Strategic Monitoring of Networked Systems with Heterogeneous Security Levels

Jezdimir Milošević, Mathieu Dahan, Saurabh Amin, and Henrik Sandberg, *Fellow, IEEE*

***Abstract*— We consider a strategic network monitoring problem involving the operator of a networked system and an attacker. The operator aims to randomize the placement of multiple protected sensors to monitor and protect components that are vulnerable to attacks. We account for the heterogeneity in the components' security levels and formulate a large-scale maximin optimization problem. After analyzing its structure, we propose a three-step approach to approximately solve the problem. First, we solve a generalized covering set problem and run a combinatorial algorithm to compute an approximate solution. Then, we compute approximation bounds by solving a nonlinear set packing problem. To evaluate our solution approach, we implement two classical solution methods based on column generation and multiplicative weights updates, and test them on real-world water distribution and power systems. Our numerical analysis shows that our solution method outperforms the classical methods on large-scale networks, as it efficiently generates solutions that achieve a close to optimal performance and that are simple to implement in practice.**

***Index Terms*— Strategic Network Monitoring, Optimization, Game Theory, Networked Control Systems, Other Applications**

I. Introduction

LARGE-SCALE networked systems, such as water distribution or power systems, are lucrative targets for malicious attackers. Indeed, several attacks have already occurred [1], [2], demonstrating the need for rapid development of effective defense strategies for these systems. An essential component of every defense strategy is attack detection [3], which can be accomplished by allocating multiple sensors for system monitoring. Yet, for large-scale systems, the number of available sensors is likely to be insufficient to ensure complete coverage. A strategic adversary can exploit this deficiency to target the most critical components of the system that are unmonitored. While the literature on sensor placement has studied fixed [4], [5] and randomized sensing strategies [6], [7], none of these works has accounted for the heterogeneous criticality of the network components, and its impact on monitoring decisions. Hence, the question that arises is how to strategically allocate multiple sensors to monitor a large-scale system of critical networked components.

Specifically, we focus on determining a randomized scheduling of monitoring operations to maximize the protection of a networked system against a strategic adversary. In this network monitoring problem, the network operator seeks to randomize the placement of multiple protected sensors to monitor the components and protect them against an attacker who aims to sabotage a system component while remaining undetected. To better account for the attacker's behavior and preference, we consider that the components have possibly heterogeneous security levels, which represent the corresponding levels of effort required by an adversary to target them. This model is motivated by risk management, where one first identifies critical system components in a risk assessment, and then allocates protection resources accordingly [3]. By monitoring a component with a protected sensor, the operator maximizes the security level of that component. Then, to protect the system against a strategic adversary, we assume that the operator aims to maximize the lowest expected security level across all components.

Contributions: Our contributions are summarized as follows:

1. We analyze the structure of the network monitoring problem and gather new insights on how the number of sensors and the heterogeneity of the component security levels impact the operator's optimal strategy. We then derive closed-form expressions of bounds on the optimal value of the network monitoring problem that depend on subsets of sensor locations and network components (Theorem 1).

2. To optimize the above-mentioned bounds and compute an approximate monitoring strategy, we derive a three-step approach. First, we formulate and solve a mixed-integer program (MIP) to compute a subset of sensor locations that forms a generalized covering set and to obtain marginal probabilities of placing sensors at each location (Proposition 1). Second, we employ the combinatorial algorithm of [8] to construct a randomized monitoring strategy that is consistent with the

Manuscript received April 6, 2023; revised August 2, 2023; accepted October 17, 2023. The work of JM and HS was supported in part by the Swedish Civil Contingencies Agency (project CERCES2) and Digital Futures (project DEMOCRITUS). The work of MD was supported by the Georgia Tech new faculty start up grant. The work of SA was supported in part by the AFOSR grant "Building Attack Resilience into Complex Networks" and NSF grant # 2039771.

Jezdimir Milošević was with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden. He is now with Scania Autonomous Transport Solutions (e-mail: jezdimirmilosevic@gmail.com).

Mathieu Dahan is with the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: mathieu.dahan@isye.gatech.edu).

Saurabh Amin is with the Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: amins@mit.edu).

Henrik Sandberg is with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden (e-mail: hsan@kth.se).

marginal probabilities computed from the MIP. Finally, we formulate a MIP to compute a subset of components that forms a nonlinear set packing and to obtain an optimality gap on our approximate solution to the network monitoring problem (Proposition 2).

3. We show that our solution approach optimally solves the network monitoring problem when the sets of components that are monitored from each sensor location are mutually disjoint (Proposition 3). Furthermore, when security levels are identical, our solution approach can be simplified by solving the minimum set cover and maximum set packing problems (Proposition 4).

4. We evaluate our solution approach and compare it with two classical solution methods based on column generation and multiplicative weights updates using two security applications: (i) detecting contaminants injected in a water distribution system and (ii) protecting actuators in a power system against extended replay attacks. Our computational results show that our approach provides solutions to the network monitoring problem with low optimality gaps for large-scale systems. Furthermore, our approach outperforms the two classical methods in terms of scalability and practical implementability of the generated monitoring strategies.

Related work: Optimization and game theoretic models have been developed for studying various security related problems, including the development of defense strategies [9]–[12], the design of anomaly detectors [13], the allocation of security budget [14], and the placement of sensors [4]–[7]. For instance, [15]–[18] consider bilevel and trilevel optimization problems to model defender-attacker interactions where each player selects a pure strategy. In contrast, we focus on randomized strategies, which are recognized to be more effective when the number of sensors to deploy is limited [9], [14].

Other models that have been studied include inspection problems [19]–[21] and search games [22], [23], in which a searcher is concerned with the optimal way of looking for a hidden adversary in a search space; see for example [24]–[27]. In particular, the recent literature on hide-and-seek games investigates the problem of coordinating multiple inspection resources for detecting hidden objects in boxes [8], [28], [29]. However, such problems do not account for the impact of sensing range and the network topology on the monitoring strategies. Still, we leverage the algorithm of [8] as part of our solution approach for computing randomized monitoring strategies with small supports.

Finally, the existing literature on sensor placement considers developing both fixed [4], [5] and randomized monitoring strategies [6], [7]. The placement problems considered previously aimed to minimize estimation error [30], achieve optimal coverage [31], detect faults [32], or improve the system's security [33]. The study [33] is related to ours and proposes two static sensor placement strategies to improve the security of actuators in large-scale systems. Our focus is instead on randomized strategies, which can be translated into random scheduling of inspections that can be performed on a day-to-day basis by utility or security personnel [34].

Our problem is also related to the simultaneous zero-sum game in [6], where the operator places multiple sensors to maximize the number of detected (homogeneous) attacks. Recent work extended this model by considering the imperfect detection by sensors [35] and the different types of sensing resources [36], for which they derive heuristic approaches. In addition, [37] formulated a game in which sensors are positioned in the nodes of a networked control system to detect attacks on them, and considered imperfect detection through a linear filter that processes the sensors' measurements to detect attacks. The authors derived equilibrium results using tools from structured systems and graph theory. However, to the best of our knowledge, no article has investigated the coordination of *multiple* sensing resources to *strategically* monitor a *networked* system with components of *heterogeneous* criticalities.

Preliminary results (without proofs) have been presented in [38]. Specifically, [38] investigate a simpler solution approach based on minimal set covers, maximal set packings, and column generation. They also provide approximation bounds under restrictive assumptions. By further leveraging the problem structure, we generalize their results and obtain a more efficient solution approach with stronger approximation bounds. We also provide a more comprehensive evaluation of our solution approach through an extended computational study that involves two other solution algorithms and two security applications.

Organization: Section II introduces and formulates the strategic network monitoring problem. In Section III, we derive and analyze our solution approach based on generalized covering sets and nonlinear set packings. We then present two alternative solution methods based on column generation and multiplicative weights updates in Section IV. Section V evaluates and compares these solution approaches through a computational study. Finally, Section VI summarizes our findings and provides avenues for future work.

II. Problem Formulation

We consider a networked system consisting of a set of components $\mathcal{U}$. Each component $u \in \mathcal{U}$ faces a risk of attack by an adversary who wishes to disrupt or take control of the network. For every $u \in \mathcal{U}$, let $\varphi_u \in [0,1]$ denote its *security level*, i.e., the effort that the attacker needs to spend in order to successfully target u and compromise its operation. Security levels can for example be estimated based on previously deployed security measures, configuration of the system, or monetary cost for conducting an attack against the components. A component u for which $\varphi_u = 0$ (resp. $\varphi_u = 1$) is unsafe (resp. highly secure).

The network operator can increase the security levels of some components by positioning r sensors within a set of locations $\mathcal{X}$. Specifically, a sensor positioned at location $x \in \mathcal{X}$ can monitor and secure a subset of components $U_x \subseteq \mathcal{U}$. For every $x \in \mathcal{X}$, we refer to U_x as the monitoring set of x. Typically, monitoring sets are computed from the characteristics of the sensors and the dynamics of the networked system. If multiple sensors are positioned at locations $X \subseteq \mathcal{X}$, then the monitored and secured components are given by $U_X := \cup_{x \in X} U_x$.

We let $\mathcal{A} := \{X \subseteq \mathcal{X} \mid |X| \leq r\}$ denote the set of actions of the operator. Then, for every sensor positioning $X \in \mathcal{A}$ and

every component $u \in \mathcal{U}$, the resulting *post-security level* of u is given by $f(X,u) := \varphi_u \mathbb{1}_{\{u \notin U_X\}} + \mathbb{1}_{\{u \in U_X\}}$. In particular, the post-security levels of all monitored components are maximal and equal to one. Next, we illustrate the monitoring model with an example.

Example 1. *Consider the networked system and monitoring model illustrated in Fig. 1.*

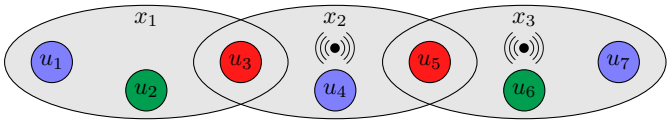

Fig. 1. Networked system with 3 sensor locations and 7 components.

