Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Network Inspection for Detecting Strategic Attacks

Mathieu Dahan, Lina Sela, Saurabh Amin

To cite this article:
Mathieu Dahan, Lina Sela, Saurabh Amin (2022) Network Inspection for Detecting Strategic Attacks. Operations Research 70(2):1008-1024. https://doi.org/10.1287/opre.2021.2180

http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 70, No. 2, March–April 2022, pp. 1008–1024
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

Network Inspection for Detecting Strategic Attacks

Mathieu Dahan,[a] Lina Sela,[b] Saurabh Amin[c]

[a] School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332; [b] Department of Civil, Architectural and Environmental Engineering, The University of Texas at Austin, Austin, Texas 78712; [c] Department of Civil and Environmental Engineering, Laboratory for Information and Decision Systems and Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Contact: mathieu.dahan@isye.gatech.edu, https://orcid.org/0000-0003-0973-6973 (MD); linasela@utexas.edu, https://orcid.org/0000-0002-5834-8451 (LS); amins@mit.edu, https://orcid.org/0000-0003-1554-015X (SA)

Received: October 11, 2018
Revised: June 24, 2020; March 21, 2021
Accepted: April 21, 2021
Published Online in Articles in Advance: January 28, 2022

Area of Review: Policy Modeling and Public Sector OR

https://doi.org/10.1287/opre.2021.2180

Abstract. This article studies a problem of strategic network inspection, in which a defender (agency) is tasked with detecting the presence of multiple attacks in the network. An inspection strategy entails monitoring the network components, possibly in a randomized manner, using a given number of detectors. We formulate the network inspection problem ($\mathcal{P}$) as a large-scale bilevel optimization problem, in which the defender seeks to determine an inspection strategy with minimum number of detectors that ensures a target expected detection rate under worst-case attacks. We show that optimal solutions of ($\mathcal{P}$) can be obtained from the equilibria of a large-scale zero-sum game. Our equilibrium analysis involves both game-theoretic and combinatorial arguments and leads to a computationally tractable approach to solve ($\mathcal{P}$). First, we construct an approximate solution by using solutions of minimum set cover (MSC) and maximum set packing (MSP) problems and evaluate its detection performance. In fact, this construction generalizes some of the known results in network security games. Second, we leverage properties of the optimal detection rate to iteratively refine our MSC/MSP-based solution through a column generation procedure. Computational results on benchmark water networks demonstrate the scalability, performance, and operational feasibility of our approach. The results indicate that utilities can achieve a high level of protection in large-scale networks by strategically positioning a small number of detectors.

Funding: This work was supported by the National Science Foundation–Division of Computer and Network Systems [Grant 1239054 and CAREER Award 1453126], University of Texas at Austin New Faculty Start Up Grant, and a Massachusetts Institute of Technology Schoettler Fellowship.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/opre.2021.2180.

Keywords: strategic network inspection • attack detection • multiple resources • large-scale bilevel optimization • equilibrium analysis

1. Introduction

Ensuring the security of critical infrastructures such as water, oil and gas, and power distribution systems is crucial for the welfare and prosperity of our society. These infrastructure networks span huge geographical areas and are inherently vulnerable to both intentional and unintentional threats. In most jurisdictions, public utilities and municipalities are the primary entities responsible for ensuring the infrastructure reliability and service quality and use various degrees of oversight to manage and respond to emergency situations. In recent years, numerous incidents have been reported that highlight the inherent vulnerability of infrastructure networks to adversarial events (Owolabi 2016, Dancy and Dancy 2017, Naureen et al. 2018). Such events often result in recurrent service interruptions and in some cases even pose significant danger to human lives (Yuhas 2016). In response, governments and public utility commissions are developing new policies and regulations that charge the utilities to proactively recognize the security risks to their infrastructure and develop specific capabilities to reduce them (Barrett 2018). Public sector operations research can provide new solutions to guide and support the utilities in such risk assessment and mitigation activities.

In this article, we study a network inspection problem that exploits the capabilities of modern sensing and event detection technology to monitor an infrastructure network against strategic attacks. Our objective is to design inspection strategies that can effectively detect adversarial failure events in a large-scale network and hence limit and reduce the operational losses faced by utilities because of undetected events. The underlying technological motivation is the commercial availability of smart detectors that can be easily operated by the utility personnel and flexibly positioned at different locations in the network (PG&E 2010). These detectors are integrated systems with advanced capabilities (Phillips et al. 2013,

Xing and Sela 2019) such as (i) on-board sensing to collect state measurements at fine temporal resolution; (ii) accurate and timely detection of faulty events using data analytics; and (iii) real-time communication to transmit data and event alerts to remote utility personnel. Our approach can help utilities leverage these capabilities in detecting targeted and/or random disruptions by providing useful guidelines for selecting the number of detectors, monitoring locations, and schedule of inspection operations in order to satisfy a target detection performance.

Specifically, we consider a bilevel optimization formulation of the strategic network inspection problem and focus on the question: *How many detectors are required and how to position them in the network to detect multiple adversarial attacks?* In our formulation, we assume that the set of locations that can be used for monitoring the network and the set of network components that can be accessed and targeted by the attacker (malicious entity) are predefined. The utility (defender) aims to minimize the number of detectors to achieve a desirable attack-detection performance, whereas the attacker seeks to avoid detection (i.e., maximize the number of undetected attacks). Importantly, we allow the choice of randomized inspection strategies by the defender, which is a departure from the traditional fixed sensing paradigm. Practically, randomized inspection entails shifting and/or mobilizing the available detectors over a subset of locations in the network. In fact, randomized strategies are known to be an effective defense mechanism in various applications, as listed in Table 1. However, such strategies cannot be adopted in practice unless they are simple and cost-effective to execute, and provide strong performance guarantees. In this article, we focus on inspection strategies that have these desirable features.

1.1. Main Contributions

In Section 2, we introduce a generic detection model that captures the key features of modern inspection systems with respect to sensing technology for event detection and flexibility of positioning. We use this detection model to formulate the bilevel optimization problem, denoted ($\mathcal{P}$), in which the defender first selects a randomized positioning of detectors, and the attacker responds by targeting one or more network components. The attacker seeks to maximize the expected number of undetected attacks, whereas the defender aims to minimize the number of detectors required to ensure that the expected detection rate under worst-case attacks is above a prespecified threshold.

Our approach to solve the problem ($\mathcal{P}$) involves analyzing the equilibrium properties of a zero-sum game, denoted Γ, where the defender (respectively, attacker) seeks to minimize (respectively, maximize) the expected number of undetected attacks (Proposition 2). However, the sets of players' actions in Γ grow combinatorially with the size of the network, thus making the equilibrium computation challenging in itself. In Section 3, we derive structural properties that are satisfied by all Nash equilibria of Γ. We present these properties for the most conservative case when the attacker has the ability to spread her attacks across the network. In particular, we show that in any equilibrium of Γ, both players must randomize their actions and use all available resources, and every network component must be monitored with positive probability (Theorem 1 and Proposition 3). Additionally, we prove the important, and rather surprising property, that the expected detection rate and the inspection strategies in equilibrium do not depend on the attacker's number of resources (Theorem 2). This implies that the defender does not need to know precisely the amount of attack resources in order to monitor the network. The proofs of these game-theoretic results rely on linear programming duality in zero-sum games, submodularity of the detection function, as well as the minimum set cover (MSC) and maximum set packing (MSP) problems, which, respectively, capture the "coverage" and "spread" of the network.

Our equilibrium analysis leads to a novel approach to solve the inspection problem ($\mathcal{P}$). First, we obtain lower and upper bounds on the optimal expected detection rate in terms of the number of available detectors, and the optimal values of the MSC and MSP problems. A preliminary and specialized version of this result was presented in Dahan et al. (2016). Furthermore, we construct an inspection strategy that randomizes the positioning of detectors over an MSC, and derive guarantees on the resulting expected detection performance. This provides us with an approximate solution to ($\mathcal{P}$) and optimality gap that can be

Table 1. Applications of the Network Inspection Problem

Inspection setting	Network	Type of detector	Type of attacks
Urban patrolling	City streets	Police unit	Robbery
Network security	Information network	Firewall	Cyberattack on server
Sensing of gas/water networks	Gas/water pipelines	Leak/pressure sensor	Pipe disruption
Interdiction of illegal goods	Transportation network	Police officer	Drug trafficking
Infiltration game	Water channel	Electric cable	Malicious infiltration

computed by solving the MSC and MSP problems. It turns out this solution is optimal for the special case when the MSCs and MSPs are of same size. Second, a consequence of the equilibrium properties is that a column generation–based procedure (Algorithm 1) can be used to iteratively improve our MSC/MSP-based solution to optimality.

Although our approach to solve problem ($\mathcal{P}$) relies on the MSC and MSP problems that are known to be NP-hard, we find that modern integer programming solvers can solve large instances of these problems. In Section 5, we demonstrate the benefits of our solution approach in monitoring large-scale urban water networks facing adversarial disruptions. Our computational study shows that the MSC/MSP-based solution is scalable, provides good performance guarantees, and is easily implementable by the defender. On the other hand, we find that implementing the optimal inspection strategy requires a much higher number of monitored locations and a more complex scheduling of operations. Furthermore, it only provides a marginal improvement in comparison with the simpler MSC-based inspection strategy. Thus, our approach can be used for designing inspection strategies that achieve a desired tradeoff between detection performance and operational feasibility.

The complete proofs of our results, as well as additional examples, are provided in the online appendix.

1.2. Related Work

Our detection model is inspired by modern sensing technology used in detecting leaks and other failure events in pipeline networks for distribution of natural gas (Phillips et al. 2013) and water (Ostfeld and Salomons 2004, Sela Perelman et al. 2016). The dominant paradigm in sensing of these infrastructure networks is to optimally place a limited number of sensors for maximizing a metric of detection performance (Berry et al. 2006, Krause et al. 2008a, Chakrabarti et al. 2009). A myriad of models for the sensor placement problem have been proposed in the literature, including robust formulations (Sela and Amin 2018); for example, Krause et al. (2008b) proposed an efficient approximation algorithm to maximize the worst-case detection performance against a set of possible failure scenarios. More recently, Tzoumas et al. (2017) and Orlin et al. (2018) designed approximation algorithms to find a sensor placement that is robust against a subset of sensors' failures. The main feature of this line of work is *fixed sensing*, that is, continuous operation of sensors placed at fixed locations in the network.

However, in large-scale networks, a fixed strategy implemented by a resource-constrained utility inevitably leaves some parts of the network unmonitored. In an adversarial situation, a strategic attacker will target these unmonitored parts to avoid or make their detection more difficult. It follows that a fixed inspection strategy can lead to a significant loss of detection performance and, in turn, compromise the overall security of the infrastructure system. In contrast, it has been shown that *randomized strategies* can significantly improve the defender's performance against worst-case disruptions or adversarial failure events (Washburn and Wood 1995, Bertsimas et al. 2016). Practically, randomized inspection strategies can be translated into random scheduling of inspections that can be performed on a day-to-day basis by utility personnel. For example, Pita et al. (2008) use randomized strategies for the scheduling of checkpoints and for generating patrolling schedules for canine units to assist the police at the Los Angeles International Airport. Hochbaum and Fishbain (2011) investigate the allocation of mobile sensors to detect transported nuclear weapons based on related radiological dispersion devices. Finally, water utilities routinely sample water quality at random locations in the distribution system to comply with safety standards such as the Safe Drinking Water Act (SDWA) rules (Tiemann 2017).

In the context of network security, several models have been proposed for the strategic allocation of defense resources (Zhuang and Bier 2007, Baykal-Gürsoy et al. 2014, Goyal and Vigier 2014). For instance, Brown et al. (2006), Bier and Haphuriwat (2011), and Alderson et al. (2015, 2018) consider bilevel and trilevel optimization problems to model defender-attacker interactions where each player selects a pure strategy. In contrast, our setting involves randomized strategies, and the combinatorial size of the sets of players' actions does not enable us to solve problem ($\mathcal{P}$) using mixed-integer linear programming techniques. Other models that have been studied include search games (Gal and Casas 2014) and inspection problems (Washburn and Wood 1995, Cormican et al. 1998, Smith and Lim 2008). In addition, Powell (2007), Bier et al. (2008), and Zhuang et al. (2010) investigate equilibria in security games with an asymmetric information structure. Conversely, our model considers detection capabilities of the defender and multiple attacks spread over a general network under a complete (and symmetric) information structure.

