J. Pita, M. Jain, J. Marecki, F. Ordóñez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus, "Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport," in *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Industrial Track, AAMAS '08*, pp. 125–132, International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [11] A. Washburn and K. Wood, "Two-person zero-sum games for network interdiction," *Operations Research*, vol. 43, pp. 243–251, 04 1995.

- [12] D. Bertsimas, E. Nasrabadi, and J. B. Orlin, "On the power of randomization in network interdiction," *Operations Research Letters*, vol. 44, no. 1, pp. 114–120, 2016.
- [13] J. Milošević, M. Dahan, S. Amin, and H. Sandberg, "A network monitoring game with heterogeneous component criticality levels," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 4379–4384, 2019.
- [14] M. Dahan, L. Sela, and S. Amin, "Network inspection for detecting strategic attacks," *Operations Research (To appear)*. Available: <http://arxiv.org/abs/1705.00349>, 2021.
- [15] M. Pirani, J. A. Taylor, and B. Sinopoli, "Strategic sensor placement on graphs," *Systems & Control Letters*, vol. 148, 2021.
- [16] A. Krause, A. Roper, and D. Golovin, "Randomized sensing in adversarial environments," in *Proc. Twenty-Second International Joint Conference on Artificial Intelligence - Volume Three, IJCAI11*, pp. 2133–2139, AAAI Press, 2011.
- [17] S. Alpern, A. Morton, and K. Papadaki, "Patrolling games," *Oper. Res.*, vol. 59, no. 5, pp. 1246–1257, 2011.
- [18] A. Garnaev, *Search Games and Other Applications of Game Theory*. Lecture Notes in Economics and Mathematical Systems, Springer, 2000.
- [19] M. Mavronicolas, V. Papadopoulou, A. Philippou, and P. Spirakis, "A network game with attackers and a defender," *Algorithmica*, vol. 51, no. 3, pp. 315–341, 2008.
- [20] A. Garnaev, G. Garnaeva, and P. Goutal, "On the infiltration game," *Int. J. Game Theory*, vol. 26, no. 2, pp. 215–221, 1997.
- [21] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*, vol. 160. SIAM, 1998.
- [22] D. S. Hochbaum and B. Fishbain, "Nuclear threat detection with mobile distributed sensor networks," *Ann. Oper. Res.*, vol. 187, no. 1, pp. 45–63, 2011.
- [23] L. Sela Perelman, W. Abbas, X. Koutsoukos, and S. Amin, "Sensor placement for fault location identification in water networks: A minimum test cover approach," *Automatica*, vol. 72, pp. 166–176, 2016.
- [24] A. Deshpande, S. E. Sarma, K. Youcef-Toumi, and S. Mekid, "Optimal coverage of an infrastructure network using sensors with distance-decaying sensing quality," *Automatica*, vol. 49, no. 11, pp. 3351–3358, 2013.