In this example, the set of sensor locations is $\mathcal{X} = \{x_1, x_2, x_3\}$*, the set of components is* $\mathcal{U} = \{u_1, \dots, u_7\}$*, and the monitoring sets are* $U_{x_1} = \{u_1, u_2, u_3\}$*,* $U_{x_2} = \{u_3, u_4, u_5\}$*, and* $U_{x_3} = \{u_5, u_6, u_7\}$*. The security levels are given by* $\varphi_{u_3} = \varphi_{u_5} = 0.2$ *(red in the figure),* $\varphi_{u_1} = \varphi_{u_4} = \varphi_{u_7} = 0.5$ *(blue in the figure), and* $\varphi_{u_2} = \varphi_{u_6} = 0.8$ *(green in the figure). We assume that the operator places two sensors on* $X = \{x_2, x_3\}$*. Then, each component* u *in* $U_X = \{u_3, u_4, u_5, u_6, u_7\}$ *is monitored and has a post-security level equal to* $f(X,u) = 1$*. The remaining components* u_1 *and* u_2 *are not monitored and their post-security levels are equal to their original security levels:* $f(X,u_1) = 0.5$ *and* $f(X,u_2) = 0.8$*. Note that* u_1 *has the lowest resulting post-security level and is the most vulnerable to attacks.* △

In such combinatorial security settings, the network operator can significantly benefit from randomizing the positioning of their sensors [34], [39], [40]. When sensors are carried by inspection crews [41], randomized sensor positionings can be implemented via daily or weekly schedules of inspections. We then consider the problem where the operator is interested in positioning their sensors in a randomized manner to protect the network against an attacker who would target the most vulnerable component. In other words, the operator aims to maximize the lowest expected post-security level across all network components. Let $\Delta_{\mathcal{A}} := \{\sigma \in [0,1]^{|\mathcal{A}|} \mid \sum_{X \in \mathcal{A}} \sigma_X = 1\}$ denote the set of mixed strategies of the operator, i.e., the set of probability distributions over $\mathcal{A}$; here, σ_X denotes the probability that the operator positions sensors on locations $X \in \mathcal{A}$. Then, the problem of strategic network monitoring can be formulated as

$$\max_{\sigma \in \Delta_{\mathcal{A}}} \min_{u \in \mathcal{U}} \mathbb{E}_{X \sim \sigma}[f(X,u)]. \tag{$\mathcal{M}$}$$

Problem ($\mathcal{M}$) can be formulated as a linear program with $|\mathcal{A}| + 1$ variables and $|\mathcal{U}| + 1$ constraints. However, the cardinality of $\mathcal{A}$ grows combinatorially with respect to the number of available sensors. Thus, standard methods for solving linear programs cannot be used for large-scale systems, and different approaches are needed to compute an optimal or approximate monitoring strategy. In Section III, we develop an approach based on mixed-integer programs (MIPs) to compute an approximate solution of ($\mathcal{M}$), while in Section IV we discuss alternative approaches based on column generation and multiplicative weights updates.

III. Main Results

In this section, we analyze the structure of the network monitoring problem ($\mathcal{M}$) and derive an approximate solution approach. Our approach first solves a generalized covering set problem and runs a combinatorial algorithm to construct a monitoring strategy with small support. Then, it computes an optimality gap by solving a nonlinear set packing problem. We further investigate problem instances with special structures, for which we refine our approach and optimality gaps.

A. Preliminaries

To carry out our analysis, we first introduce additional quantities and notations: For any optimization problem ($\mathcal{O}$), we denote by $z^*_{(\mathcal{O})}$ its optimal value. Then, for every location $x \in \mathcal{X}$, let $\overline{\varphi}_x := \min_{u \in U_x} \varphi_u$ denote the lowest security level of a component in the monitoring set of x. This quantity will impact the operator's monitoring strategy, as the objective is to maximize the worst-case component post-security level. We quantify the *criticality* of a location $x \in \mathcal{X}$ (resp. component $u \in \mathcal{U}$) from the value of $\overline{\varphi}_x$ (resp. φ_u): If $\overline{\varphi}_x$ (resp. φ_u) is low, we say that $x \in \mathcal{X}$ (resp. $u \in \mathcal{U}$) is more critical. We additionally define, for every sensor positioning $X \in \mathcal{A}$ and every location $x \in \mathcal{X}$, $\bar{f}(X,x) := \overline{\varphi}_x \mathbb{1}_{\{x \notin X\}} + \mathbb{1}_{\{x \in X\}}$ to be the lowest post-security level within the monitoring set U_x.

For any subset of locations $C \subseteq \mathcal{X}$ (resp. any subset of components $T \subseteq \mathcal{U}$), we denote as $\overline{\varphi}_C$ (resp. φ_T) the following vector $(\overline{\varphi}_x)_{x \in C}$ (resp. $(\varphi_u)_{u \in T}$), sorted in nondecreasing order. For every subset of components $T \subseteq \mathcal{U}$, we also define the following quantity $S_T := \sum_{u \in T} (1 - \varphi_u)^{-1}$.

For any vector $y \in \mathbb{R}^n$, we denote its support as $\operatorname{supp}(y) := \{i \in [\![1,n]\!] \mid y_i \neq 0\}$, where $[\![1,n]\!] := \{1, \dots, n\}$. Then, given a monitoring strategy $\sigma \in \Delta_{\mathcal{A}}$, we define its *node basis* as $\mathcal{X}(\sigma) := \cup_{X \in \operatorname{supp}(\sigma)} X$, that is, the set of nodes that are monitored with positive probability under the monitoring strategy σ. Note that node bases provide a good indicator of the practical implementability of a monitoring strategy: A strategy that monitors a small number of locations can easily be implemented in practice via a randomized scheduling of operations [34].

Without loss of generality, we assume the following: (i) $U_x \neq \emptyset$ for every $x \in \mathcal{X}$; (ii) every component can be monitored from at least one sensor location; and (iii) $\varphi_u < 1$ for every $u \in \mathcal{U}$. If any of these assumptions does not hold, then the sets $\mathcal{X}$ and $\mathcal{U}$ can be adjusted to satisfy these assumptions without impacting the optimal solutions of ($\mathcal{M}$).

We then consider the collection of set covers $\mathcal{S}$ of our networked system. Specifically, a subset of locations $C \subseteq \mathcal{X}$ is a *set cover* if $U_C = \mathcal{U}$, i.e., every component is monitored from at least one location in C. Let n^* denote the minimum set cover size. We note that if $r \geq n^*$, then an optimal solution of the network monitoring problem ($\mathcal{M}$) consists of placing sensors on a set cover and ensuring a post-security level of 1 for every component. Henceforth, we assume that $r < n^*$.

Finally, we consider the collection of set packings $\mathcal{T}$ of our networked system. Specifically, a subset of components $T \subseteq \mathcal{U}$ is a *set packing* if $|T \cap U_x| \leq 1$ holds for every $x \in \mathcal{X}$. Set packings indicate the *spread* of the network, which

directly impacts the monitoring strategy of the operator: Since components in a set packing must be monitored from distinct sensor locations, a network with large set packings is more challenging to protect. We denote the maximum set packing size as m^*.

B. Structural Analysis

We first analyze the structure of the network monitoring problem $(\mathcal{M})$ and derive bounds on its optimal value. To obtain a lower bound on $z^*_{(\mathcal{M})}$, we simplify the problem by supposing that when the operator places a sensor at location $x \in \mathcal{X}$, they protect the component u within the monitoring set U_x that is the most critical. The premise is that the post-security level achieved for any component in U_x will be at least that of component u. Formally, we relax the objective function of $(\mathcal{M})$ and consider the following problem:

$$\max_{\sigma\in\Delta_{\mathcal{A}}} \min\left\{ \min_{x\in\mathcal{X}(\sigma)} \mathbb{E}_{X\sim\sigma}[\bar{f}(X,x)], \min_{u\notin U_{\mathcal{X}(\sigma)}} \varphi_u \right\}. \qquad (\overline{\mathcal{M}})$$

Problem $(\overline{\mathcal{M}})$ selects a monitoring strategy σ that partitions the set of components $\mathcal{U}$ into two subsets: the subset of components $\mathcal{U} \setminus U_{\mathcal{X}(\sigma)}$ that are never monitored by σ and for which the post-security levels are given by the corresponding original security levels; and the subset of components $U_{\mathcal{X}(\sigma)}$ that are monitored with positive probability. In this latter set, σ guarantees a post-security level for every component in $U_{\mathcal{X}(\sigma)}$ by improving the lowest security level within each monitoring set U_x for $x \in \mathcal{X}(\sigma)$.

In fact, we provide in the next lemma a different perspective on problem $(\overline{\mathcal{M}})$ to guide our analysis:

Lemma 1. *Problem $(\overline{\mathcal{M}})$ satisfies the following property:*

$$z^*_{(\overline{\mathcal{M}})} = \max_{C\subseteq\mathcal{X}} \min\left\{ z^*_{(\mathcal{R}_C)}, \min_{u\notin U_C} \varphi_u \right\},$$

where for every $C \subseteq \mathcal{X}$, $z^*_{(\mathcal{R}_C)}$ is the optimal value of

$$\max_{\sigma\in\Delta_{\mathcal{A}}} \min_{x\in C} \mathbb{E}_{X\sim\sigma}[\bar{f}(X,x)]. \qquad (\mathcal{R}_C)$$

Thus, $(\overline{\mathcal{M}})$ can be solved by first determining the set of locations C over which to randomize the placement of sensors, and then determining the monitoring strategy that optimizes $(\mathcal{R}_C)$. Given a preselected subset of n locations $C = \{x_1, \ldots, x_n\} \subseteq \mathcal{X}$, indexed such that $\overline{\varphi}_{x_1} \leq \cdots \leq \overline{\varphi}_{x_n}$, we provide some intuition on the monitoring strategy σ that maximizes $\min_{x\in C} \mathbb{E}_{X\sim\sigma}[\bar{f}(X,x)]$. Initially, the goal is to randomize the placement of the sensors among the locations in C so as to equalize the post-security level of the most critical component in each of their monitoring sets. However, because of the heterogeneity of the security levels, the operator should instead focus on the most critical locations and not monitor the most secure locations. The number of most critical locations in C to monitor is given by

$$k^*_C := \max\left\{ k \in [\![1,n]\!] \;\middle|\; \overline{\varphi}_{x_k} \leq 1 - \frac{k-r}{\sum_{l=1}^{k}(1-\overline{\varphi}_{x_l})^{-1}} \right\}.$$

For simplicity, we also denote $S^{k^*}_C := \sum_{l=1}^{k^*_C}(1-\overline{\varphi}_{x_l})^{-1}$.