The zero-sum game Γ we analyze for solving problem ($\mathcal{P}$) is more general than the classical *hide-and-seek* game first introduced by Von Neumann (1953) and further discussed in chapter 3.2 of Karlin and Peres (2016). In this game, a robber hides in one of a set of "safe houses" located at intersections of vertical and horizontal roads, and a police unit simultaneously chooses to travel along one road to find the robber. Our equilibrium analysis can be applied to solve the generalized hide-and-seek game, which involves multiple police units patrolling in a complex street network to find multiple robbers. Related to our setting is also the work by

Mavronicolas et al. (2008), who consider a security game on a bipartite information network in which servers are vulnerable to multiple attacks and the defender can install a firewall to protect a subnetwork. In fact, our analysis approach can be used to derive a more sophisticated defense strategy that installs multiple firewalls to secure more complex information networks against multiple simultaneous attacks. Our game similarly generalizes the *patrolling game* studied in Alpern et al. (2011), and the *infiltration games* defined in Garnaev et al. (1997) and in chapter 2.1 of Garnaev (2000) by considering multiple player resources and more complex network systems.

2. Problem Description

In this section, we introduce a generic formulation for the strategic network inspection problem. Our formulation is a bilevel optimization model of sequential defender-attacker interaction on an infrastructure network, with each player having access to multiple resources.

2.1. Defender and Attacker Models

We consider the setting where a defender (utility) is tasked with inspecting an infrastructure network that transports a commodity (e.g., water, oil, natural gas). The network faces a risk of disruptions by an attacker (malicious entity) who can compromise the operational functionality of the set of network components, denoted $\mathcal{E}$. To inspect the network and monitor its components, the defender positions a set of *detectors* on a set of locations (or nodes), denoted $\mathcal{V}$. Each detector is an integrated system comprising an on-board sensing unit, detection software, and communication unit (Chong and Kumar 2003). The defender (or the utility's employees) can flexibly mobilize the detectors from one node to another. For our purposes, the sets $\mathcal{E}$ and $\mathcal{V}$ are predefined.

For example, in the context of a municipal water network, the set $\mathcal{E}$ represents the system's components that can be accessed and targeted by an attacker, and the set $\mathcal{V}$ represents the access points where detectors can be deployed (e.g., manholes, valves, or fire hydrants). Targeted physical or remote attacks to the network can induce damage to pipelines and valves or backflow at fire hydrants (Monroe et al. 2018, Hassanzadeh et al. 2020). Such disruption events typically result in local perturbations in the state variables (water flow rate and pressure) that progressively propagate to other parts of the network. If the water utility has positioned a detector at a node that experiences perturbations from a disruption event, the on-board sensing unit can measure the change in state variables (Allen et al. 2011, Srirangarajan et al. 2013). These measurements can then be processed to detect the occurrence of the event. Clearly, the ability of the defender to detect such disruption events depends on how the available detectors are positioned in the network.

Formally, when a detector is positioned at node $i \in \mathcal{V}$ by the defender, the following steps govern its attack detection capability: First, the sensing unit collects relevant state measurements from node i. These measurements capture the state of a subset of components $\mathcal{C}_i \in 2^{\mathcal{E}}$, that is, the detector at node i monitors the components in $\mathcal{C}_i$. Second, the detection software processes these measurements and generates a diagnostic signal indicating the number of disruption events (or attacks) present within the component set $\mathcal{C}_i$. Third, the communication unit transmits the diagnostic signal to the defender. We assume that the cost of data collection, processing, and transmission is negligible in comparison with the cost of procuring the detector. For a detector positioned at node $i \in \mathcal{V}$, we refer to the set $\mathcal{C}_i$ as a *monitoring set*, because under the aforementioned setting, an attack to any component $e \in \mathcal{E}$ can be detected if and only if $e \in \mathcal{C}_i$. The tuple $\mathcal{G} := (\mathcal{V}, \mathcal{E}, \{\mathcal{C}_i, i \in \mathcal{V}\})$ represents the *detection model* of the network. Without loss of generality, we assume that each component in $\mathcal{E}$ can be monitored from at least one node in $\mathcal{V}$.

Importantly, we consider that the defender has access to only a limited number of detectors for network inspection. This limitation results from the economic and operational constraints of the defender. For simplicity, we also suppose that all detectors are homogeneous in terms of their monitoring and detection capabilities, and cost. Let $b_1 \in \mathbb{N}$ be the number of available detectors that can be simultaneously positioned on distinct nodes in $\mathcal{V}$. We denote a *detector positioning* by a set $S \in 2^{\mathcal{V}}$ of nodes that receive detectors. The set of feasible detector positionings is then defined by $\mathcal{A}_1 := \{S \in 2^{\mathcal{V}} \mid |S| \le b_1\}$. For a given detector positioning $S \in \mathcal{A}_1$, let $\mathcal{C}_S := \cup_{i \in S} \mathcal{C}_i$ denote the set of components that are monitored by at least one detector in S.

To count the number of components in any given subset of components of $\mathcal{E}$ that can be monitored using an arbitrary detector positioning, we define a *detection function* $\mathrm{F} : 2^{\mathcal{V}} \times 2^{\mathcal{E}} \to \mathbb{N}$. For a detector positioning $S \in 2^{\mathcal{V}}$ and a subset of components $T \in 2^{\mathcal{E}}$, the value of $\mathrm{F}(S, T)$ is the number of components of T that are monitored by at least one detector positioned in S, that is,

$$\forall (S, T) \in 2^{\mathcal{V}} \times 2^{\mathcal{E}}, \quad \mathrm{F}(S, T) := |\mathcal{C}_S \cap T| \,. \tag{1}$$

Under our detection model, if the components of T face attack-induced disruptions, the number of attacks detected by the detector positioning S is

$\mathrm{F}(S,T)$. The detection function satisfies two natural properties:

i. For any subset of components $T \in 2^{\mathcal{E}}$, $\mathrm{F}(\cdot,T)$ is submodular and monotone:

$$\forall T \in 2^{\mathcal{E}},\ \forall (S,S') \in (2^{\mathcal{V}})^2,$$
$$\begin{cases} \mathrm{F}(S \cup S',T) + \mathrm{F}(S \cap S',T) \le \mathrm{F}(S,T) + \mathrm{F}(S',T), \\ S \subseteq S' \Rightarrow \mathrm{F}(S,T) \le \mathrm{F}(S',T). \end{cases}$$

That is, adding a detector to a smaller detector positioning increases the number of monitored components in T by at least as many as when adding that detector to a larger detector positioning.

ii. For any detector positioning $S \in 2^{\mathcal{V}}$, $\mathrm{F}(S,\cdot)$ is finitely additive (a direct consequence of (1)):

$$\forall S \in 2^{\mathcal{V}},\ \forall (T,T') \in (2^{\mathcal{E}})^2 \mid T \cap T' = \emptyset,$$
$$\mathrm{F}(S,T \cup T') = \mathrm{F}(S,T) + \mathrm{F}(S,T').$$

Similar to the defender, the attacker is also resource constrained, in that she can attack a subset of network components $T \in 2^{\mathcal{E}}$ of size no larger than $b_2 \in \mathbb{N}$; we refer to such a subset as an *attack plan*. This constraint models the attacker's limited ability to gain access to network components and disrupt them. The set of all attack plans is given by $\mathcal{A}_2 := \{T \in 2^{\mathcal{E}} \mid |T| \le b_2\}$.

In fact, our solution approach and results can be extended to the model of imperfect detection, where each detector only detects a disruption in its monitoring set with independent probability $\lambda \in [0,1]$. Given a detector positioning $S \in 2^{\mathcal{V}}$, and an attack plan $T \in 2^{\mathcal{E}}$, the average number of detected attacks would be given by $\sum_{e \in T}(1-(1-\lambda)^{|\{i \in S \mid e \in C_i\}|})$, which is also submodular and monotone with respect to S (see chapter 2 of Fujishige 2005) and finitely additive with respect to T. For ease of exposition, we henceforth assume the model of perfect detection, given by (1).

2.2. Network Inspection Problem

We are now in the position to introduce our network inspection problem, which we define as a bilevel optimization model. In this problem, the defender (referred to as player 1 or P1) first chooses an inspection strategy to monitor network components using a minimum number of detectors. After observing the defender's action, the attacker (referred to as player 2 or P2) targets one or more components to induce disruption events. A typical assumption in infrastructure defense is that of an informed attacker who knows the defender's capabilities. Thus, we assume that both players know the detection model $\mathcal{G}$. At this stage, we also assume that the defender knows the number of attack resources b_2, although we will later show that our solution to the network inspection problem does not depend on it.

The detector positionings (respectively, attack plans) are realized from a chosen probability distribution on the set $\mathcal{A}_1$ (respectively, $\mathcal{A}_2$). Specifically, the defender and attacker respectively choose an inspection strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$ and an attack strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$, where $\Delta(\mathcal{A}_1) := \{\sigma^1 \in [0,1]^{|\mathcal{A}_1|} \mid \sum_{S \in \mathcal{A}_1} \sigma^1_S = 1\}$ and $\Delta(\mathcal{A}_2) := \{\sigma^2 \in [0,1]^{|\mathcal{A}_2|} \mid \sum_{T \in \mathcal{A}_2} \sigma^2_T = 1\}$ denote the mixed strategy sets. Here, σ^1_S (respectively, σ^2_T) represents the probability assigned to the detector positioning S (respectively, attack plan T) by the defender's strategy σ^1 (respectively, the attacker's strategy σ^2).

For ease of exposition, we denote $\mathrm{F}(i,e) := \mathrm{F}(\{i\},\{e\})$ for all $(i,e) \in \mathcal{V} \times \mathcal{E}$. We will also refer to the degenerate mixed-strategies $\mathbb{1}_{\{S\}} \in \Delta(\mathcal{A}_1)$ and $\mathbb{1}_{\{T\}} \in \Delta(\mathcal{A}_2)$ as S and T, respectively. The *support* of $\sigma^1 \in \Delta(\mathcal{A}_1)$ (respectively, $\sigma^2 \in \Delta(\mathcal{A}_2)$) is defined as $\mathrm{supp}(\sigma^1) = \{S \in \mathcal{A}_1 \mid \sigma^1_S > 0\}$ (respectively, $\mathrm{supp}(\sigma^2) = \{T \in \mathcal{A}_2 \mid \sigma^2_T > 0\}$). The *node basis* of a strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$, denoted $\mathcal{V}_{\sigma^1} := \{i \in \mathcal{V} \mid \mathbb{P}_{\sigma^1}(i \in S) > 0\}$, is the set of nodes that are inspected with nonzero probability by the defender. Analogously, the *component basis* of a strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$, denoted $\mathcal{E}_{\sigma^2} := \{e \in \mathcal{E} \mid \mathbb{P}_{\sigma^2}(e \in T) > 0\}$, is the set of components that are targeted with positive probability by the attacker.

We now present the inner and outer problem of our bilevel optimization model.

2.2.1. Inner Problem.

In our model, the attacker responds to the defender's inspection strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$ by choosing an attack strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$, with the objective of maximizing the expected number of attacks that remain undetected by the defender, given by

$$U(\sigma^1,\sigma^2) := \mathbb{E}_{(\sigma^1,\sigma^2)}[|T| - \mathrm{F}(S,T)]. \tag{2}$$

We denote $B_2(\sigma^1,b_2) := \arg\max_{\sigma^2 \in \Delta(\mathcal{A}_2)} U(\sigma^1,\sigma^2)$ the set of attack strategies that are best responses to σ^1. We note that for every inspection strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$, at least one attack plan $T \in \mathcal{A}_2$ is a best response to σ^1.

2.2.2. Outer Problem.

The defender seeks to minimize the number of detectors and also ensure that her chosen inspection strategy provides a certain level of detection performance against the attacker's best response strategy. We use the following metric of detection performance: For a given strategy profile $\sigma \in \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2)$, the *expected detection rate*, denoted $r(\sigma)$, is the expectation (under σ) of the ratio between the number of attacks that are detected and the total number of attacks:

$$r(\sigma) := \mathbb{E}_\sigma\left[\frac{\mathrm{F}(S,T)}{|T|}\right]. \tag{3}$$

Thus, the defender aims to find a minimum-resource inspection strategy while ensuring that the expected

detection rate is no less than a prespecified threshold level $\alpha \in [0,1]$ against worst-case attack strategies. This can be written as the following network inspection problem:

$$(\mathcal{P}): \quad \underset{b_1,\,\sigma^1}{\text{minimize}} \quad b_1$$
$$\text{subject to} \quad r(\sigma^1,\sigma^2) \geq \alpha, \quad \forall \sigma^2 \in B_2(\sigma^1,b_2) \qquad (4)$$
$$\sigma^1 \in \Delta(\mathcal{A}_1)$$
$$b_1 \in \mathbb{N}.$$

Specifically, Constraints (4) ensure that for a given number of attack resources b_2, the expected detection rate induced by the chosen number of detectors b_1 and their randomized positioning σ^1 is at least α under the attacker's best response to σ^1. The target detection rate α reflects the performance requirement that the defender's inspection strategy must satisfy (e.g., because of a regulatory imposition). We denote b_1^* the optimal value of $(\mathcal{P})$.

More generally, the problem $(\mathcal{P})$ captures some of the key features of network inspection in strategic settings; see Table 1 for a comparison of various applications. First, the detection model $\mathcal{G}$ is generic in that it represents the detection capability of the defender, without making further modeling assumptions on the dependence of the monitoring sets $\mathcal{C}_i$ $(i \in \mathcal{V})$ on specific aspects such as the sensing technology used by detectors, the different means that the attacker may use in targeting a component, and the network's topological structure. Second, it considers multiple resources on the part of both players. This is a particularly desirable feature when the attacker can simultaneously attack multiple components across the network, and the defender's inspection involves positioning multiple detectors in order to monitor a large number of critical components. However, $(\mathcal{P})$ is a challenging problem to solve. Indeed, bilevel optimization problems are known to be NP-hard (Hansen et al. 1992), and in our case the number of possible detector positionings grows combinatorially with the number of available detectors and the size of $\mathcal{V}$. Thus, we must leverage structural properties of the problem to solve it in a scalable manner.

3. Game-Theoretic Analysis

In this section, we derive the key properties satisfied by optimal solutions of our network inspection problem $(\mathcal{P})$. We start by studying P2's best response function B_2 and analyze the corresponding zero-sum game. This will in turn help derive a scalable solution approach to $(\mathcal{P})$.

3.1. Zero-Sum Game

Given the detection model $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \{\mathcal{C}_i, i \in \mathcal{V}\})$, and the players' resources b_1 and b_2, we consider the zero-sum game in normal form $\Gamma(b_1,b_2) := \langle \{1,2\}, (\Delta(\mathcal{A}_1), \Delta(\mathcal{A}_2)), (-U,U)\rangle$. In this game, P1 (respectively, P2) selects an inspection strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$ (respectively, an attack strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$) and seeks to minimize (respectively, maximize) the expected number of undetected attacks (2).

A strategy profile $(\sigma^{1^*}, \sigma^{2^*}) \in \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2)$ is a mixed strategy *Nash equilibrium* (NE) of the game $\Gamma(b_1,b_2)$ if for every $(\sigma^1,\sigma^2) \in \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2)$,

$$U(\sigma^{1^*},\sigma^2) \leq U(\sigma^{1^*},\sigma^{2^*}) \leq U(\sigma^1,\sigma^{2^*}). \qquad (5)$$

We denote the set of NE of the game $\Gamma(b_1,b_2)$ as $\Sigma(b_1,b_2)$. Also, when there is no confusion, we simply refer to $\Gamma(b_1,b_2)$, $\Sigma(b_1,b_2)$, and $B_2(\sigma^1,b_2)$ as Γ, Σ, and $B_2(\sigma^1)$, respectively.

Because Γ is a zero-sum game, the set of NE Σ can be obtained by solving the following pair of dual linear programming problems:

$$(\text{LP}_1) \quad \min_{\sigma^1 \in \Delta(\mathcal{A}_1)} \max_{T \in \mathcal{A}_2} U(\sigma^1,T)$$
$$(\text{LP}_2) \quad \max_{\sigma^2 \in \Delta(\mathcal{A}_2)} \min_{S \in \mathcal{A}_1} U(S,\sigma^2).$$

We refer to the optimal value of (LP_1) and (LP_2) as the value of the game $\Gamma(b_1,b_2)$, denoted by $U^*(b_1,b_2)$. In principle, linear programming techniques can be used to compute NE of Γ. However, the computation of (LP_1) and (LP_2) quickly becomes intractable as the size of the network increases. In particular, because of the size of the players' sets of actions ($|\mathcal{A}_1| = \sum_{k=0}^{b_1} \binom{|\mathcal{V}|}{k}$ and $|\mathcal{A}_2| = \sum_{l=0}^{b_2} \binom{|\mathcal{E}|}{l}$), the number of variables and constraints in both linear programs can be huge. For example, for a network consisting of 200 nodes and components, and $b_1 = b_2 = 10$, computing the equilibria of game $\Gamma(b_1,b_2)$ entails solving linear programs containing $2.37 \cdot 10^{16}$ variables and constraints. For large bimatrix games, Lipton et al. (2003) provide an algorithm to compute an ϵ−NE in $n^{O(\ln n/\epsilon^2)}$ time, where n is the number of pure strategies available to each player. However, for realistic instances of the game Γ, this number can easily reach values for which their algorithm is practically inapplicable. Similarly, multiplicative weights update algorithms (Freund and Schapire 1999, Hellerstein et al. 2019) cannot be used to solve the large-scale game Γ.

Next, we develop new results to study the equilibrium characteristics of the game $\Gamma(b_1,b_2)$, given any parameters b_1 and b_2. Our equilibrium characterization uses two combinatorial optimization problems, formulated as minimum set cover and maximum set packing problems. This characterization enables us to analyze the detection performance in equilibrium, which in turn reveals properties satisfied by optimal solutions of our problem $(\mathcal{P})$.

3.2. Set Cover and Set Packing Problems

We say that a set of nodes $S \in 2^{\mathcal{V}}$ is a *set cover* if and only if every component in $\mathcal{E}$ can be monitored by at least one detector positioned in S, that is, $\mathrm{F}(S,e) = 1$, for all $e \in \mathcal{E}$. A set of nodes $S \in 2^{\mathcal{V}}$ is a *minimal set cover* if S is a set cover that is minimum with respect to inclusion; that is, if any node of S is removed, the resulting set is not a set cover anymore. A set of nodes $S \in 2^{\mathcal{V}}$ is a *minimum set cover* (MSC) if and only if it is an optimal solution of the following problem:

$$(\mathcal{I}_{\mathrm{MSC}}): \quad \underset{S\in 2^{\mathcal{V}}}{\text{minimize}} \ |S| \quad \text{subject to} \ \mathrm{F}(S,e) = 1, \quad \forall\, e \in \mathcal{E}. \tag{6}$$

Solving $(\mathcal{I}_{\mathrm{MSC}})$ amounts to determining the minimum number of detectors and their positioning to monitor all network components. We denote the set (respectively, the size) of MSCs by $\mathcal{S}$ (respectively, n^*). Because we assumed that each component can be monitored from at least one node in the network (Section 2.1), $(\mathcal{I}_{\mathrm{MSC}})$ is feasible and n^* exists.

We say that a set of components $T \in 2^{\mathcal{E}}$ is a *set packing* if and only if a detector positioned at any node i can monitor at most one component in T, that is, $\mathrm{F}(i,T) \le 1$, for all $i \in \mathcal{V}$. A set of components $T \in 2^{\mathcal{E}}$ is a maximum set packing (MSP) if and only if it optimally solves the following problem:

$$(\mathcal{I}_{\mathrm{MSP}}): \quad \underset{T\in 2^{\mathcal{E}}}{\text{maximize}} \ |T| \quad \text{subject to} \ \mathrm{F}(i,T) \le 1, \quad \forall\, i \in \mathcal{V}. \tag{7}$$

Solving $(\mathcal{I}_{\mathrm{MSP}})$ amounts to finding the maximum number of independent components, that is, a set of components of maximum size such that monitoring each component requires a unique detector. We denote the set (respectively, the size) of MSPs by $\mathcal{M}$ (respectively, m^*).

Although $(\mathcal{I}_{\mathrm{MSC}})$ and $(\mathcal{I}_{\mathrm{MSP}})$ are known to be NP-hard problems, modern mixed-integer optimization solvers can be used to optimally solve them for realistic problem instances; see Section 5. Furthermore, their integer programming formulations have linear programming relaxations that are dual of each other (see chapter 13.1 of Vazirani 2001). This implies that $m^* \le n^*$.

MSCs and MSPs represent the network's *coverage* and *spread*, respectively: n^* represents the minimum number of detectors required by P1 to completely monitor the network, and m^* represents the maximum number of attack resources for which P2 can spread her attacks across the network. In fact, solving $\Gamma(b_1,b_2)$ is trivial when $b_1 \ge n^*$, because P1 can monitor all network components by deterministically positioning the detectors on an MSC. Such a detector positioning satisfies Constraints (4) for any target detection rate α. A direct consequence is that the optimal number of detectors in $(\mathcal{P})$, b_1^*, is at most n^*.

On the other hand, a practically relevant (and interesting) case is when P2's number of attack resources is less than the size of MSPs, that is, $b_2 < m^*$. This case captures the situations in which the network is large enough in that P2 can exhaust her ability to spread attacks, thereby making it most challenging for P1 to detect the attacks using her inspection strategy. Furthermore, when $b_2 \ge m^*$, a larger number of attack resources improves P1's ability to detect some of the attacks. Thus, an inspection strategy that ensures the target detection performance for the case $b_2 < m^*$ can also be applied when $b_2 \ge m^*$. Henceforth, our analysis primarily focuses on the case when $b_1 < n^*$ and $b_2 < m^*$ (Figure 1). We discuss the other cases whenever relevant.

3.3. Equilibrium Analysis of Game $\Gamma(b_1,b_2)$

We proceed in three steps. First, we derive bounds on the value of the game Γ based on exact or approximate solutions to the MSC and MSP problems (Proposition 1). Second, we show that every NE satisfies certain structural properties (Theorem 1 and Proposition 3); these properties establish a connection between the zero-sum game Γ and problem $(\mathcal{P})$ (Proposition 2). Finally, we derive properties satisfied by the expected detection rate in equilibrium of Γ (Theorem 2).

3.3.1. Step 1: MSC/MSP-Based Bounds on the Value of the Game Γ. Recall that NE and the value of the game Γ are respectively given by the optimal solutions and optimal value of the linear programs (LP_1) and (LP_2). To derive bounds on the optimal value of (LP_1) and (LP_2), along with mixed strategies that achieve these bounds, we use the following construction.