Using this property and the collection of set packings $\mathcal{T}$, we can then derive the following main theorem:

Theorem 1. *The optimal value of the network monitoring problem $(\mathcal{M})$ is bounded as follows:*

$$z^*_{(\overline{\mathcal{M}})} = \max_{C\subseteq\mathcal{X}} \min\left\{ 1 - \frac{k^*_C - r}{S^{k^*}_C}, \min_{u\notin U_C} \varphi_u \right\}$$
$$\leq z^*_{(\mathcal{M})} \leq \min\left\{ 1, \min_{T\in\mathcal{T}} \left(1 - \frac{|T|-r}{S_T} \right) \right\}.$$

Furthermore, for any set of $n > r$ locations $C = \{x_1, \ldots, x_n\} \subseteq \mathcal{X}$ indexed in nondecreasing order of their criticalities, any monitoring strategy σ' satisfying the following conditions is an optimal solution of $(\mathcal{R}_C)$:

$$\mathbb{P}_{X\sim\sigma'}(x \in X) = \begin{cases} 1 - \dfrac{k^*_C - r}{(1-\overline{\varphi}_x)S^{k^*}_C}, & \text{if } x \in \{x_1, \ldots, x_{k^*_C}\} \\ 0, & \text{otherwise.} \end{cases} \qquad (1)$$

Interestingly, from this theorem, we find that solving problem $(\overline{\mathcal{M}})$ indeed provides a lower bound on the optimal value of $(\mathcal{M})$. Furthermore, given a subset of sensor locations $C \subseteq \mathcal{X}$ of size greater than r, the post-security level that can be guaranteed for any component in U_C is given by $1 - \frac{k^*_C - r}{S^{k^*}_C}$. To achieve this post-security level, the network operator must ensure that the probability that each location in C is monitored follows (1). Importantly, we find that the k^*_C most critical locations must be monitored, with a probability that is decreasing with the associated lowest security level. In contrast, the remaining locations should not be monitored, as their security levels are already larger than the post-security level achieved for the k^*_C most critical locations (by definition of k^*_C). We note that a monitoring strategy satisfying (1) is guaranteed to exist by Farkas' lemma [6], [39].

From Theorem 1, we also find that upper bounds on the optimal value of the network monitoring problem $(\mathcal{M})$ can be computed from set packings. We will leverage this finding to efficiently compute optimality gaps for our solutions.

C. Solution Approach

From our structural analysis and Theorem 1, we next derive a three-step approach to approximately solve the network monitoring problem $(\mathcal{M})$. Steps 1 and 2 compute a monitoring strategy that optimally solves the relaxed problem $(\overline{\mathcal{M}})$, and Step 3 computes a set packing that achieves the upper bound from Theorem 1.

In the first step, we formulate a MIP to compute $z^*_{(\overline{\mathcal{M}})}$ as well as the marginal probabilities of monitoring each location at optimality of $(\overline{\mathcal{M}})$. In particular, we leverage the fact that the objective function in $(\mathcal{R}_C)$ can be expressed using these marginal probabilities. Indeed, for every $\sigma \in \Delta_{\mathcal{A}}$ and $x \in \mathcal{X}$,

$$\mathbb{E}_{X\sim\sigma}[\bar{f}(X,x)] = \overline{\varphi}_x + (1-\overline{\varphi}_x)\mathbb{P}_{X\sim\sigma}(x \in X). \qquad (2)$$

Using Lemma 1, we then formulate the following MIP,

which we refer to as a *generalized covering set problem*:

$$\max_{y,z,\rho} \; z \tag{$\mathcal{C}$}$$
$$\text{s.t. } z \le \overline{\varphi}_x + (1-\overline{\varphi}_x)\rho_x + M_1(1-y_x), \; \forall x \in \mathcal{X} \tag{3}$$
$$z \le \varphi_u + M_1 \sum_{\{x\in\mathcal{X} \,|\, u\in U_x\}} y_x, \quad \forall u \in \mathcal{U} \tag{4}$$
$$0 \le \rho_x \le y_x, \quad \forall x \in \mathcal{X} \tag{5}$$
$$\sum_{x\in\mathcal{X}} \rho_x = r, \tag{6}$$
$$y_x \in \{0,1\}, \quad \forall x \in \mathcal{X}.$$

Here, for every $x \in \mathcal{X}$, the binary variable y_x determines if x can receive a sensor, while the continuous variable ρ_x determines the marginal probability that x receives a sensor. M_1 is a "big M", which we can set to $1 - \min_{u\in\mathcal{U}} \varphi_u$. Given $C = \operatorname{supp}(y)$, we then observe that Constraints (3) (combined with (2)) and Constraints (4) ensure that the auxiliary variable z is upper bounded by $\min_{x\in C} \mathbb{E}_{X\sim\sigma}[\bar{f}(X,x)]$ and $\min_{u\notin U_C} \varphi_u$, respectively. In fact, we show the equivalence between $(\overline{\mathcal{M}})$ and $(\mathcal{C})$:

Proposition 1. *Problems $(\overline{\mathcal{M}})$ and $(\mathcal{C})$ have identical optimal values: $z^*_{(\overline{\mathcal{M}})} = z^*_{(\mathcal{C})}$. Furthermore, a monitoring strategy $\sigma^* \in \Delta_{\mathcal{A}}$ is an optimal solution of $(\overline{\mathcal{M}})$ if and only if there exists an optimal solution (y^*, z^*, ρ^*) of $(\mathcal{C})$ satisfying $\rho^*_x = \mathbb{P}_{X\sim\sigma^*}(x \in X)$ for every $x \in \mathcal{X}$.*

Thus, the optimal value of $(\overline{\mathcal{M}})$ can be computed by solving $(\mathcal{C})$. In addition, given an optimal solution (y^*, z^*, ρ^*) of $(\mathcal{C})$, ρ^* provides the marginal probabilities that each node must be monitored at optimality of $(\overline{\mathcal{M}})$.

The second step of our approach consists of reconstructing an optimal solution of $(\overline{\mathcal{M}})$ from ρ^*. To this end, we utilize the combinatorial algorithm derived in [8] that computes a probability distribution $\sigma^* \in \Delta_{\mathcal{A}}$ satisfying $\mathbb{P}_{X\sim\sigma^*}(x \in X) = \rho^*_x$ for every $x \in \mathcal{X}$. From Proposition 1, this guarantees the optimality of σ^* for $(\overline{\mathcal{M}})$.

Specifically, the algorithm iteratively expresses a vector $\rho^* \in [0,1]^{|\mathcal{X}|}$ satisfying (6) as a convex combination of the incidence vector of a sensor positioning $X \in \mathcal{A}$, and another vector $\rho' \in [0,1]^{|\mathcal{X}|}$ satisfying (6) and with at least one additional integral component compared to ρ^*. We refer the reader to [8] for the detailed algorithm, which we henceforth refer to as the *coordination algorithm*.

Importantly, given $n = |\operatorname{supp}(\rho^*)|$ the number of locations that must receive a sensor with positive probability, this algorithm runs in time $O(n^2)$ and constructs a monitoring strategy of support size $n+1$. We recall that a monitoring strategy with small support and/or small node basis can more easily be implemented in practice. We will numerically evaluate the node basis sizes of our solutions in our computational study.

The final step of our approach consists of evaluating the optimality gap of our solution for the network monitoring problem. To this end, we derive a MIP to efficiently compute the upper bound on $z^*_{(\mathcal{M})}$ from Theorem 1. Specifically, we aim to compute a set packing $T \in \mathcal{T}$ that minimizes $1 - \frac{|T|-r}{S_T}$, which is equivalent to finding a set packing of size greater than r that minimizes $\frac{S_T}{|T|-r}$. Although this is a *nonlinear set packing problem*, we can reformulate it as the following MIP:

$$\min_{t,y,z} \sum_{l=r+1}^{|\mathcal{U}|} \frac{1}{l-r} t_l \tag{$\mathcal{P}$}$$
$$\text{s.t. } \sum_{u\in U_x} y_u \le 1, \quad \forall x \in \mathcal{X} \tag{7}$$
$$\sum_{u\in\mathcal{U}} y_u = \sum_{l=r+1}^{|\mathcal{U}|} l z_l, \tag{8}$$
$$\sum_{l=r+1}^{|\mathcal{U}|} z_l = 1, \tag{9}$$
$$t_l \ge \sum_{u\in\mathcal{U}} (1-\varphi_u)^{-1} y_u - M_2(1-z_l), \forall l \in [\![r+1, |\mathcal{U}|]\!] \tag{10}$$
$$t_l \ge 0, \quad \forall l \in [\![r+1, |\mathcal{U}|]\!] \tag{11}$$
$$z_l \in \{0,1\}, \quad \forall l \in [\![r+1, |\mathcal{U}|]\!]$$
$$y_u \in \{0,1\}, \quad \forall u \in \mathcal{U}.$$

Specifically, the binary variables y_u, for $u \in \mathcal{U}$, select the components to be part of the set packing. To parametrize the nonlinearity of the problem, we introduced for every $l \in [\![r+1, |\mathcal{U}|]\!]$ the binary variable z_l, which is equal to 1 if the set packing is of size l. Finally, for every $l \in [\![r+1, |\mathcal{U}|]\!]$, the continuous variable t_l represents the product $z_l \sum_{u\in\mathcal{U}} (1-\varphi_u)^{-1} y_u$. Constraints (7) ensure that $\operatorname{supp}(y)$ is a set packing. Constraints (8)-(9) select the set packing size. Then, constraints (10)-(11) ensure that at optimality, the objective value is equal to $\sum_{u\in\operatorname{supp}(y)} (1-\varphi_u)^{-1} / (|\operatorname{supp}(y)| - r)$. M_2 is also a "big M", which we can set to $|\mathcal{U}|(1 - \max_{u\in\mathcal{U}} \varphi_u)^{-1}$.

As a consequence, the upper bound on the optimal value of $(\mathcal{M})$ from Theorem 1 can be computed by solving $(\mathcal{P})$:

Proposition 2. *The following equality holds:*

$$\min\left\{1, \min_{T\in\mathcal{T}} \left(1 - \frac{|T|-r}{S_T}\right)\right\} = 1 - \frac{1}{z^*_{(\mathcal{P})}}.$$

We note that the number of variables in $(\mathcal{P})$ can be reduced in practice by computing an upper bound on the maximum set packing size in the networked system. Given the maximum set packing size m^* (or an upper bound), we can reduce the variables in $(\mathcal{P})$ from $|\mathcal{U}| + 2(|\mathcal{U}| - r)$ to $|\mathcal{U}| + 2(m^* - r)$, and set M_2 to $m^*(1 - \max_{u\in\mathcal{U}} \varphi_u)^{-1}$, which can provide significant computational benefits.

Thus, we obtain a three-step approach to compute an approximate solution to the network monitoring problem $(\mathcal{M})$ and evaluate its optimality gap. Next, we further study our solution for problem instances with particular characteristics.