Lemma 1. *Consider a set of nodes $S \in 2^{\mathcal{V}}$ of size $n \ge b_1$, and a set of components $T \in 2^{\mathcal{E}}$ of size $m \ge b_2$. Then, there exists a strategy profile, denoted $(\sigma^1(S,b_1),\sigma^2(T,b_2)) \in \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2)$, whose node basis and component basis are S and T, respectively, and such that*

$$\forall i \in S,\ \mathbb{P}_{\sigma^1(S,b_1)}(i \text{ is inspected by P1}) = \frac{b_1}{n}, \tag{8}$$

$$\forall e \in T,\ \mathbb{P}_{\sigma^2(T,b_2)}(e \text{ is targeted by P2}) = \frac{b_2}{m}. \tag{9}$$

Figure 1. Three Cases Based on the Magnitude of b_1 and b_2 Relative to n^* and m^*

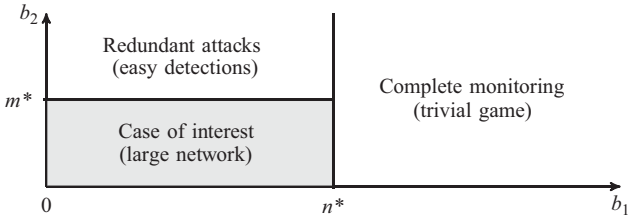

For details on the construction of $(\sigma^1(S,b_1),\ \sigma^2(T,b_2))$, we refer to Lemma EC.3 in the online appendix. The main idea behind the construction of the inspection strategy $\sigma^1(S,b_1)$ is to cycle over size-b_1 subsets of S, such that every node of S is inspected with an identical probability given by (8). A similar idea is used for the construction of the attack strategy $\sigma^2(T,b_2)$. We can use Lemma 1 to derive bounds on the value of the game $\Gamma(b_1,b_2)$ using set covers and set packings:

Proposition 1. *The value of the game $\Gamma(b_1,b_2)$ is upper-bounded by $b_2(1-b_1/|S'|)$ for every minimal set cover $S'\in 2^{\mathcal{V}}$, and is lower-bounded by $\max\{0,b_2(1-b_1/|T'|)\}$ for every set packing $T'\in 2^{\mathcal{E}}$ of size at least b_2. Furthermore, these bounds are achieved by $\sigma^1(S',b_1)$ and $\sigma^2(T',b_2)$, respectively:*

$$\max\left\{0,b_2\left(1-\frac{b_1}{|T'|}\right)\right\}=\min_{S\in\mathcal{A}_1}U(S,\sigma^2(T',b_2))\le U^*(b_1,b_2)$$
$$\le\max_{T\in\mathcal{A}_2}U(\sigma^1(S',b_1),T)=b_2\left(1-\frac{b_1}{|S'|}\right).$$

Recall from Section 3.2 that if P1 had at least n^* detectors (i.e., $b_1\ge n^*$), an equilibrium inspection strategy would be to position n^* detectors on an MSC. Proposition 1 shows that, even for the case when P1 has strictly less than n^* detectors, a set cover is a good candidate for node basis. Analogously, a good candidate for component basis is a set packing. Indeed, if P2 targets components that are spread apart, then it will be difficult for P1 to detect many of these attacks using the available detectors. Thus, by targeting a set packing, P2 can ensure that a single detector can detect at most one attack. We observe that decreasing the size of the minimal set cover and increasing the size of the set packing tighten the bounds on the value of the game Γ. Thus, the best lower (respectively, upper) bound on $U^*(b_1,b_2)$ is $\max\{0,b_2(1-b_1/m^*)\}$ (respectively, $b_2(1-b_1/n^*)$).

3.3.2. Step 2: Equilibrium Properties. The second step consists of deriving structural properties satisfied by every NE of Γ. An important property is that when $b_1<n^*$ and $b_2<m^*$, *any* equilibrium strategy for each player necessarily randomizes over actions that use all available resources.

Theorem 1. *In any equilibrium of $\Gamma(b_1,b_2)$, where $b_1<n^*$ and $b_2<m^*$, P1 must choose an inspection strategy that randomizes over detector positionings of size exactly b_1, and P2 must randomize her attacks over sets of b_2 components.*

$$\forall(\sigma^{1^*},\sigma^{2^*})\in\Sigma,\ \forall S\in\mathrm{supp}(\sigma^{1^*}),\ |S|=b_1, \tag{10}$$

$$\forall(\sigma^{1^*},\sigma^{2^*})\in\Sigma,\ \forall T\in\mathrm{supp}(\sigma^{2^*}),\ |T|=b_2. \tag{11}$$

Then, the NE of Γ can be obtained by solving the following two linear programs:

$$(\overline{\mathrm{LP}_1})\quad\min_{\sigma^1\in\Delta(\overline{\mathcal{A}_1})}\max_{T\in\overline{\mathcal{A}_2}}U(\sigma^1,T)$$
$$(\overline{\mathrm{LP}_2})\quad\max_{\sigma^2\in\Delta(\overline{\mathcal{A}_2})}\min_{S\in\overline{\mathcal{A}_1}}U(S,\sigma^2),$$

where $\overline{\mathcal{A}_1}:=\{S\in 2^{\mathcal{V}}\mid |S|=b_1\}$ and $\overline{\mathcal{A}_2}:=\{T\in 2^{\mathcal{E}}\mid |T|=b_2\}$.

Although it is intuitive that both players *should* use all available resources, this result shows that both players *must necessarily* do so. Property (10) is proven by showing that any additional detector can be used by P1 to strictly improve her payoff, which holds because the network is large (captured by the inequality $b_1<n^*$). Similarly, Property (11) is proven by showing that any additional attack resource can be used by P2 to strictly improve her payoff. This argument combines the fact that P1 cannot monitor all network components with a single detector positioning, and that P2 can spread her attacks across the network (because $b_2<m^*$). In addition, showing (11) involves using the features of the detection function F, Proposition 1, and the properties of $(\mathcal{I}_{\mathrm{MSC}})$ and $(\mathcal{I}_{\mathrm{MSP}})$. Theorem 1 also holds when $b_1<n^*$ and $b_2=m^*$. However, counterexamples can be found when $b_1\ge n^*$ or $b_2>m^*$ (see Section EC.4 in the online appendix).

From (10) and (11), we conclude that the NE of the game Γ can be obtained by solving smaller linear programs. Particularly, the number of variables and constraints can be reduced from $1+\sum_{k=0}^{b_1}\binom{|\mathcal{V}|}{k}$ and $1+\sum_{l=0}^{b_2}\binom{|\mathcal{E}|}{l}$ for (LP_1), to $1+\binom{|\mathcal{V}|}{b_1}$ and $1+\binom{|\mathcal{E}|}{b_2}$ for $(\overline{\mathrm{LP}_1})$; similar reduction applies between (LP_2) and $(\overline{\mathrm{LP}_2})$. Although $(\overline{\mathrm{LP}_1})$ and $(\overline{\mathrm{LP}_2})$ can be used to compute NE for small-sized networks, this approach is still not scalable to large-sized networks.

Importantly, when the network is large enough, that is, when $b_1<n^*$ and $b_2<m^*$, we can build on Proposition 1 and Theorem 1 to establish the following result, which connects $(\mathcal{P})$ and Γ.

Proposition 2. *For an inspection strategy $\sigma^{1^*}\in\Delta(\mathcal{A}_1)$ that maximizes $\min_{\sigma^2\in B_2(\sigma^1,b_2)}r(\sigma^1,\sigma^2)$, any best response to σ^{1^*} randomizes over attack plans of size b_2:*

$$\forall\sigma^{1^*}\in\arg\max_{\sigma^1\in\Delta(\mathcal{A}_1)}\min_{\sigma^2\in B_2(\sigma^1,b_2)}r(\sigma^1,\sigma^2),\ \forall\,\sigma^2\in B_2(\sigma^{1^*},b_2),$$
$$\forall T\in\mathrm{supp}(\sigma^2),\quad |T|=b_2. \tag{12}$$

Then, the optimal value of $(\mathcal{P})$ is the smallest number of detectors for which the expected detection rate in equilibrium of Γ is at least α:

$$b_1^*=\arg\min\{b_1\in\mathbb{N}\mid r(\sigma^*)\ge\alpha,\ \ \forall\sigma^*\in\Sigma(b_1,b_2)\}.$$

Furthermore, an equilibrium inspection strategy of $\Gamma(b_1^*,b_2)$ is an optimal inspection strategy of $(\mathcal{P})$.

Thus, an optimal solution of ($\mathcal{P}$) can be obtained by solving the game Γ and determining the expected detection rate in equilibrium. Next, we focus on the node bases of inspection strategies in equilibrium, which represent the sets of nodes that are inspected with positive probability by P1.

Proposition 3. *In any NE* $(\sigma^{1^*},\sigma^{2^*})\in\Sigma$, *the node basis* $\mathcal{V}_{\sigma^{1^*}}$ *is a set cover. Furthermore, both players must necessarily randomize their actions in equilibrium.*

The proof of this result is based on a best-response argument and uses the fact that from any inspection strategy that leaves one or more components completely unmonitored, we can construct another strategy that strictly improves P1's payoff. This argument is completed by repositioning some detectors and evaluating the resulting change in P1's payoff, which involves exploiting the submodularity of the detection function F, the lower bound on the value of the game Γ (Proposition 1), and the fact that the players must use all resources in equilibrium (Theorem 1). Interestingly, this result may not hold when $b_2\geq m^*$: In that case, P2 may target components that are close to each other, which can result in P1 leaving some components completely unmonitored to focus on the ones for which targeted attacks are easier to detect (see Section EC.4 in the online appendix for an example).

Proposition 3 provides an important insight for planning network inspection operations. To position and operate detectors on the network, the defender typically needs to prepare a subset of locations (nodes). For example, such locations need a secure connection between the detectors' sensing unit and the infrastructure network, as well as reliable power supply for sensing and transmitting the measurements. The number of distinct locations that need to be prepared can be minimized by finding an equilibrium inspection strategy that has a node basis of minimum size. From Proposition 3, we deduce that this number is *at least* n^*.

3.3.3. Step 3: Properties of the Expected Detection Rate in Equilibrium. We now conclude our game-theoretic analysis of Γ by focusing on the equilibrium expected detection rate. In particular, we combine Proposition 1 and Theorem 1 to obtain the following parametric bounds:

$$\forall\sigma^*\in\Sigma(b_1,b_2),\quad \frac{b_1}{n^*}\leq r(\sigma^*)\leq \min\left\{\frac{b_1}{m^*},1\right\}. \tag{13}$$

These bounds are nonincreasing with respect to n^* and m^*. Indeed, as the network size becomes larger, both n^* and m^* increase because each monitoring set covers a smaller fraction of the network. Thus, it is more difficult for P1 to detect attacks (with the same number of detectors) in larger-sized networks, reducing her detection performance.

Next, (13) and Proposition 1 imply that given an MSC $S^{min}\in\mathcal{S}$, the expected detection rate by positioning b_1 detectors according to $\sigma^1(S^{min},b_1)$ provides the following detection guarantee:

$$\min_{\sigma^2\in\Delta(\mathcal{A}_2)} r(\sigma^1(S^{min},b_1),\sigma^2)=\frac{b_1}{n^*}\geq\frac{\max\{b_1,m^*\}}{n^*}r(\sigma^*),\quad \forall\sigma^*\in\Sigma(b_1,b_2). \tag{14}$$

Property (14) shows that by using $\sigma^1(S^{min},b_1)$ as inspection strategy, P1 is guaranteed an expected detection rate of at least b_1/n^*, regardless of the attack strategy chosen by P2. Thus, this guarantee applies even if the disruptions are caused by random failures or by an attacker who does not always select a best response strategy. In fact, the relative difference between the expected detection rate in equilibrium and when P1 chooses $\sigma^1(S^{min},b_1)$ is upper bounded by $1-\max\{b_1,m^*\}/n^*$; we refer to this bound as the *relative loss of performance*.

We note that when n^* and m^* become closer to each other (or equivalently, as the duality gap between ($\mathcal{I}_{\mathrm{MSC}}$) and ($\mathcal{I}_{\mathrm{MSP}}$) decreases), the gaps between the upper and lower bounds in Proposition 1 and (13) also become narrower. When $n^*=m^*$, the results in Proposition 1 and (13) can be tightened as follows: If $n^*=m^*$ (in addition to $b_1<n^*$ and $b_2<m^*$), then $(\sigma^1(S^{min},b_1),\sigma^2(T^{max},b_2))$ is a NE of $\Gamma(b_1,b_2)$, and the value of the game Γ and the equilibrium expected detection rate are given by:

$$\forall(\sigma^{1^*},\sigma^{2^*})\in\Sigma(b_1,b_2),\quad U(\sigma^{1^*},\sigma^{2^*})=b_2\left(1-\frac{b_1}{n^*}\right), \tag{15}$$

$$\forall\sigma^*\in\Sigma(b_1,b_2),\quad r(\sigma^*)=\frac{b_1}{n^*}. \tag{16}$$

In fact, this result generalizes the equilibrium characterization of prior results on a class of security games. Indeed, our MSC/MSP-based characterization of NE applies to any detection model for which $n^*=m^*$ holds. In Table 2, we list some of the classical models reported in the literature that fall in this category, and compare their features with those of the game Γ. The table also compares the combinatorial objects underlying our equilibrium characterization with their settings. Thus, our analysis generalizes the equilibrium analysis of these security games for more complex networks and when both players have multiple resources.

Finally, we observe in (16) that when $n^*=m^*$, the expected detection rate in equilibrium *does not* depend on the attack resources b_2. We are able to generalize this result and show that the expected detection rate in equilibrium satisfies an important and rather surprising property:

Theorem 2. *Given P1's resources* $b_1\in\mathbb{N}$, *the expected detection rate in equilibrium is identical in any game* $\Gamma(b_1,b_2)$, *with* $b_2<m^*$; *we denote it as* $r^*_{b_1}$:

Table 2. Comparison of Security Games

	Detection model (monitoring locations, components, monitoring sets)			Resources (defender, attacker)	Combinatorial objects used for equilibrium characterization	
	$\mathcal{V}$	$\mathcal{E}$	$\mathcal{C}_i, i \in \mathcal{V}$	$b_1 \geq 1, b_2 \geq 1$	MSC	MSP
Karlin and Peres (2016)	Street roads	Safe-houses	Safe-houses located on road i	$b_1 = 1, b_2 = 1$	Minimum line cover	Maximum matching
Alpern et al. (2011)	Walks	(Node, time) tuples	(Node, time) tuples belonging to walk i	$b_1 = 1, b_2 = 1$	Minimum covering set	Maximum independent set
Garnaev (2000)	Cable locations	Channel sections	Sections covered from the cable's location i	$b_1 = 1, b_2 = 1$	Minimum interval cover	Maximum independent infiltration set
Mavronicolas et al. (2008)	Network edges	Network nodes	End nodes of edge i	$b_1 = 1, b_2 \geq 1$	Minimum edge cover	Maximum independent set
Garnaev et al. (1997)	(Node, time) tuples	Walks	Walks containing the (node, time) tuple i	$b_1 \geq 1, b_2 = 1$	Set of covering calendars	Set of wait-and-run walks

$$\forall b_1 \in \mathbb{N}, \exists r^*_{b_1} \in [0,1] \mid \forall b_2 < m^*, \ \forall \sigma^* \in \Sigma(b_1, b_2), \quad r(\sigma^*) = r^*_{b_1}. \tag{17}$$

Furthermore, inspection strategies in equilibrium of the game $\Gamma(b_1, 1)$ *are also inspection strategies in equilibrium of any game* $\Gamma(b_1, b_2)$ *with* $b_2 < m^*$.

Property (17) is the result of both game-theoretic and combinatorial aspects of our problem. The proof starts by upper bounding the attack probabilities of each component in equilibrium. This is done by accounting for P2's ability to spread her attacks in the network, which is evaluated by Proposition 1 and MSPs, and by exploiting the submodularity of the detection function F with respect to the first variable. These bounds enable us to further characterize P2's equilibrium strategies for any game $\Gamma(b_1, b_2)$, with $b_2 < m^*$. Specifically, consider an attack strategy σ^{2^*} in equilibrium of $\Gamma(b_1, 1)$, and let $\rho_{\sigma^{2^*}}(e)$ denote the resulting probability with which each component $e \in \mathcal{E}$ is targeted. Then, for any $b_2 < m^*$, by applying Farkas' lemma, we show the existence of an attack strategy in equilibrium of $\Gamma(b_1, b_2)$ such that the probability that $e \in \mathcal{E}$ is targeted is given by $b_2 \rho_{\sigma^{2^*}}(e)$. From the additivity of the detection function F with respect to the second variable, we deduce that the value of the game Γ is linear with respect to b_2, and can be expressed as follows:

$$\forall b_1 < n^*, \ \forall b_2 < m^*, \quad U^*(b_1, b_2) = (1 - r^*_{b_1}) b_2. \tag{18}$$

Finally, Theorem 1 implies that the equilibrium expected detection rate is independent of b_2. This whole argument holds because the network is large in comparison with P2's resources, that is, $b_2 < m^*$. Although Theorem 2 also holds when $b_2 = m^*$, Section EC.4 in the online appendix illustrates a counterexample when $b_2 > m^*$.

For the special case when $n^* = m^*$, $r^*_{b_1} = b_1/n^*$ (from (16)) and we find again (15) from (18). Other implications of Theorem 2 and Proposition 2 are that the optimal value of ($\mathcal{P}$) does not depend on b_2, and that equilibrium inspection strategies in the game $\Gamma(b^*_1, 1)$ are optimal inspection strategies of ($\mathcal{P}$). Therefore, we can solve the problem ($\mathcal{P}$) by considering that $b_2 = 1$. Thus, the defender *does not* need to know the actual number of attack resources, *so long as* $b_2 < m^*$.

This conclusion provides a significant advantage from a computational viewpoint. Recall from Theorem 1 that equilibrium inspection strategies of $\Gamma(b_1, b_2)$ are the optimal solutions of ($\overline{\text{LP}_1}$). Now, given $b_1 < n^*$, and by considering that $b_2 = 1$, the optimal value of ($\overline{\text{LP}_1}$) is the expected *undetection* rate in equilibrium $1 - r^*_{b_1}$ (see (18)), and its optimal solutions are inspection strategies in equilibrium of *any* game $\Gamma(b_1, b_2)$ with $b_2 < m^*$. Thus, ($\overline{\text{LP}_1}$) can now be reformulated with $\binom{|\mathcal{V}|}{b_1} + 1$ variables and only $|\mathcal{E}| + 1$ constraints and

can be solved using column generation (Dantzig and Wolfe 1960). In fact, the additivity of the detection function F implies that equilibrium attack strategies can also be computed with a second column generation algorithm. In the Appendix, we present a procedure for computing NE of the game $\Gamma(b_1, b_2)$ with $b_1 < n^*$ and $b_2 < m^*$ in the general case $m^* \le n^*$. Next, we leverage Theorem 2 and the MSC/MSP-based bounds on the expected detection rate in equilibrium (Property (13)) to derive a scalable solution approach to ($\mathcal{P}$).

4. Solution of the Network Inspection Problem

In this section, we use the equilibrium properties of the game Γ to solve the problem ($\mathcal{P}$). Our approach provides an approximate solution to ($\mathcal{P}$) based on MSCs and MSPs that can be further improved with a refinement procedure.

4.1. MSC/MSP-Based Solution

To motivate our approach, let us again consider the special case of $n^* = m^*$. From (16), we conclude that the minimum number of detectors that are needed for the expected detection rate to be at least α in equilibrium is $b_1^* = \lceil \alpha n^* \rceil$. Besides, for an MSC S^{min}, Proposition 1 implies that $\sigma^1(S^{min}, b_1^*)$ is an equilibrium inspection strategy of $\Gamma(b_1^*, b_2)$. Thus, when $n^* = m^*$, we know from Proposition 2 that an optimal solution of the network inspection problem ($\mathcal{P}$) is given by $(\lceil \alpha n^* \rceil, \sigma^1(S^{min}, \lceil \alpha n^* \rceil))$.

For the general case $m^* \le n^*$, we make the following observations. First, the lower bound on the equilibrium expected detection rate, given in (13), ensures the target detection rate α is satisfied with $\lceil \alpha n^* \rceil$ detectors. Second, the upper bound in (13) implies that P1 needs at least $\lceil \alpha m^* \rceil$ detectors to meet the target detection performance. Consequently, the optimal value of ($\mathcal{P}$) satisfies $\lceil \alpha m^* \rceil \le b_1^* \le \lceil \alpha n^* \rceil$. Finally, (14) ensures that our inspection strategy constructed over an MSC (according to Lemma 1) satisfies Constraints (4). These observations lead to the following *MSC/MSP-based solution*:

For any MSC $S^{min} \in \mathcal{S}$ and any number of attack resources $b_2 < m^*$, $(\lceil \alpha n^* \rceil, \sigma^1(S^{min}, \lceil \alpha n^* \rceil))$ is an approximate solution of ($\mathcal{P}$), with optimality gap given by $\lceil \alpha n^* \rceil - \lceil \alpha m^* \rceil$.

We now summarize the main advantages of this MSC/MSP-based solution. First, it reduces to a significant extent the size of the optimization problems that are involved in computing a solution. Indeed, although ($\overline{\text{LP}_1}$) can be used to solve ($\mathcal{P}$) with $b_2 = 1$, the number of variables and constraints required is equal to $\binom{|\mathcal{V}|}{b_1} + 1$ and $|\mathcal{E}| + 1$, respectively. On the other hand, the number of variables and constraints of ($\mathcal{I}_{\text{MSC}}$) is only $|\mathcal{V}|$ and $|\mathcal{E}|$, respectively. Similarly, the number of variables and constraints of ($\mathcal{I}_{\text{MSP}}$) is only $|\mathcal{E}|$ and $|\mathcal{V}|$, respectively.

Second, solving a single instance of ($\mathcal{I}_{\text{MSC}}$) and ($\mathcal{I}_{\text{MSP}}$) enables us to derive a solution to problem ($\mathcal{P}$) for any target detection rate α. Furthermore, for any $b_2 < m^*$, the loss in detection performance and the optimality gap associated with our solution can directly be computed from n^* and m^*. In contrast, for a given target detection rate α, computing an optimal solution of ($\mathcal{P}$) using ($\overline{\text{LP}_1}$) requires solving it for each value of b_1.

Third, the MSC-based inspection strategy derived from our approach is desirable from a practical viewpoint. Because $\sigma^1(S^{min}, b_1)$ is a uniform distribution, it can easily be translated into a schedule that determines how the detectors are positioned in the network and mobilized between the locations (nodes) that have been prepared for the purpose of monitoring. Furthermore, the number of distinct locations being monitored by $\sigma^1(S^{min}, b_1)$, which is represented in our model by the node basis size, is only n^*. From Proposition 3, we recall that the node basis size in equilibrium is at least n^*, and this number is optimal when $n^* = m^*$. This suggests that our MSC-based inspection strategy is simple to implement.

Finally, we note that, although the previously mentioned results require computing an MSC and an MSP (both NP-hard problems), modern mixed-integer optimization solvers can be used to optimally solve them (see Section 5). For extremely large-sized problems, these solvers may not be able to solve ($\mathcal{I}_{\text{MSC}}$) and ($\mathcal{I}_{\text{MSP}}$) to optimality. Still, we can extend our results based only on the computation of a set cover and a set packing. Given a set cover S' and a set packing T' obtained from a heuristic or greedy algorithm (Chvatal 1979, Hifi 1997), we can conclude that $(\lceil \alpha |S'| \rceil, \sigma^1(S', \lceil \alpha |S'| \rceil))$ is an approximate solution of ($\mathcal{P}$). The associated optimality gap is given by $\lceil \alpha |S'| \rceil - \lceil \alpha |T'| \rceil$, which decreases as the size of the set cover decreases and the size of the set packing increases.

4.2. Refinement Procedure

Despite the above-listed advantages of our MSC/MSP-based solution, finding an optimal solution to ($\mathcal{P}$) (i.e., an inspection strategy using less number of detectors than $\lceil \alpha n^* \rceil$) might be desirable. Next, we develop a procedure that iteratively refines the MSC/MSP-based solution proposed in Section 4.1 to provide a stronger performance guarantee, until it reaches optimality of ($\mathcal{P}$). This procedure relies on a column generation algorithm to optimally solve ($\overline{\text{LP}_1}$) for $b_2 = 1$, and obtains an inspection strategy

in equilibrium of any game $\Gamma(b_1, b_2)$ with $b_2 < m^*$ (Theorem 2).