D. Special Cases

1) Non-overlapping monitoring sets: We consider problem instances for which monitoring sets are pairwise disjoint, i.e., $U_x \cap U_{x'} = \emptyset$ for every $x \neq x' \in \mathcal{X}$. Such instances occur when it is desirable to reduce sensor interference or the network's energetic cost [42], [43]. In security games, disjoint monitoring is naturally satisfied [44]–[46]. For such instances,

we show that our approach in Section III-C optimally solves the network monitoring problem:

Proposition 3. *If the monitoring sets are pairwise disjoint, then* $z^*_{(\overline{\mathcal{M}})} = z^*_{(\mathcal{M})}$. *Furthermore, any monitoring strategy satisfying* (1) *for* $C = \mathcal{X}$ *is an optimal solution of* $(\mathcal{M})$.

For instances when monitoring sets do not overlap, we find that problem $(\mathcal{M})$ can be optimally solved by first selecting the $k^*_{\mathcal{X}}$ most critical sensor locations and monitoring them with probabilities given by (1); and then by utilizing the coordination algorithm to construct a monitoring strategy that randomizes the placement of r sensors and is consistent with the marginal probabilities in (1).

2) Homogeneous security levels: We now consider instances where the security levels are identical across the components in $\mathcal{U}$. In such cases, we refine the bounds on the optimal value of $(\mathcal{M})$ using the minimum set cover and maximum set packing sizes (n^* and m^*):

Proposition 4. *If all security levels are equal to* φ, *then:*

$$z^*_{(\overline{\mathcal{M}})} = \varphi + (1-\varphi)\frac{r}{n^*} \leq z^*_{(\mathcal{M})} \leq \min\left\{1, \varphi + (1-\varphi)\frac{r}{m^*}\right\}.$$

In such instances, we find that $(\mathcal{C})$ and $(\mathcal{P})$ can be simplified and the bounds on $z^*_{(\mathcal{M})}$ can be computed by solving minimum set cover and maximum set packing problems, as in [6]. If C^* denotes a minimum set cover, then a monitoring strategy that optimizes $(\overline{\mathcal{M}})$ monitors each location in C^* with identical probability (given by (1)), and can be constructed by cycling the r sensors over C^* (see [6] for reference).

IV. Alternative Approaches

In order to compare our solution approach derived in Section III-C, we next consider two alternative approaches for solving the network monitoring problem. In particular, we show how the column generation algorithm and multiplicative weights update algorithm can be applied to solve $(\mathcal{M})$.

A. Column Generation Algorithm

Problem $(\mathcal{M})$ can be formulated as a linear program with a large number $(\binom{|\mathcal{X}|}{r} + 1)$ of variables and a small number $(|\mathcal{U}| + 1)$ of constraints. Such problems can potentially be solved using the column generation algorithm, as long as the pricing problem can be solved efficiently.

Specifically, at each iteration of the column generation algorithm, the following restricted master problem is solved:

$$\max_{\sigma, z} \; z \qquad (\mathcal{M}_{\mathcal{I}})$$

$$\text{s.t. } z - \sum_{X \in \mathcal{I}} \sigma_X f(X,u) \leq 0, \; \forall\, u \in \mathcal{U} \qquad (12)$$

$$\sum_{X \in \mathcal{I}} \sigma_X = 1, \qquad (13)$$

$$\sigma_X \geq 0, \qquad \forall\, X \in \mathcal{I},$$

where $\mathcal{I} \subseteq \mathcal{A}$ is a subset of feasible sensor positionings. To determine if the optimal solution of $(\mathcal{M}_{\mathcal{I}})$ is optimal for $(\mathcal{M})$, the algorithm then determines the variable in $(\mathcal{M})$ with highest reduced cost. Let (α^*, β^*) denote the optimal dual variables of $(\mathcal{M}_{\mathcal{I}})$, associated with Constraints (12)-(13). Then, the variable with highest reduced cost can be determined by solving the following pricing problem:

$$\max_{X \in \mathcal{A}} -\beta^* + \sum_{u \in \mathcal{U}} \alpha^*_u f(X,u).$$

In fact, the pricing problem for $(\mathcal{M})$ is a *maximum weighted covering problem*, and can be solved using the following integer program (IP):

$$\max_{y,z} \; -\beta^* + \sum_{u \in \mathcal{U}} \alpha^*_u(\varphi_u + (1-\varphi_u)y_u) \qquad (\mathcal{W})$$

$$\text{s.t. } y_u \leq \sum_{\{x \in \mathcal{X} \,|\, u \in U_x\}} z_x, \qquad \forall\, u \in \mathcal{U} \qquad (14)$$

$$\sum_{x \in \mathcal{X}} z_x = r, \qquad (15)$$

$$y_u \in \{0,1\}, \qquad \forall\, u \in \mathcal{U} \qquad (16)$$

$$z_x \in \{0,1\}, \qquad \forall\, x \in \mathcal{X}. \qquad (17)$$

In this IP, the binary variables z_x, for $x \in \mathcal{X}$, represent the nodes that receive a sensor, and the binary variables y_u, for $u \in \mathcal{U}$, represent the components that are monitored by at least one sensor.

If $z^*_{(\mathcal{W})} = 0$, then the optimal solution of $(\mathcal{M}_{\mathcal{I}})$ is optimal for $(\mathcal{M})$ and the algorithm terminates. If instead $z^*_{(\mathcal{W})} > 0$, then given an optimal solution (y^*, z^*) of $(\mathcal{W})$, the algorithm adds $X^* = \text{supp}(z^*)$ to the set $\mathcal{I}$ of variable indices, and solves the new restricted master problem $(\mathcal{M}_{\mathcal{I}})$.

B. Multiplicative Weights Update Algorithm

Problem $(\mathcal{M})$ is also equivalent to a simultaneous two-person zero-sum game, where the first player is the network operator who places r sensors across the networked system and the second player is an attacker who targets a component. The payoff function is given by f, which the operator (resp. attacker) wants to maximize (resp. minimize).

In [47], a multiplicative weights update algorithm is devised for computing approximate mixed Nash equilibria of such simultaneous games, as long as one player has a small number of actions and the other player can efficiently compute best responses. The algorithm runs for N iterations, and at each iteration $t \in [\![1, N]\!]$, generates a probability distribution π^t over $\mathcal{U}$ for the attacker as well as a feasible sensor positioning $X^t \in \mathcal{A}$ for the network operator. The algorithm starts by setting π^1 to the uniform distribution over $\mathcal{U}$. Then, at each iteration $t \in [\![1, N]\!]$, the algorithm determines the operator's best response to π^t by solving the following problem:

$$X^t \in \arg\max_{X \in \mathcal{A}} \mathbb{E}_{u \sim \pi^t}[f(X,u)].$$

In our setting, this problem can also be formulated as a maximum weighted covering problem and can be solved using the following IP:

$$\max_{y,z} \; \sum_{u \in \mathcal{U}} \pi^t_u(\varphi_u + (1-\varphi_u)y_u) \qquad (\mathcal{W}')$$

$$\text{s.t. } (14) - (17).$$

Then, the algorithm updates the attacker's strategy using the following multiplicative weights update:

$$\forall\, u \in \mathcal{U},\ \pi_u^{t+1} = \pi_u^t \frac{\eta^{f(X^t,u)}}{\sum_{u' \in \mathcal{U}} \pi_{u'}^t \eta^{f(X^t,u')}},$$

where $\eta = \left(1 + \sqrt{2\frac{\ln|\mathcal{U}|}{N}}\right)^{-1}$. The authors in [47] showed that the monitoring strategy $\bar{\sigma} \in \Delta_{\mathcal{A}}$ that uniformly randomizes over the N best responses $X^1, \ldots, X^N$ satisfies:

$$\min_{u \in \mathcal{U}} \mathbb{E}_{X \sim \bar{\sigma}}[f(X,u)] \geq z^*_{(\mathcal{M})} - \sqrt{\frac{2\ln|\mathcal{U}|}{N}} - \frac{\ln|\mathcal{U}|}{N}.$$

Thus, if $N = 4\lceil \frac{\ln|\mathcal{U}|}{\epsilon^2} \rceil$ for $\epsilon > 0$, then the (absolute) optimality gap associated with the monitoring strategy $\bar{\sigma}$ generated by this algorithm is upper bounded by ϵ.

In fact, [7] showed that each iteration of this algorithm can be conducted in polynomial time by approximately solving $(\mathcal{W}')$ using a greedy algorithm. However, the resulting guarantee on the monitoring strategy $\bar{\sigma}$ becomes

$$\min_{u \in \mathcal{U}} \mathbb{E}_{X \sim \bar{\sigma}}[f(X,u)] \geq (1 - 1/e) z^*_{(\mathcal{M})} - \epsilon,$$

for $N = 4\lceil \frac{\ln|\mathcal{U}|}{\epsilon^2} \rceil$. Since commercial solvers are now highly efficient at solving classical IPs, we instead optimally solve $(\mathcal{W}')$ at each iteration of the algorithm.

We now move to computational experiments, where we evaluate the three approaches presented in this article and compare their running times, as well as the performance and implementability of the monitoring strategies they generate.

V. Computational Study

In this section, we describe how our strategic network monitoring problem can be applied for contamination detection in water distribution systems and actuator protection in networked control systems. Our experiments are coded in Julia and all optimization problems are solved using the Gurobi solver v9.0 on a computer with 2.3-GHz, 8-Core Intel Core i9 processor and 32 GB of RAM.

A. Application I: Contamination Detection in Water Distribution Systems

We first consider the security problem where the operator of a water distribution system aims to allocate sensors to monitor their network against adversarially-induced contamination events. We model the water distribution system as a directed graph, with the vertices representing pumps, junctions, and water tanks, while the edges represent pipes. The direction of every edge is adopted to be in the direction of the water flow. We assume that the operator has access to contaminant detection sensors that can be deployed on valve access points or fire hydrants and measure water quality indicators such as electrical conductivity, free and total chlorine, turbidity, and oxygen reduction potential in the water [48]–[50].

In this application, the set of vulnerable components $\mathcal{U}$ (resp. set of sensor locations $\mathcal{X}$) is given by the set of edges (resp. nodes) of the graph modeling the water network. Monitoring sets are constructed through simulations using a hydraulic network solver [51] that tracks the advection and reaction dynamics from a contaminant intrusion event [52]. In practice, the security levels of components can be assessed based on previously deployed security measures and their accessibility to an adversary. One can then use different security scales to determine quantitative values for the security levels based on that assessment [53]. In this application, we adopt the scale $\{0.2, 0.4, 0.6, 0.8\}$ for security levels and assign them to each component uniformly at random.