Each iteration of the column generation algorithm involves solving a master problem and a subproblem. Essentially, the master problem is a restricted version of ($\overline{\text{LP}_1}$), where only a subset of variables are considered. Once the master problem is solved, the optimal dual variables are used to construct the subproblem, which involves finding the variable in the unrestricted ($\overline{\text{LP}_1}$) with lowest reduced cost. Specifically, given a subset $\mathcal{I} \subseteq \overline{\mathcal{A}_1}$ of indices, the master problem is given by

$$
\begin{aligned}
(M_{\text{CG}}(\mathcal{I})): \quad & \text{minimize} && z \\
& \text{subject to} && z + \sum_{S \in \mathcal{I}} \mathrm{F}(S,e)\sigma_S^1 \geq 1, \quad \forall e \in \mathcal{E} \\
& && \sum_{S \in \mathcal{I}} \sigma_S^1 = 1 \\
& && \sigma_S^1 \geq 0, \qquad\qquad \forall S \in \mathcal{I}.
\end{aligned}
$$

For notational simplicity, for any feasible solution $(\sigma^1, z) \in \mathbb{R}_+^{|\mathcal{I}|} \times \mathbb{R}$ of $(M_{\text{CG}}(\mathcal{I}))$, we also use $(\sigma^1, z) \in \mathbb{R}_+^{|\overline{\mathcal{A}_1}|} \times \mathbb{R}$ to represent the corresponding feasible solution of ($\overline{\text{LP}_1}$). Let $(\sigma^{1^*}, z^*) \in \mathbb{R}_+^{|\mathcal{I}|} \times \mathbb{R}$ (respectively, $(\rho^*, z'^*) \in \mathbb{R}_+^{|\mathcal{E}|} \times \mathbb{R}$) denote the optimal primal (respectively, dual) solution of $(M_{\text{CG}}(\mathcal{I}))$. The reduced cost associated with each $S \in \overline{\mathcal{A}_1}$ is given by $-\sum_{e \in \mathcal{E}} \mathrm{F}(S,e)\rho_e^* - z'^*$. Therefore, the detector positioning with the lowest reduced cost can be obtained by solving a maximum weighted covering set problem, where the component weights are the optimal dual variables obtained from the master problem. The subproblem can be formulated as the following integer program:

$$
\begin{aligned}
(S_{\text{CG}}(\rho^*)): \quad & \text{maximize} && \sum_{e \in \mathcal{E}} \rho_e^* y_e \\
& \text{subject to} && y_e \leq \sum_{\{i \in \mathcal{V} \mid e \in \mathcal{C}_i\}} x_i, \qquad \forall e \in \mathcal{E} \\
& && \sum_{i \in \mathcal{V}} x_i = b_1 \\
& && x_i, y_e \in \{0,1\}, \qquad \forall i \in \mathcal{V},\ \forall e \in \mathcal{E}.
\end{aligned}
$$

If the optimal value of $(S_{\text{CG}}(\rho^*))$ is no more than $-z'^*$, then this proves that $(\sigma^{1^*}, z^*) \in \mathbb{R}_+^{|\overline{\mathcal{A}_1}|} \times \mathbb{R}$ is an optimal solution of ($\overline{\text{LP}_1}$). However, if the optimal value of $(S_{\text{CG}}(\rho^*))$ is more than $-z'^*$, then we add the detector positioning corresponding to the optimal solution of $(S_{\text{CG}}(\rho^*))$ to the set of indices $\mathcal{I}$. The master problem $(M_{\text{CG}}(\mathcal{I}))$ is then solved with the new set of indices $\mathcal{I}$. In fact, this algorithm can be warm-started by considering $\mathcal{I} = \text{supp}(\sigma^1(S^{min}, b_1))$, and repeated until an optimal solution of $(M_{\text{CG}}(\mathcal{I}))$ is found.

Thus, we arrive at the computational procedure shown below to *exactly* solve problem ($\mathcal{P}$).

Algorithm 1 (Optimal Solution of ($\mathcal{P}$))

Input: Detection model $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \{\mathcal{C}_i, i \in \mathcal{V}\})$, target detection rate $\alpha \in [0,1]$, MSC $S^{min} \in \mathcal{S}$ of size n^*, and MSP $T^{max} \in \mathcal{M}$ of size m^*.

Output: Number of detectors b_1^* and inspection strategy σ^{1^*}.

A1: Run a binary search method in the discrete interval $[\![\lceil \alpha m^* \rceil, \lceil \alpha n^* \rceil]\!]$:
A2: Select b_1
A3: $\mathcal{I} \leftarrow \text{supp}(\sigma^1(S^{min}, b_1))$
A4: Solve ($\overline{\text{LP}_1}$) by considering $b_2 = 1$ using column generation:
A5: $(\sigma^{1^*}, z^*), (\rho^*, z'^*) \leftarrow$ optimal primal and dual solutions of $(M_{\text{CG}}(\mathcal{I}))$
A6: $(x^*, y^*) \leftarrow$ optimal solution of $(S_{\text{CG}}(\rho^*))$
A7: **If** $-\sum_{e \in \mathcal{E}} \rho_e^* y_e^* - z'^* < 0$, **then**
A8: $\mathcal{I} \leftarrow \mathcal{I} \cup \text{supp}(x^*)$ and go to (A5)
A9: **else**
A10: Output σ^{1^*} and $r_{b_1}^* = 1 - z^*$
A11: **end if**
A12: Terminate the binary search with $b_1^* = \arg\min \{b_1 \in [\![\lceil \alpha m^* \rceil, \lceil \alpha n^* \rceil]\!] \mid r_{b_1}^* \geq \alpha\}$

After each iteration of the column generation algorithm (A5)–(A11) on ($\overline{\text{LP}_1}$) for a given $b_1 \in [\![\lceil \alpha m^* \rceil, \lceil \alpha n^* \rceil]\!]$, let $\sigma^{1'}$ and $1 - r'$ denote the current inspection strategy and value of the objective function, respectively; note that $r' = \min_{e \in \mathcal{E}} r(\sigma^{1'}, e)$. Then, one can derive performance guarantees for $\sigma^{1'}$ by solving ($\mathcal{I}_{\text{MSP}}$), similar to (14). Indeed, given m^*, an upper bound on the relative loss in detection performance is given by $\ell' = 1 - r' \max\{b_1, m^*\}/b_1$. When the support of the MSC-based inspection strategy $\sigma^1(S^{min}, b_1)$ is used to warm-start the column generation algorithm, the first iteration of the master problem will give $r' = b_1/n^*$, for which we find again the expression of the loss in detection performance in (14). Then, ℓ' decreases as the number of iterations of the column generation algorithm increases. If $r' \geq \alpha$, then $(b_1, \sigma^{1'})$ is a feasible solution of ($\mathcal{P}$), with optimality gap given by $b_1 - \lceil \alpha m^* \rceil$.

When ($\overline{\text{LP}_1}$) is solved to optimality for a given b_1, the optimal dual variables of ($\overline{\text{LP}_1}$) represent the probabilities with which the network components are targeted by P2 in an equilibrium of the game $\Gamma(b_1, 1)$. In the Appendix, we derive a procedure that uses these probabilities to construct an attack strategy in equilibrium of $\Gamma(b_1, b_2)$ for $b_2 < m^*$.

A downside of Algorithm 1 is that it can output a significantly complex inspection strategy; for example, one that randomizes over $|\mathcal{E}|$ detector positionings on a node basis of size $b_1 |\mathcal{E}|$, as opposed to our MSC-based inspection strategy $\sigma^1(S^{min}, b_1)$ that uniformly randomizes over n^* detector positionings on a node basis of size n^*. Thus, scheduling a network inspection

according to this new strategy would likely require a larger level of preparation and operational capability on the part of the defender. Our approach enables the defender to compute and choose an inspection strategy with a tradeoff between detection performance and ease of implementation.

5. Computational Results

In this section, we demonstrate the scalability and performance guarantees of our approach for large-scale networks. We consider a batch of benchmark water distribution networks varying in their size and complexity that are typically used to test network monitoring algorithms. Table 3 lists the characteristics of the 13 networks considered in our study. The data for these networks can be found in Perelman et al. (2008), University of Exeter (2014), and Jolly et al. (2014). The water networks in our study range from medium-sized to very-large-sized networks serving populations from 3,000 to 250,000 consumers (U.S. Environmental Protection Agency 2007). All network simulations were implemented in Matlab, and all optimization problems were solved using the Gurobi solver on a computer with a 2.3-GHz, 8-Core Intel Core i9 processor and 32 GB of RAM.

We consider an application of problem ($\mathcal{P}$), in which pipelines are subject to attack-induced disruptions. To detect these attacks, we consider that the water utility has access to the relevant sensing technology, such as pressure loggers that can easily be mounted at various nodes (e.g., access points such as valves and hydrants), and shifted from one node to another (Allen et al. 2011, Wright et al. 2015, Xing and Sela 2019). For this application, the set of monitoring locations $\mathcal{V}$ is given by the set of network nodes, and the set $\mathcal{E}$ of critical components is the set of pipes. Then, for each possible monitoring location $i \in \mathcal{V}$, we compute the monitoring set $\mathcal{C}_i$ (defined in Section 2.1). In our study, monitoring sets are computed through simulations using a threshold-based detection model, as proposed in Deshpande et al. (2013) and Sela Perelman et al. (2016).

We then apply our solution approach for each network (see Section EC.3.1 in the online appendix for an illustrative example using a small-sized pipeline network from Giustolisi et al. (2008)): We solve ($\mathcal{I}_{\text{MSC}}$) to compute the number of detectors $\lceil \alpha n^* \rceil$ that are sufficient to achieve the target detection rate α, and determine an MSC S^{min} that should be prepared by the water utility for inspection. Then, the utility's schedule of inspections can be generated from the inspection strategy $\sigma^1(S^{min}, \lceil \alpha n^* \rceil)$. Next, we solve ($\mathcal{I}_{\text{MSP}}$), which enables us to evaluate the performance of our solution, i.e., we compute the optimality gap $\lceil \alpha n^* \rceil - \lceil \alpha m^* \rceil$ given in Section 4.1 and the relative loss of performance $1 - \max\{\lceil \alpha n^* \rceil, m^*\}/n^*$ derived from (14). The computational results are summarized in Table 3.

We observe that the sizes of MSCs and MSPs are equal for 6 of the 13 networks. Thus, for these six networks, our MSC/MSP-based solution is optimal for ($\mathcal{P}$). For the remaining seven networks, we note that the relative difference between n^* and m^* is small, which implies that our estimate of the optimal value of ($\mathcal{P}$), $\lceil \alpha n^* \rceil$, is close to the optimal value b_1^*. Additionally, we can see from Table 3 that the loss in detection performance by choosing $\sigma^1(S^{min}, \lceil \alpha n^* \rceil)$ in comparison with the optimal performance is also small (2.7% on average over all networks).

We note that ($\mathcal{I}_{\text{MSC}}$) and ($\mathcal{I}_{\text{MSP}}$) can be solved quickly. For networks with less than 1,500 nodes and components, Gurobi computes an optimal solution in less than 0.5 seconds, which directly enables us to construct an approximate solution to ($\mathcal{P}$). For larger networks, we can obtain n^* and m^* in approximately one minute. Thus, our MSC/MSP-based solution is scalable to large-scale networks.

Next, we run the refinement procedure (Algorithm 1) to improve our solution for the seven networks for which $m^* < n^*$. Table 4 summarizes the computational results for the four networks for which the procedure terminated in a reasonable time.

For instance, we find that for network ky13, an optimal solution of ($\mathcal{P}$) requires one fewer detector than the MSC/MSP-based solution does to satisfy the target detection performance α. Furthermore, the equilibrium expected detection rate $r^*_{b_1^*}$ improves by 3.39% the detection performance of the MSC-based inspection strategy $\sigma^1(S^{min}, b_1^*)$. Finally, we find that the optimal inspection strategy σ^{1^*} has a support of size 47 and randomizes over 205 distinct locations.

Table 4 shows that for four of the remaining seven networks, Algorithm 1 optimally solves ($\mathcal{P}$). First, we observe that the optimal solutions require only one fewer detector than the MSC-based solution. Second, given b_1^* detectors, the improvement between the optimal and MSC-based inspection strategies is between 1.79% and 3.39% and is achieved under 50 minutes. We note that solving ($\mathcal{I}_{\text{MSP}}$) significantly reduces the runtime of the refinement procedure by limiting the binary search to the interval $[\![\lceil \alpha m^* \rceil, \lceil \alpha n^* \rceil]\!]$: ($\overline{\text{LP}_1}$) is solved for only one value of b_1 for networks ky5, ky2, and ky4 and is solved for two values of b_1 for network ky13.

However, for the three larger networks (dover, bswn2, mnsr), the refinement procedure did not terminate after 72 hours of runtime. One of the main reasons is that the restricted master problem ($M_{\text{CG}}(\mathcal{I})$) faces degeneracy issues when the subset of variables $\mathcal{I}$ is small: Even if the subproblem ($S_{\text{CG}}(\rho^*)$) finds a

Table 3. Network Data and Computational Results of the MSC/MSP-Based Solution: $\alpha = 0.75$

Network	Total length (km)	No. of pipes	No. of nodes	Running time (s) ($\mathcal{I}_{\text{MSP}}$)	Running time (s) ($\mathcal{I}_{\text{MSC}}$)	m^*	n^*	$\lceil \alpha n^* \rceil$	Optimality gap	Relative loss of performance
bwsn1	37.56	168	126	0.05	0.11	7	7	6	0%	0%
ky3	91.29	366	269	0.01	0.03	15	15	12	0%	0%
ky5	96.58	496	420	0.02	0.05	18	19	15	1 (7.14%)	5.3%
ky7	137.05	603	481	0.09	0.08	28	28	21	0%	0%
ky6	123.20	644	543	0.08	0.06	24	24	18	0%	0%
ky1	166.60	907	791	0.03	0.08	31	31	24	0%	0%
ky13	153.30	940	778	0.06	0.08	28	30	23	2 (9.52%)	6.7%
ky2	152.25	1,124	811	0.39	0.41	18	19	15	1 (7.14%)	5.3%
ky4	260.24	1,156	959	0.03	0.05	62	64	48	1 (2.13%)	3.1%
ky8	247.34	1,614	1,325	0.14	0.22	45	45	34	0%	0%
dover	779.86	16,000	14,965	4.34	8.36	119	121	91	1 (1.11%)	1.7%
bswn2	1,844.04	14,822	12,523	0.77	4.06	352	361	271	7 (2.65%)	2.5%
mnsr	476.67	25,484	24,681	58.89	68.67	50	52	39	1 (2.63%)	3.8%

variable with negative reduced cost, adding that variable to $\mathcal{I}$ does not change the new optimal solution of ($M_{\text{CG}}(\mathcal{I})$). On the other hand, when the subset of variables $\mathcal{I}$ is large, the runtime of one iteration of the column generation algorithm (A5)–(A11) is large.

In Figure 2, we compare the column generation algorithm applied to ($\overline{\text{LP}_1}$) for network ky4 with and without the MSC-based warm-start. The left plot (respectively, right plot) represents the worst-case detection rate (respetively, node basis size) of the inspection strategy found by column generation, as a function of the algorithm's runtime. We recall that the worst-case detection rate of an inspection strategy is given by one minus its objective value in ($\overline{\text{LP}_1}$). Additional figures are presented in Section EC.3.2 in the online appendix.

Figure 2 shows that as the column generation algorithm (A4) runs, the objective value of the inspection strategy obtained by solving the master problem ($M_{\text{CG}}(\mathcal{I})$) decreases, which results in an increase in the worst-case detection rate. Furthermore, for this network, initiating the column generation algorithm (A4) with the variables corresponding to the detector positionings in the support of the MSC-based strategy reduces the runtime by half. Interestingly, we observe a peak in the size of the node basis, that is, the number of distinct monitored locations, of the solution σ^1 to the restricted master problem ($M_{\text{CG}}(\mathcal{I})$). In the first iterations, P2's best response (given by the dual variables of ($M_{\text{CG}}(\mathcal{I})$)) targets parts of the network that are not monitored by the inspection strategy. Thus, the algorithm first "explores" the network and positions the detectors in a greedy-like manner on locations that are more spread out. As the inspection strategy improves, P2 selects attack strategies that are more evenly spread in the network. This in turn forces the algorithm to consolidate the support of the inspection strategy and position detectors on more strategic locations.

Finally, we note that MSC-based strategies are significantly simpler than the optimal inspection strategies. For instance, for network ky4, the MSC-based strategy randomizes over $n^* = 64$ different locations, whereas the optimal strategy randomizes over 311 different locations. Similarly, for network ky2, the MSC-based strategy has a support of size $n^* = 19$, whereas the support of the optimal strategy is of size 39.

In conclusion, our computational results show a tradeoff between the optimal and MSC-based strategies. Specifically, the optimal strategies only provide a marginal improvement in terms of number of utilized detectors and detection performance. Conversely, implementing the optimal strategies would require a much higher level of effort in preparing the detectors' locations and in scheduling the inspections. Thus, depending on her operational capabilities, the defender can decide to implement a simple MSC-based strategy with good performance guarantees, or a more complex optimal strategy.

Table 4. Computational Results of the Refinement Procedure (Algorithm 1): $\alpha = 0.75$

Network	Running time (s)	b_1^*	$r_{b_1^*}^*$	No. of detectors improvement $\lceil \alpha n^* \rceil - b_1^*$	Detection performance improvement $r_{b_1^*}^* - b_1^*/n^*$	$\lvert \text{supp}(\sigma^{1^*}) \rvert$	$\lvert \mathcal{V}_{\sigma^{1^*}} \rvert$
ky5	22.74	14	0.75	1	0.0132 (1.79%)	27	94
ky13	643.86	22	0.7582	1	0.0248 (3.39%)	47	205
ky2	153.11	14	0.75	1	0.0132 (1.79%)	39	133
ky4	2,901.90	47	0.7510	1	0.0165 (2.26%)	73	311

Figure 2. (Color online) Results of Column Generation Applied to ($\overline{\text{LP}_1}$) for Network ky4 ($b_1 = 47$)

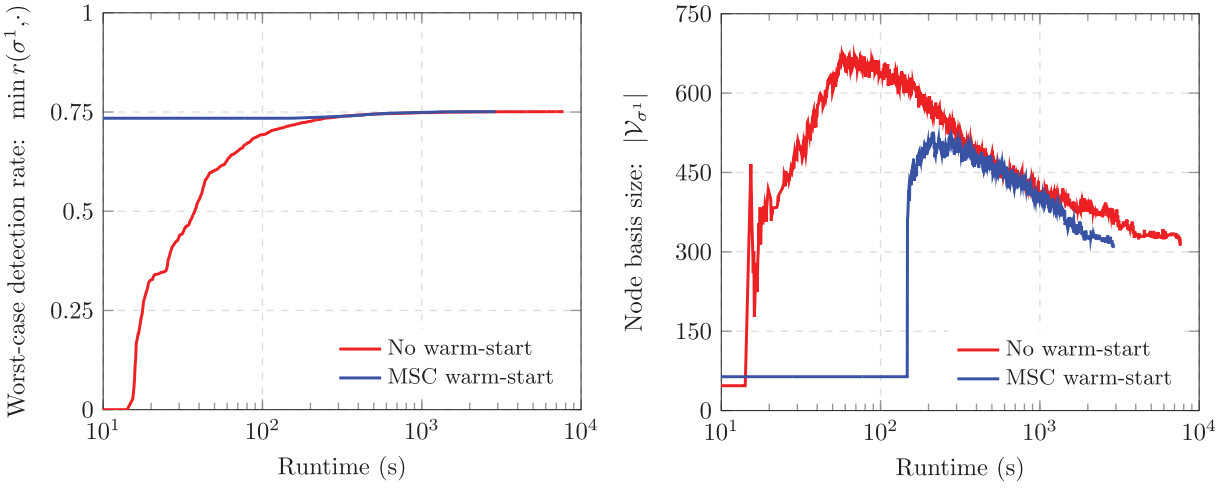

6. Final Remarks

In this article, we studied a generic yet practically relevant formulation of a large-scale bilevel optimization problem for strategic network inspection. In this problem, the defender seeks a randomized inspection strategy that uses a minimum number of detectors while ensuring that the expected detection performance against worst-case attack strategies is above a desirable threshold. We developed a novel approach that analyzes the equilibria of a zero-sum game, which enables us to solve the inspection problem for large-scale networks along with performance guarantees.

Our equilibrium analysis involves (i) deriving useful qualitative properties satisfied by all NE of the zero-sum game; (ii) obtaining bounds on the expected detection rate in equilibrium based on solutions of the MSC and MSP problems; and (iii) showing that, in equilibrium, the expected detection rate and inspection strategies are independent of the attack resources.

Our equilibrium analysis leads to a tractable approach to solve the inspection problem: First, the MSC and MSP problems are solved to obtain an approximate solution that estimates the required number of detectors (with optimality gap) and provides an inspection strategy with guarantees on the expected detection performance. Then, a column generation–based procedure further improves the guarantees of our solution. We demonstrated the scalability and performance of our approach for the allocation of sensing resources in large-scale urban water networks facing security attacks. Our results highlight an important tradeoff between the optimal and MSC/MSP-based solutions in terms of performance guarantees and ease of implementation.

A future research question is to solve the inspection problem under a more refined detection model that accounts for the imperfect detection of attacks (and other types of compromises). Typically, the diagnostic ability of the sensing technology can be represented by a probabilistic detection rate for any given false alarm rate. In fact, as mentioned in Section 2.1, the guarantees provided by our approach can be extended (via simple scaling) to the case when the detection probability is a priori known and homogeneous for all detectors. The general case of heterogeneous detection rates can be addressed by extending our detection model; in particular, by adding a weight to each inspected node to represent the probability of detecting an attack within the node's monitoring set.

Another research question is to extend our solution approach to account for the heterogeneity of the network components in terms of their criticality to the overall network functionality. In principle, this case can be addressed by adding weights to the detection function. However, in many practical situations, the defender can only qualitatively distinguish the criticality of various components (high versus low). In such cases, our approach for strategic network inspection can be applied to each group of components with homogeneous criticality levels, and the inspection strategies for individual groups can be then integrated based on the defender's operational constraints.

Overall, the outcomes of the proposed approach can be used to inform and guide public utilities to design inspection strategies for protecting critical infrastructures against intentional threats. The results indicate that a small number of defense resources, if allocated in a strategic manner, can be sufficient to achieve a high level of protection in large-scale networks, which is especially appealing for budget-constrained utilities. With advances in sensing and detection technologies, randomized inspection strategies, such as the ones proposed in this work, are expected to be essential for reducing risks and for building greater resilience in critical infrastructure systems.

Acknowledgments

The authors thank Özlem Ergun, the associate editor, and the reviewers for useful suggestions that improved the quality of this article and Ali Jadbabaie, Patrick Jaillet, Asuman E. Ozdaglar, Georgia Perakis, and Zuo-Jun Max Shen for feedback.

Appendix. Column Generation and NE of Γ

In Section 4.2, we presented a column generation algorithm to solve the linear program ($\overline{\text{LP}_1}$) for $b_2 = 1$ and obtain an inspection strategy in equilibrium of any game $\Gamma(b_1, b_2)$ with $b_2 < m^*$. Next, we discuss how a second column generation algorithm can be applied to derive attack strategies in equilibrium of any game $\Gamma(b_1, b_2)$ with $b_2 < m^*$.