Next, we implement the three solution approaches described in Sections III and IV on the water distribution system ky5 [54]. This anonymized real-world water network from Kentucky comprises 496 pipes and 420 nodes (junctions, water tanks, pumps), satisfies a demand of 2.28 million gallons of water per day, and spans 52.3 miles. In Fig. 2, we illustrate the bounds derived in Theorem 1 and computed by solving the generalized covering set problem $(\mathcal{C})$ and the nonlinear set packing problem $(\mathcal{P})$, respectively.

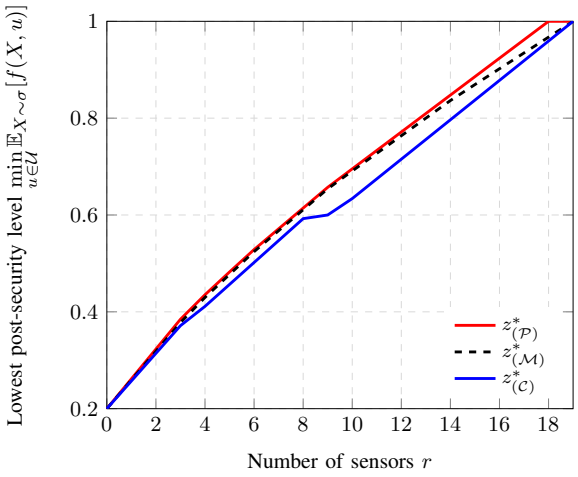

Fig. 2. Illustration of the bounds from Theorem 1 for varying numbers of sensors.

From Fig. 2, we observe that the bounds obtained from $(\mathcal{C})$ and $(\mathcal{P})$ are close to the optimal value of $(\mathcal{M})$. On average, the optimality gap given by the bounds does not exceed 4.7%, while the actual optimality gap is equal to 3.6% on average.

In general, we observe that $z^*_{(\mathcal{M})}$ increases nonlinearly with respect to the number of sensors. When r is small, the optimal monitoring strategy focuses on protecting the most vulnerable components (with a security level of 0.2). As the operator has access to more sensors, they randomize their placement to increase the post-security level of less critical components. With 19 sensors, the operator can monitor all components and ensure a post-security level of 1 for the entire network.

Next, we compare the running times of the various approaches. Specifically, we illustrate the time to (a) solve $(\mathcal{P})$ and get an upper bound on $z^*_{(\mathcal{M})}$, (b) solve $(\mathcal{C})$ and construct a monitoring strategy using the coordination algorithm (CA), (c) solve $(\mathcal{M})$ using column generation (CG), and (d) solve $(\mathcal{M})$ using multiplicative weights updates (MWU) with $N = 2{,}480$ iterations to ensure a theoretical (absolute) optimality gap of $\epsilon = 0.1$. Their running times are illustrated in Fig. 3.

Fig. 3 shows that the upper bound on $z^*_{(\mathcal{M})}$ derived in

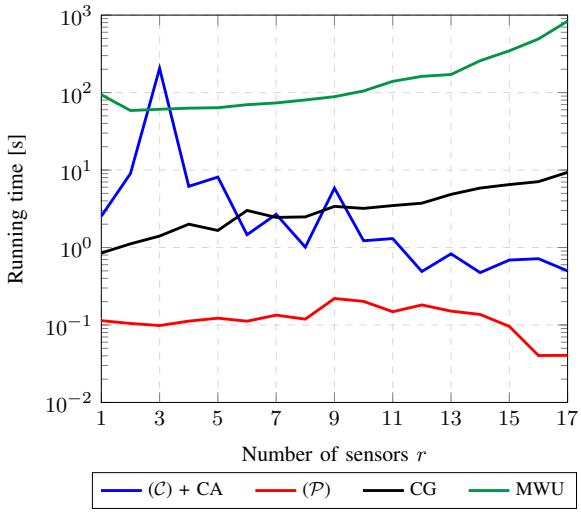

Fig. 3. Approaches' running times for varying numbers of sensors.

Theorem 1 can be efficiently computed; by solving ($\mathcal{P}$), we obtain the upper bound in less than 0.2 seconds. We also observe that the number of sensors has limited impact on the running time.

Furthermore, we find that CG is approximately 100 times faster than MWU. This is in part due to the number of iterations that are required for MWU to theoretically guarantee a small optimality gap. Interestingly, we find that CG can efficiently compute an optimal solution of ($\mathcal{M}$) when the number of sensors is small. However, as r increases, the running time of CG increases exponentially. Finally, we find that our approximate monitoring strategy, which achieves the lower bound in Theorem 1 and is obtained by solving ($\mathcal{C}$) + CA, can be efficiently computed.

We next compare the implementability of the solutions generated by the three approaches by illustrating their node basis sizes in Fig. 4.

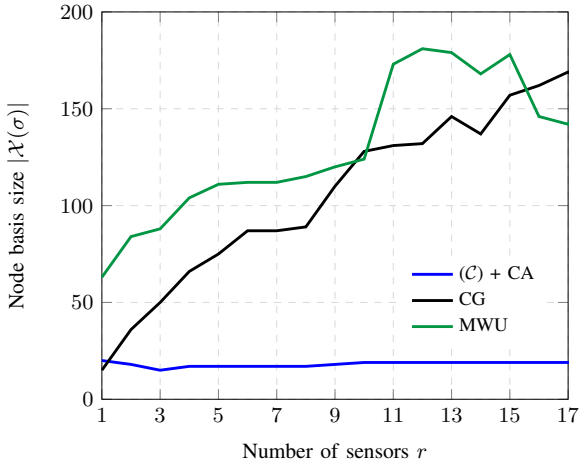

Fig. 4. Node basis sizes for varying numbers of sensors.

Fig. 4 clearly shows the implementability advantage of our monitoring strategy. We find that our solution monitors at most 20 different locations and achieves an excellent worst-case post-security level (from Fig. 2). In contrast, the optimal post-security level achieved using CG or MWU requires monitoring up to 180 different locations. From an implementation perspective, our monitoring strategy is considerably easier to translate into a daily or weekly scheduling of monitoring operations.

B. Application II: Protecting Actuators from Extended Replay Attacks

We now show how our findings can be used to protect actuators in a networked control system from extended replay attacks. The networked system consists of a set of physical states $\mathcal{X}$ and actuators $\mathcal{U}$, as well as fixed unprotected sensors. We consider an attacker who aims to conduct an undetectable attack on an actuator using an extended replay strategy. More precisely, an extended replay attack on an actuator is undetectable if the sensor measurements remain in the steady state during the attack. To this end, the attacker must also compromise a subset of components (i.e., unprotected sensors and other actuators) to ensure that the attack is undetected. Such subset of components typically depends on the dynamics of the networked system, and can be efficiently identified in large-scale systems [55, Chapter 7], [33].

Importantly, the vulnerability of an actuator u with respect to extended replay attacks can be determined based on its security index δ_u, which represents the minimum number of sensors and actuators, including u itself, that the attacker must compromise to conduct an undetectable extended replay attack against u, in the absence of protected sensors [55, Chapter 7]. If no undetectable extended replay attack against u is possible, then $\delta_u = +\infty$. We can then scale these security indices and define security levels for this application, following our modeling framework. Specifically, we adopt an identical scale as in the previous application, and define the security level of each actuator $u \in \mathcal{U}$ as follows:

$$\begin{aligned}\varphi_u =& 0.2 \cdot \mathbb{1}_{\{\delta_u \in (0,5]\}} + 0.4 \cdot \mathbb{1}_{\{\delta_u \in (5,15]\}} + 0.6 \cdot \mathbb{1}_{\{\delta_u \in (15,20]\}} \\ &+ 0.8 \cdot \mathbb{1}_{\{\delta_u \in (20,+\infty)\}} + 1.0 \cdot \mathbb{1}_{\{\delta_u = +\infty\}}. \end{aligned} \tag{18}$$

To defend the networked system, we assume the operator possesses protected sensors that cannot be compromised by the attacker and that can be positioned within the network. We assume that each positioned sensor can only measure one physical state, which is a commonly adopted model of large–scale systems [30], [56], [57]. In particular, a protected sensor can detect if the state it measures deviates from its steady state value during an extended replay attack. Then, actuator $u \in \mathcal{U}$ belongs to the monitoring set of state $x \in \mathcal{X}$ if placing a protected sensor at x prevents the attacker from conducting an undetectable extended replay attack on u.

We now consider the IEEE 2,383 bus power system. We model the system using linearized swing equations where the generators are represented by two states (rotor angle and frequency) and load buses with one state (voltage angle) [58]. We assume that all $|\mathcal{X}| = 3{,}037$ states are measurable, and that the attacker can conduct an attack using $|\mathcal{U}| = 1{,}042$ loads (we randomly selected 30% of the loads to be attackable) [59]. Fig. 5 illustrates the distribution of security levels, computed from security indices according to (18). Table I compares

TABLE I
COMPARISON OF SOLUTION APPROACHES ON IEEE 2,383 BUS POWER SYSTEM.

# sensors	Optimality gap [%]			Running time [s]			Node basis size		
r	$(\mathcal{C})$ + CA	CG	MWU	$(\mathcal{C})$ + CA	CG	MWU	$(\mathcal{C})$ + CA	CG	MWU
50	0.0	0.0	0.3	1.6	71.3	227.7	223	373	430
100	0.0	0.4	0.4	1.9	600	245.1	318	868	737
150	0.0	1.1	0.4	13.2	600	302.6	328	981	824
200	0.0	8.0	0.3	5.1	600	412.2	328	1,009	827
250	0.0	15.0	0.3	174.2	600	581.2	344	1,024	1,115
300	0.0	24.0	0.3	39.2	600	591.0	344	1,086	612

the performance of the different solution approaches on this networked system. Note that we set a time limit of 10 minutes for each approach.

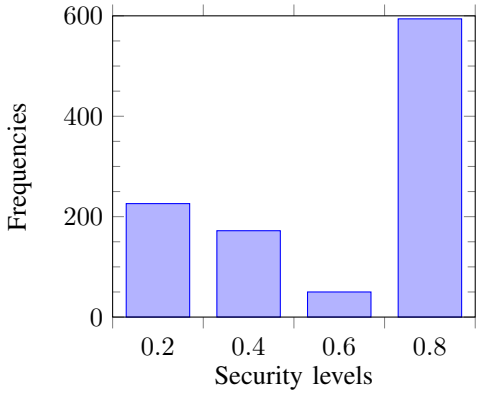

Fig. 5. Security levels distribution for the IEEE 2,383 bus power system.

We observe that for this network, our solution approach optimally solves the network monitoring problem $(\mathcal{M})$. Indeed, the upper bound obtained by solving $(\mathcal{P})$ guarantees that the monitoring strategy computed from $(\mathcal{C})$ + CA is optimal for $(\mathcal{M})$ for every number of sensors. Although we observe some variance in the running time of our approach, the average running time for this network is 29 seconds. On the other hand, MWU achieves a very high performance, but requires 393 seconds on average. This illustrates the fact that the empirical performance of MWU is better than the theoretical guarantees [47]. In contrast, we find that CG scales poorly with the number of sensors, which is interesting given that MWU and CG solve the very similar IPs $(\mathcal{W})$ and $(\mathcal{W}')$.

At optimality, the node basis size (i.e., the number of locations that are monitored with positive probability) exhibits plateaus as the number of sensors r increases. This observation is a consequence of the distribution of security levels that incentivizes the operator to secure the most critical components as r is small, and to gradually secure components with higher security levels as r increases. Finally, we note that the node basis of our monitoring strategy is up to 3 times smaller than when utilizing CG or MWU, thus providing a solution that is more easily implementable in practice.

C. Summary and Extension of Solution Methodology

In summary, we find that by extracting the structure of the network monitoring problem, we derived an efficient solution approach when facing heterogeneous security levels. Our approach solves two MIPs and runs the coordination algorithm to directly obtain a monitoring strategy with strong performance guarantees and that is easier to implement in practice. This contrasts with the classical CG and MWU algorithms that iteratively refine their solutions and generally run more slowly.

Interestingly, more flexibility can be embedded in our solution approach. For instance, if the operator is interested in an even simpler monitoring strategy, we can add the cardinality constraint $\sum_{x\in\mathcal{X}} y_x \leq s$ in $(\mathcal{C})$ to find the best monitoring strategy of $(\mathcal{M})$ that monitors at most s locations. This permits the operator to achieve the desired tradeoff between solution implementability and guaranteed post-security level.

Finally, we note that the three solution approaches can be combined for an even better performance. Indeed, for the instances for which our approach does not provide an optimal solution to the network monitoring problem $(\mathcal{M})$, we can utilize our solution as a warm start for CG and/or MWU, which will then iteratively improve it.

VI. CONCLUSION

In this article, we studied the problem of strategically monitoring large-scale networked systems with heterogeneous component criticalities using multiple sensors. Specifically, we formulated a large-scale maximin problem where the network operator selects a randomized placement of their protected sensors to improve the post-security level of the network's most critical components against an attacker.

We developed a three-step solution approach that leverages the structure of the network monitoring problem. First, we solved a generalized covering set problem to determine the marginal probabilities to monitor each location. Then, we utilized a combinatorial algorithm to randomize the positioning of the multiple sensors and obtain a monitoring strategy that is consistent with the marginal probabilities. Finally, we solved a nonlinear set packing problem to evaluate our solution's optimality gap. We also investigated problem instances with special structures, for which we refined our solution approach and optimality gaps.

We then adapted two classical solution methods for solving our monitoring problem based on column generation and multiplicative weights updates, and compared them on real-world benchmark water distribution and power systems. Our computational study revealed that our solution approach outperforms the classical methods as the size of the problem increases. In

particular, we find that our approach can efficiently compute near optimal solutions for large instances and are significantly simpler to implement.

As part of future work, it would be interesting to investigate how the solution method developed in this article (combining smaller IPs and combinatorial algorithms) can be applied to solve problems where an attacker targets multiple components simultaneously, or where sensors are faulty.

APPENDIX

Proof of Lemma 1. We consider an optimal solution σ^* of $(\overline{\mathcal{M}})$ and we set $C = \mathcal{X}(\sigma^*)$. Then, $z^*_{(\mathcal{R}_C)} \geq \min_{x\in\mathcal{X}(\sigma^*)} \mathbb{E}_{X\sim\sigma^*}[\bar{f}(X,x)]$. Furthermore, $\min_{u\notin U_C} \varphi_u = \min_{u\notin U_{\mathcal{X}(\sigma^*)}} \varphi_u$. Therefore, $z^*_{(\overline{\mathcal{M}})} \leq \max_{C\subseteq\mathcal{X}} \min\{z^*_{(\mathcal{R}_C)}, \min_{u\notin U_C} \varphi_u\}$.

To show the reverse inequality, let $C^* \in \arg\max_{C\subseteq\mathcal{X}} \min\{z^*_{(\mathcal{R}_C)}, \min_{u\notin U_C} \varphi_u\}$. We also consider an optimal solution σ^* of $(\mathcal{R}_{C^*})$. Without loss of generality, we can assume that $\mathcal{X}(\sigma^*) \subseteq C^*$ (i.e., σ^* only places sensors within C^*). Then, $\min_{x\in\mathcal{X}(\sigma^*)} \mathbb{E}_{X\sim\sigma^*}[\bar{f}(X,x)] \geq z^*_{(\mathcal{R}_{C^*})}$. Furthermore, let $u' \notin U_{\mathcal{X}(\sigma^*)}$. If $u' \in U_{C^*} \setminus U_{\mathcal{X}(\sigma^*)}$, let $x' \in C^* \setminus \mathcal{X}(\sigma^*)$ such that $u' \in U_{x'}$. We then obtain:

$$\varphi_{u'} \geq \overline{\varphi}_{x'} = \mathbb{E}_{X\sim\sigma^*}[\bar{f}(X,x')] \geq z^*_{(\mathcal{R}_{C^*})}.$$

If instead $u' \notin U_{C^*}$, then $\varphi_{u'} \geq \min_{u\notin U_{C^*}} \varphi_u$. In conclusion, $z^*_{(\overline{\mathcal{M}})} = \max_{C\subseteq\mathcal{X}} \min\{z^*_{(\mathcal{R}_C)}, \min_{u\notin U_C} \varphi_u\}$. □

Proof of Theorem 1. We first show that $z^*_{(\mathcal{M})} \geq z^*_{(\overline{\mathcal{M}})}$: Let $\sigma^* \in \Delta_{\mathcal{A}}$ be an optimal solution of $(\overline{\mathcal{M}})$ and let $u' \in \mathcal{U}$. If $u' \in U_{\mathcal{X}(\sigma^*)}$, then let $x' \in \mathcal{X}(\sigma^*)$ such that $u' \in U_{x'}$. Since $\varphi_{u'} \geq \overline{\varphi}_{x'}$ and for every $X \in \mathcal{A}$, $\mathbb{1}_{\{x'\notin X\}} \geq \mathbb{1}_{\{u'\notin U_X\}}$, then

$$\begin{aligned}\mathbb{E}_{X\sim\sigma^*}[f(X,u')] &\geq \mathbb{E}_{X\sim\sigma^*}[1-(1-\overline{\varphi}_{x'})\mathbb{1}_{\{x'\notin X\}}] \\ &\geq \min_{x\in\mathcal{X}(\sigma^*)} \mathbb{E}_{X\sim\sigma^*}[1-(1-\overline{\varphi}_x)\mathbb{1}_{\{x\notin X\}}].\end{aligned}$$

If instead $u' \notin U_{\mathcal{X}(\sigma^*)}$, then:

$$\mathbb{E}_{X\sim\sigma^*}[f(X,u')] = \varphi_{u'} \geq \min_{u\notin U_{\mathcal{X}(\sigma^*)}} \varphi_u.$$

Thus, $z^*_{(\mathcal{M})} \geq z^*_{(\overline{\mathcal{M}})}$. We next show that for every $C \subseteq \mathcal{X}$ of size greater than r, $z^*_{(\mathcal{R}_C)} = 1 - \frac{k^*_C - r}{S^{k^*}_C}$.

We rewrite $C = \{x_1, \ldots, x_n\}$ with $n > r$ and indices satisfying $\overline{\varphi}_{x_1} \leq \cdots \leq \overline{\varphi}_{x_n}$. Then necessarily, $k^*_C \geq r+1$ since

$$1 - \frac{r+1-r}{\sum_{l=1}^{r+1}(1-\overline{\varphi}_{x_l})^{-1}} \geq 1-(1-\overline{\varphi}_{x_{r+1}}) = \overline{\varphi}_{x_{r+1}}.$$

Let $\sigma' \in \Delta_{\mathcal{A}}$ be a monitoring strategy satisfying (1). The existence of such a monitoring strategy follows from Farkas' lemma (for instance, see [6, Lemma EC.4]), since

$$\sum_{l=1}^{k^*_C} \left(1 - \frac{k^*_C - r}{(1-\overline{\varphi}_{x_l})S^{k^*}_C}\right) = r,$$

and for every $l \in [\![1, k^*_C]\!]$,

$$1 \geq 1 - \frac{k^*_C - r}{(1-\overline{\varphi}_{x_l})S^{k^*}_C} \geq 1 - \frac{k^*_C - r}{(1-\overline{\varphi}_{x_{k^*_C}})S^{k^*}_C} \geq 0.$$

We next consider $x \in C$. If $x \in \{x_1, \ldots, x_{k^*_C}\}$, then:

$$\begin{aligned}\mathbb{E}_{X\sim\sigma'}[\bar{f}(X,x)] &= \overline{\varphi}_x + (1-\overline{\varphi}_x)\mathbb{P}_{X\sim\sigma'}(x\in X) \\ &= 1 - \frac{k^*_C - r}{S^{k^*}_C}.\end{aligned}$$

If instead $x \in C \setminus \{x_1, \ldots, x_{k^*_C}\}$, then:

$$\begin{aligned}\mathbb{E}_{X\sim\sigma'}[\bar{f}(X,x)] &= \overline{\varphi}_x \geq \overline{\varphi}_{x_{k^*_C+1}} \\ &> 1 - \frac{k^*_C + 1 - r}{S^{k^*}_C + (1-\overline{\varphi}_{x_{k^*_C+1}})^{-1}},\end{aligned}$$

which is equivalent to

$$\mathbb{E}_{X\sim\sigma'}[\bar{f}(X,x)] > 1 - \frac{k^*_C - r}{S^{k^*}_C}.$$

Thus, $z^*_{(\mathcal{R}_C)} \geq 1 - \frac{k^*_C - r}{S^{k^*}_C}$.

Now, let us assume by contradiction that $z^*_{(\mathcal{R}_C)} > 1 - \frac{k^*_C - r}{S^{k^*}_C}$. Let σ^* be an optimal solution of $z^*_{(\mathcal{R}_C)}$. Then, we deduce that:

$$\forall\, x \in C,\ \mathbb{P}_{X\sim\sigma^*}(x\in X) > 1 - \frac{k^*_C - r}{(1-\overline{\varphi}_x)S^{k^*}_C},$$

which induces the following contradiction:

$$\begin{aligned}r &= \sum_{X\in\mathcal{A}} \sigma^*_X |X| = \sum_{x\in\mathcal{X}} \sum_{X\in\mathcal{A}} \sigma^*_X \mathbb{1}_{\{x\in X\}} \\ &= \sum_{x\in\mathcal{X}} \mathbb{P}_{X\sim\sigma^*}(x\in X) > \sum_{l=1}^{k^*_C} \left(1 - \frac{k^*_C - r}{(1-\overline{\varphi}_{x_l})S^{k^*}_C}\right) = r.\end{aligned}$$

In conclusion, $z^*_{(\mathcal{R}_C)} = 1 - \frac{k^*_C - r}{S^{k^*}_C}$ and an optimal solution is given by σ' satisfying (1).

We note that for every $C \subseteq \mathcal{X}$ of size no more than r, $z^*_{(\mathcal{R}_C)} = 1$ (which is obtained by placing a sensor on each location in C). Also in such cases, $\mathcal{U}\setminus U_C \neq \emptyset$ and $1 - \frac{k^*_C - r}{S^{k^*}_C} \geq 1 \geq \min_{u\notin U_C} \varphi_u$. Thus, from Lemma 1, we deduce that

$$z^*_{(\overline{\mathcal{M}})} = \max_{C\subseteq\mathcal{X}} \min\left\{1 - \frac{k^*_C - r}{S^{k^*}_C}, \min_{u\notin U_C} \varphi_u\right\}.$$

Finally, we prove the upper bound on $z^*_{(\mathcal{M})}$. Since $f(X,u) \leq 1$ for every $X \in \mathcal{A}$ and every $u \in \mathcal{U}$, then $z^*_{(\mathcal{M})} \leq 1$. Furthermore, by strong duality, we know that $z^*_{(\mathcal{M})} = \min_{\pi\in\Delta_{\mathcal{U}}} \max_{X\in\mathcal{A}} \mathbb{E}_{u\sim\pi}[f(X,u)]$, where $\Delta_{\mathcal{U}} := \{\pi \in [0,1]^{|\mathcal{U}|} \mid \sum_{u\in\mathcal{U}} \pi_u = 1\}$. Then, given a set packing $T \in \mathcal{T}$, we define $\pi' \in \Delta_{\mathcal{U}}$ as follows:

$$\pi'_u = \begin{cases} \frac{1}{(1-\varphi_u)S_T}, & \text{if } u \in T \\ 0, & \text{otherwise.}\end{cases}$$

We can then derive the following upper bound.

$$\begin{aligned}z^*_{(\mathcal{M})} &\leq \max_{X\in\mathcal{A}} \mathbb{E}_{u\sim\pi'}[f(X,u)] \\ &= \max_{X\in\mathcal{A}} \sum_{u\in T} \pi'_u (1-(1-\varphi_u)\mathbb{1}_{\{u\notin U_X\}}) \\ &= 1 - \frac{1}{S_T} \min_{X\in\mathcal{A}} \sum_{u\in T} \mathbb{1}_{\{u\notin U_X\}} \leq 1 - \frac{|T|-r}{S_T}.\end{aligned}$$

Thus, $z^*_{(\mathcal{M})} \leq \min\{1, \min_{T\in\mathcal{T}}(1 - \frac{|T|-r}{S_T})\}$. □

Proof of Proposition 1. Let (y^*, z^*, ρ^*) be an optimal solution of $(\mathcal{C})$, and we set $C = \operatorname{supp}(y^*)$. Then, $(\mathcal{C})$ can be rewritten as follows:

$$\max_y \min \left\{ \max_\rho \min_{x \in \operatorname{supp}(y)} (\overline{\varphi}_x + (1-\overline{\varphi}_x)\rho_x), \min_{u \notin U_{\operatorname{supp}(y)}} \varphi_u \right\}.$$

$$\text{s.t. } 0 \le \rho_x \le y_x, \qquad \forall x \in \mathcal{X} \tag{19}$$

$$\sum_{x \in \mathcal{X}} \rho_x = r, \tag{20}$$

$$y_x \in \{0,1\}, \qquad \forall x \in \mathcal{X}.$$

The equivalence between $(\overline{\mathcal{M}})$ and $(\mathcal{C})$ then follows from Lemma 1, Equality (2), and the existence of monitoring strategies satisfying Constraints (19)-(20). □

Proof of Proposition 2. We argued in the main text that Problem $(\mathcal{P})$ is indeed a MIP formulation of the following problem $\min_{T \in \mathcal{T}} \frac{S_T}{|T|-r}$. In particular, Constraints (9) and M_2 ensure that at optimality all t_l variables are equal to 0 except for the index corresponding to the set packing size, thus ensuring the validity of the formulation. Then, $\min\{1, \min_{T \in \mathcal{T}}(1 - \frac{|T|-r}{S_T})\} = 1 - \frac{1}{z^*_{(\mathcal{P})}}$. Note that we used the convention that if $(\mathcal{P})$ is infeasible (which occurs if there is no set packing of size greater than r), then $z^*_{(\mathcal{P})} = +\infty$. □

Proof of Proposition 3. We consider an instance where $U_x \cap U_{x'} = \emptyset$ for every $x \neq x' \in \mathcal{X}$. We index the locations in $\mathcal{X} = \{x_1, \dots, x_n\}$ in nondecreasing order of their criticalities. Then, we consider the set packing $T' = \{u_1, \dots, u_{k^*_{\mathcal{X}}}\} \in \mathcal{T}$, where for every $l \in [\![1, k^*_{\mathcal{X}}]\!]$, $\varphi_{u_l} = \overline{\varphi}_{x_l}$. In other words, T' selects the most critical component within each of the first $k^*_{\mathcal{X}}$ monitoring sets. As a consequence, $|T'| = k^*_{\mathcal{X}}$ and $S_{T'} = S^{k^*}_{\mathcal{X}}$. Since $U_{\mathcal{X}} = \mathcal{U}$, then from Theorem 1, we obtain:

$$z^*_{(\mathcal{R}_{\mathcal{X}})} = 1 - \frac{k^*_{\mathcal{X}} - r}{S^{k^*}_{\mathcal{X}}} \le z^*_{(\overline{\mathcal{M}})} \le z^*_{(\mathcal{M})} \le 1 - \frac{|T'| - r}{S_{T'}} = 1 - \frac{k^*_{\mathcal{X}} - r}{S^{k^*}_{\mathcal{X}}}.$$

Thus, $z^*_{(\mathcal{M})} = z^*_{(\overline{\mathcal{M}})}$. Furthermore, any monitoring strategy $\sigma' \in \Delta_{\mathcal{A}}$ satisfying (1) for $C = \mathcal{X}$, which is optimal for $(\mathcal{R}_{\mathcal{X}})$, is then optimal for $(\mathcal{M})$. □

Proof of Proposition 4. We suppose that for every component, the security levels are identical and equal to φ. Let $C \subseteq \mathcal{X}$. Then, we find that $k^*_C = |C|$ since for every $k \in [\![1, |C|]\!]$, $\frac{k-r}{k}(1-\varphi) \le 1 - \varphi$. Therefore, $z^*_{(\overline{\mathcal{M}})} = \max_{C \subseteq \mathcal{X}} \min\{1 - \frac{|C|-r}{|C|}(1-\varphi), \min_{u \notin U_C} \varphi\}$.

Note that for any $C \subseteq \mathcal{X}$, $1 - \frac{|C|-r}{|C|}(1-\varphi) = \varphi + \frac{r}{|C|}(1-\varphi) \ge \varphi$. Therefore,

$$z^*_{(\overline{\mathcal{M}})} = \max_{C \in \mathcal{S}} \left(\varphi + \frac{r}{|C|}(1-\varphi) \right) = \varphi + \frac{r}{n^*}(1-\varphi),$$

where $\mathcal{S}$ is the collection of set covers. Similarly, $S_T = \frac{|T|}{1-\varphi}$ for every $T \in \mathcal{T}$ and we obtain the following:

$$\min_{T \in \mathcal{T}} \left(1 - \frac{|T|-r}{S_T} \right) = \min_{T \in \mathcal{T}} \left(\varphi + (1-\varphi)\frac{r}{|T|} \right) = \varphi + (1-\varphi)\frac{r}{m^*}.$$

□

REFERENCES

[1] J. Slay and M. Miller, "Lessons learned from the Maroochy water breach," in *Proceedings of the International Conference on Critical Infrastructure Protection*, 2007, pp. 73–82.

[2] J. Majkut and A. Dawes, "Responding to Russian attacks on Ukraine's power sector," *Center for Strategic & International Studies*, https://tinyurl.com/h2dm5mk8, Nov 2022.

[3] K. Stouffer, J. Falco, and K. Scarfone, "Guide to industrial control systems security," *National Institute of Standards & Technology*, 2011.

[4] A. Rahmattalabi, P. Vayanos, and M. Tambe, "A robust optimization approach to designing near-optimal strategies for constant-sum monitoring games," in *Decision and Game Theory for Security*. Springer International Publishing, 2018, pp. 603–622.

[5] M. Pirani, E. Nekouei, H. Sandberg, and K. H. Johansson, "A game-theoretic framework for security-aware sensor placement problem in networked control systems," in *Proceedings of the American Control Conference*, 2019, pp. 3699–3706.

[6] M. Dahan, L. Sela, and S. Amin, "Network inspection for detecting strategic attacks," *Operations Research*, vol. 70, no. 2, pp. 1008–1024, 2022.

[7] A. Krause, A. Roper, and D. Golovin, "Randomized sensing in adversarial environments," in *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.

[8] M. Dziubiński and J. Roy, "Hide and seek game with multiple resources," in *International Symposium on Algorithmic Game Theory*. Springer, 2018, pp. 82–86.

[9] 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*, vol. 35, no. 1, pp. 46–65, 2015.

[10] M. Baykal-Gürsoy, Z. Duan, H. V. Poor, and A. Garnaev, "Infrastructure security games," *Eur. J. Oper. Res.*, vol. 239, no. 2, pp. 469–478, 2014.

[11] S. Goyal and A. Vigier, "Attack, defense, and contagion in networks," *Rev. Econ. Stud.*, vol. 81, no. 4, pp. 1518–1542, 2014.

[12] J. Zhuang and V. M. Bier, "Balancing terrorism and natural disasters—defensive strategy with endogenous attacker effort," *Oper. Res.*, vol. 55, no. 5, pp. 976–991, 2007.

[13] 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.

[14] A. R. Hota, A. A. Clements, S. Sundaram, and S. Bagchi, "Optimal and game-theoretic deployment of security investments in interdependent assets," in *International Conference on Decision and Game Theory for Security*. Springer, 2016, pp. 101–113.

[15] G. Brown, M. Carlyle, J. Salmerón, and K. Wood, "Defending critical infrastructure," *INFORMS J. Applied Analytics*, vol. 36, no. 6, pp. 530–544, 2006.

[16] V. M. Bier and N. Haphuriwat, "Analytical method to identify the number of containers to inspect at U.S. ports to deter terrorist attacks," *Ann. Oper. Res.*, vol. 187, no. 1, pp. 137–158, 2011.

[17] D. L. Alderson, G. G. Brown, and W. M. Carlyle, "Operational models of infrastructure resilience," *Risk Anal.*, vol. 35, no. 4, pp. 562–586, 2015.

[18] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood, "Assessing and improving the operational resilience of a large highway infrastructure system to worst-case losses," *Transportation Sci.*, vol. 52, no. 4, pp. 1012–1034, 2018.

[19] A. Washburn and K. Wood, "Two-person zero-sum games for network interdiction," *Oper. Res.*, vol. 43, no. 2, pp. 243–251, 1995.

[20] K. J. Cormican, D. P. Morton, and R. K. Wood, "Stochastic network interdiction," *Oper. Res.*, vol. 46, no. 2, pp. 184–197, 1998.

[21] J. C. Smith and C. Lim, "Algorithms for network interdiction and fortification games," in *Pareto Optimality, Game Theory And Equilibria* (Springer, New York, NY), A. Chinchuluun, P. M. Pardalos, A. Migdalas, and L. Pitsoulis, Eds., 2008, pp. 609–644.

[22] S. Gal and J. Casas, "Succession of hide–seek and pursuit–evasion at heterogeneous locations," *J. Roy. Soc. Interface*, vol. 11, no. 94, p. 20140062, 2014.

[23] L. Hellerstein, T. Lidbetter, and D. Pirutinsky, "Solving zero-sum games using best-response oracles with applications to search games," *Operations Research*, vol. 67, no. 3, pp. 731–743, 2019.

[24] T. Lidbetter, "Search games with multiple hidden objects," *SIAM Journal on Control and Optimization*, vol. 51, no. 4, pp. 3056–3074, 2013.

[25] J. Clarkson, K. Y. Lin, and K. D. Glazebrook, "A classical search game in discrete locations," *Mathematics of Operations Research*, 2022.

[26] S. Alpern and S. Gal, *The theory of search games and rendezvous*. Springer Science & Business Media, 2006, vol. 55.

[27] R. Hohzaki, "Search games: Literature and survey," *Journal of the Operations Research Society of Japan*, vol. 59, no. 1, pp. 1–34, 2016.

[28] S. Gal and J. Casas, "Succession of hide–seek and pursuit–evasion at heterogeneous locations," *Journal of the Royal Society Interface*, vol. 11, no. 94, p. 20140062, 2014.

[29] B. Bahamondes and M. Dahan, "Hide-and-seek game with capacitated locations and imperfect detection," 2023.

[30] V. Tzoumas, A. Jadbabaie, and G. J. Pappas, "Sensor placement for optimal Kalman filtering: Fundamental limits, submodularity, and algorithms," in *Proceedings of the American Control Conference*, 2016.

[31] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.

[32] L. S. 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.

[33] J. Milošević, A. Teixeira, K. H. Johansson, and H. Sandberg, "Actuator security indices based on perfect undetectability: Computation, robustness, and sensor placement," *IEEE Transactions on Automatic Control*, vol. 65, no. 9, pp. 3816–3831, 2020.

[34] 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*, 2008.

[35] B. Bahamondes and M. Dahan, "Network inspection from locations with imperfect detection capabilities," in *2022 American Control Conference (ACC)*, 2022, pp. 613–620.

[36] B. McCann and M. Dahan, "Network inspection using heterogeneous sensors for detecting strategic attacks," in *55th Hawaii International Conference on System Sciences (HICSS)*, 2022, pp. 6802–6811.

[37] M. Pirani, J. A. Taylor, and B. Sinopoli, "Strategic sensor placement on graphs," *Systems & Control Letters*, vol. 148, p. 104855, 2021.

[38] J. Milošević, M. Dahan, S. Amin, and H. Sandberg, "A Network Monitoring Game with Heterogeneous Component Criticality Levels," in *Proc. IEEE 58th Conf. Decis. Control*, Dec 2019.

[39] A. Washburn and K. Wood, "Two-person zero-sum games for network interdiction," *Operations research*, vol. 43, no. 2, pp. 243–251, 1995.

[40] 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, 2016.

[41] PG&E, "PG&E improves gas leak survey capability with heath consultants' advanced detection device," Nov 2012, https://tinyurl.com/4wju87ut.

[42] M. Cardei and D.-Z. Du, "Improving wireless sensor network lifetime through power aware organization," *Wireless Networks*, vol. 11, no. 3, pp. 333–340, 2005.

[43] W.-Q. Wang and H. Shao, "Radar-to-radar interference suppression for distributed radar sensor networks," *Remote Sensing*, vol. 6, no. 1, pp. 740–755, 2014.

[44] R. Powell, "Sequential, nonzero-sum "Blotto": Allocating defensive resources prior to attack," *Games and Economic Behavior*, vol. 67, no. 2, pp. 611–615, 2009.

[45] S. Behnezhad, A. Blum, M. Derakhshan, M. HajiAghayi, M. Mahdian, C. H. Papadimitriou, R. L. Rivest, S. Seddighin, and P. B. Stark, "From battlefields to elections: Winning strategies of Blotto and auditing games," in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2018, pp. 2291–2310.

[46] M. Musegaas, L. Schlicher, and H. Blok, "Stackelberg production-protection games: Defending crop production against intentional attacks," *Eur. J. Oper. Res*, vol. 297, no. 1, pp. 102–119, 2022.

[47] Y. Freund and R. E. Schapire, "Adaptive game playing using multiplicative weights," *Games and Economic Behavior*, vol. 29, no. 1, pp. 79–103, 1999.

[48] C.-Y. Chong and S. P. Kumar, "Sensor networks: evolution, opportunities, and challenges," *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1247–1256, Aug 2003.

[49] J. C. von Fischer, D. Cooley, S. Chamberlain, A. Gaylord, C. J. Griebenow, S. P. Hamburg, J. Salo, R. Schumacher, D. Theobald, and J. Ham, "Rapid, vehicle-based identification of location and magnitude of urban natural gas pipeline leaks," *Environmental Science & Technology*, vol. 51, no. 7, pp. 4091–4099, 04 2017.

[50] A. Aisopou, I. Stoianov, and N. J. D. Graham, "In-pipe water quality monitoring in water supply systems under steady and unsteady state flow conditions: A quantitative assessment," *Water Research*, vol. 46, no. 1, pp. 235 – 246, 2012.

[51] USEPA, *EPANET 2.00.12*. Cincinnati, Ohio: U.S. Environmental Protection Agency, 2002, http://www2.epa.gov/water-research/epanet, Accessed: 2014-10-24.

[52] L. Perelman, J. Arad, M. Housh, and A. Ostfeld, "Event detection in water distribution systems from multivariate water quality time series," *Environmental Science & Technology*, vol. 46, no. 15, pp. 8212–8219, 2012.

[53] R. S. Ross, "Guide for conducting risk assessments," *NIST Special Publication*, vol. 800, no. 30, 2012.

[54] M. D. Jolly, A. D. Lothes, L. Sebastian Bryson, and L. Ormsbee, "Research database of water distribution system models," *Journal of Water Resources Planning and Management*, 2013.

[55] J. Milošević, "Security metrics and allocation of security resources for control systems," Ph.D. dissertation, KTH Royal Institute of Technology, 2020.

[56] V. Tzoumas, M. A. Rahimian, G. J. Pappas, and A. Jadbabaie, "Minimal actuator placement with bounds on control effort," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 1, pp. 67–78, 2016.

[57] F. Pasqualetti, S. Zampieri, and F. Bullo, "Controllability metrics, limitations and algorithms for complex networks," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 40–52, 2014.

[58] A. R. Bergen and D. J. Hill, "A structure preserving model for power system stability analysis," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-100, no. 1, pp. 25–35, Jan 1981.

[59] A. Mohsenian-Rad and A. Leon-Garcia, "Distributed internet-based load altering attacks against smart power grids," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 667–674, Dec 2011.

Jezdimir Milošević received the M.Sc. degree in electrical engineering and computer science from the University of Belgrade, Belgrade, Serbia, in 2015 and the Ph.D. in electrical engineering and computer science from the KTH Royal Institute of Technology, Stockholm, Sweden, in 2020.

He is currently a research and development engineer in Scania's Autonomous Transport Solutions, Södertälje, Sweden. His research interests are within cyber-security of industrial control systems.

Mathieu Dahan received the M.S. and Ph.D. degrees in computational science and engineering from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2016 and 2019, respectively.

He is currently an Assistant Professor with the School of Industrial and Systems Engineering at the Georgia Institute of Technology, Atlanta, GA, USA. His research interests are in game theory as well as network and combinatorial optimization to improve the security and resilience of large-scale systems against uncertainties, strategic attacks, and natural disasters.

Saurabh Amin received the Ph.D. degree in systems engineering from the University of California, Berkeley, Berkeley, CA, USA, in 2011.

He is currently a Professor with the Department of Civil and Environmental Engineering at the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. He is a member of the Laboratory for Information and Decision Systems (LIDS) and the Operations Research Center (ORC) at MIT. His fields of expertise include control and optimization, applied game theory, and networks. His research focuses on the design and implementation of resilient monitoring and control algorithms for networked infrastructure systems. He serves on the editorial board of IEEE Transactions on Automatic Control.

Henrik Sandberg (F'23) received the M.Sc. degree in engineering physics and the Ph.D. degree in automatic control from Lund University, Lund, Sweden, in 1999 and 2004, respectively.

He is currently a Professor with the Division of Decision and Control Systems, KTH Royal Institute of Technology, Stockholm, Sweden. His current research interests include security of cyber-physical systems, power systems, model reduction, and fundamental limitations in control.

Dr. Sandberg was a recipient of the Best Student Paper Award from the IEEE Conference on Decision and Control in 2004, an Ingvar Carlsson Award from the Swedish Foundation for Strategic Research in 2007, and a Consolidator Grant from the Swedish Research Council in 2016. He has served on the editorial boards of IEEE Transactions on Automatic Control and the IFAC Journal Automatica.