Consider ($\overline{\text{LP}_1}$) for $b_2 = 1$, and assume that the column generation algorithm (A4) finds an optimal solution. Then, the optimal dual variables $(\rho_e^*)_{e\in\mathcal{E}}$ of $(M_{\text{CG}}(\mathcal{I}))$ represent the probabilities with which each component can be targeted in equilibrium of the game $\Gamma(b_1, 1)$. In the proof of Theorem 2, we show how to reallocate these probabilities to create an attack strategy in equilibrium of $\Gamma(b_1, 1)$ with the additional property that each component is not targeted with probability more than $1/m^*$. Then, given $b_2 < m^*$, Lemma EC.4 in the online appendix and the proof of Theorem 2 show that an attack strategy in equilibrium of $\Gamma(b_1, b_2)$ can be computed by solving the following feasibility problem: Find $\sigma^2 \in \mathbb{R}_+^{|\overline{\mathcal{A}_2}|}$ such that $\sum_{\{T\in\overline{\mathcal{A}_2}\,|\,e\in T\}} \sigma_T^2 = b_2\rho_e^*$ for every $e \in \mathcal{E}$. This can be done by considering the following auxiliary linear problem:

$$(\mathcal{F}(b_2\rho^*)):\ \underset{\sigma^2, s}{\text{minimize}} \ \sum_{e\in\mathcal{E}} s_e \quad \text{subject to} \sum_{\{T\in\overline{\mathcal{A}_2}\,|\,e\in T\}} \sigma_T^2 + s_e = b_2\rho_e^*, \quad \forall e \in \mathcal{E}$$
$$\sigma^2 \geq \mathbf{0}_{|\overline{\mathcal{A}_2}|}, \quad s \geq \mathbf{0}_{|\mathcal{E}|}.$$

Problem $(\mathcal{F}(b_2\rho^*))$ can also be solved using column generation, with $(\sigma^2, s) = (\mathbf{0}_{|\overline{\mathcal{A}_2}|}, b_2\rho^*)$ as initial feasible solution. Given the current restricted master problem generated by the column generation algorithm, let $\beta^* \in \mathbb{R}^{|\mathcal{E}|}$ denote its optimal dual variables. Then, the index $T^* \in \overline{\mathcal{A}_2}$ with lowest reduced cost is given by $T^* \in \arg\max_{T\in\overline{\mathcal{A}_2}} \sum_{e\in T}\beta_e^*$. Therefore, T^* can be efficiently computed by simply finding the b_2 highest values of β_e^*. Lemma EC.4 in the online appendix guarantees that the optimal value of $(\mathcal{F}(b_2\rho^*))$ is zero, and an optimal solution is an equilibrium attack strategy of $\Gamma(b_1, b_2)$.

In conclusion, we obtain the following procedure for computing NE of the game $\Gamma(b_1, b_2)$ in the general case $m^* \leq n^*$.

Algorithm A.1 (NE of $\Gamma(b_1, b_2)$)

Input: Detection model $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \{\mathcal{C}_i, i \in \mathcal{V}\})$, and players' resources $b_1 < n^*$ and $b_2 < m^*$.

Output: Strategy profile $(\sigma^{1^*}, \sigma^{2^*})$.

Solve ($\overline{\text{LP}_1}$) by considering $b_2 = 1$ using column generation:

$(\sigma^{1^*}, z^*), (\rho^*, z'^*) \leftarrow$ optimal primal and dual solutions of ($\overline{\text{LP}_1}$)

Reallocate probabilities in ρ^* so that $\rho_e^* \leq 1/m^*$ for every $e \in \mathcal{E}$

Solve $(\mathcal{F}(b_2\rho^*))$ using column generation:

$(\sigma^{2^*}, \mathbf{0}_{|\overline{\mathcal{A}_2}|}) \leftarrow$ optimal primal solution of $(\mathcal{F}(b_2\rho^*))$

References

Alderson DL, Brown GG, Carlyle WM (2015) Operational models of infrastructure resilience. *Risk Anal.* 35(4):562–586.

Alderson DL, Brown GG, Carlyle WM, Wood RK (2018) Assessing and improving the operational resilience of a large highway infrastructure system to worst-case losses. *Transportation Sci.* 52(4): 1012–1034.

Allen M, Preis A, Iqbal M, Stitangarajan S, Lim HN, Girod L, Whittle AJ (2011) Real time in-network monitoring to improve operational efficiently. *J. Amer. Water Works Assoc.* 103(7):63–75.

Alpern S, Morton A, Papadaki K (2011) Patrolling games. *Oper. Res.* 59(5):1246–1257.

Barrett MP (2018) *Framework for Improving Critical Infrastructure Cybersecurity* (National Institute of Standards and Technology, Gaithersburg, MD).

Baykal-Gürsoy M, Duan Z, Poor HV, Garnaev A (2014) Infrastructure security games. *Eur. J. Oper. Res.* 239(2):469–478.

Berry J, Hart W, Phillips C, Uber J, Watson J (2006) Sensor placement in municipal water networks with temporal integer programming models. *J. Water Resource Planning Management* 132(4):218–224.

Bertsimas D, Nasrabadi E, Orlin JB (2016) On the power of randomization in network interdiction. *Oper. Res. Lett.* 44(1):114–120.

Bier VM, Haphuriwat N (2011) Analytical method to identify the number of containers to inspect at U.S. ports to deter terrorist attacks. *Ann. Oper. Res.* 187(1):137–158.

Bier VM, Haphuriwat N, Menoyo J, Zimmerman R, Culpen AM (2008) Optimal resource allocation for defense of targets based on differing measures of attractiveness. *Risk Anal.* 28(3):763–770.

Brown G, Carlyle M, Salmerón J, Wood K (2006) Defending critical infrastructure. *INFORMS. J. Appl. Anal.* 36(6):530–544.

Chakrabarti S, Kyriakides E, Eliades D (2009) Placement of synchronized measurements for power system observability. *IEEE Trans. Power Delivery* 24(1):12–19.

Chong CY, Kumar SP (2003) Sensor networks: evolution, opportunities, and challenges. *Proc. IEEE* 91(8):1247–1256.

Chvatal V (1979) A greedy heuristic for the set-covering problem. *Math. Oper. Res.* 4(3):233–235.

Cormican KJ, Morton DP, Wood RK (1998) Stochastic network interdiction. *Oper. Res.* 46(2):184–197.

Dahan M, Perelman LS, Amin S (2016) Network sensing for security against link disruption attacks. *Proc. 54th Allerton Conf. Comm., Control, Comput.* (IEEE, Piscataway, NJ), 808–815.

Dancy JR, Dancy VA (2017) Terrorism and oil & gas pipeline infrastructure: Vulnerability and potential liability for cybersecurity attacks. *Oil and Gas, Nat. Resources Energy J.* 2(6):579–619.

Dantzig GB, Wolfe P (1960) Decomposition principle for linear programs. *Oper. Res.* 8(1):101–111.

Deshpande A, Sarma SE, Youcef-Toumi K, Mekid S (2013) Optimal coverage of an infrastructure network using sensors with distance-decaying sensing quality. *Automatica J. IFAC* 49(11):3351–3358.

Freund Y, Schapire RE (1999) Adaptive game playing using multiplicative weights. *Games and Econom. Behav.* 29(1):79–103.

Fujishige S (2005) *Submodular Functions and Optimization*, vol. 58, 2nd ed. (Elsevier, Amsterdam).

Gal S, Casas J (2014) Succession of hide–seek and pursuit–evasion at heterogeneous locations. *J. R. Soc. Interface* 11(94):20140062.

Garnaev A (2000) *Search Games and Other Applications of Game Theory.* Lecture Notes in Economics and Mathematical Systems (Springer, Berlin).

Garnaev A, Garnaeva G, Goutal P (1997) On the infiltration game. *Internat. J. Game Theory* 26(2):215–221.

Giustolisi O, Savic D, Kapelan Z (2008) Pressure-driven demand and leakage simulation for water distribution networks. *J. Hydraulic Engrg.* 134(5):626–635.

Goyal S, Vigier A (2014) Attack, defense, and contagion in networks. *Rev. Econom. Stud.* 81(4):1518–1542.

Hansen P, Jaumard B, Savard G (1992) New branch-and-bound rules for linear bilevel programming. *SIAM J. Sci. Statist. Comput.* 13(5):1194–1217.

Hassanzadeh A, Rasekh A, Galelli S, Aghashahi M, Taormina R, Ostfeld A, Banks MK (2020) A review of cybersecurity incidents in the water sector. *J. Environment Engrg.* 146(5):03120003.

Hellerstein L, Lidbetter T, Pirutinsky D (2019) Solving zero-sum games using best-response oracles with applications to search games. *Oper. Res.* 67(3):731–743.

Hifi M (1997) A genetic algorithm-based heuristic for solving the weighted maximum independent set and some equivalent problems. *J. Oper. Res. Soc.* 48(6):612–622.

Hochbaum DS, Fishbain B (2011) Nuclear threat detection with mobile distributed sensor networks. *Ann. Oper. Res.* 187(1):45–63.

Jolly MD, Lothes AD, Bryson S, Ormsbee L (2014) Research database of water distribution system models. *J. Water Resource Planning Management* 140(4):410–416.

Karlin A, Peres Y (2016) *Game Theory, Alive* (AMS, Providence, RI).

Krause A, Singh A, Guestrin C (2008a) Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Machine Learning Res.* 9:235–284.

Krause A, McMahan B, Guestrin C, Gupta A (2008b) Robust submodular observation selection. *J. Machine Learning Res.* 9:2761–2801.

Lipton RJ, Markakis E, Mehta A (2003) Playing large games using simple strategies. *Proc. 4th ACM Conf. Elec. Commerce* (ACM, New York), 36–41.

Mavronicolas M, Papadopoulou V, Philippou A, Spirakis P (2008) A network game with attackers and a defender. *Algorithmica* 51(3):315–341.

Monroe J, Ramsey E, Berglund E (2018) Allocating countermeasures to defend water distribution systems against terrorist attack. *Reliability Engrg. System Safety* 179:37–51.

Naureen MS, Collins R, Vamburkar M (2018) Cyberattack pings data systems of at least four gas networks. *Bloomberg* (April 3), https://www.bloomberg.com/news/articles/2018-04-03/day-after-cyber-attack-a-third-gas-pipeline-data-system-shuts.

Orlin JB, Schulz AS, Udwani R (2018) Robust monotone submodular function maximization. *Math. Programming* 172(1):505–537.

Ostfeld A, Salomons E (2004) Optimal layout of early warning detection stations for water distribution systems security. *J. Water Resource Planning Management* 130(5):377–385.

Owolabi T (2016) Nigerian militant group claims attack on oil pipeline in Niger Delta. *Reuters* (September 29), https://www.reuters.com/article/us-nigeria-oil-idUSKCN11Z0XE.

Perelman L, Maslia ML, Ostfeld A, Sautner JB (2008) Using aggregation/skeletonization network models for water quality simulations in epidemiologic studies. *J. Amer. Water Works Assoc.* 100(6):122–133.

PG&E (2010) Pipeline accident report: Pacific gas and electric company natural gas transmission pipeline rupture and fire. Technical report, National Transportation Safety Board, Washington, DC.

Phillips NG, Ackley R, Crosson ER, Down A, Hutyra LR, Brondfield M, Karr JD, et al. (2013) Mapping urban pipeline leaks: Methane leaks across Boston. *Environmental Pollution* 173:1–4.

Pita J, Jain M, Marecki J, Ordóñez F, Portway C, Tambe M, Western C, et al. (2008) Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. *Proc. 7th Internat. Joint Conf. on Autonomous Agents and Multiagent Systems: Industrial Track* (IFAAMAS, Richland, SC), 125–132.

Powell R (2007) Allocating defensive resources with private information about vulnerability. *Amer. Political Sci. Rev.* 101(4):799–809.

Sela L, Amin S (2018) Robust sensor placement for pipeline monitoring: Mixed integer and greedy optimization. *Adv. Engrg. Inform.* 36:55–63.

Sela Perelman L, Abbas W, Koutsoukos X, Amin S (2016) Sensor placement for fault location identification in water networks: A minimum test cover approach. *Automatica J. IFAC* 72:166–176.

Smith JC, Lim C (2008) *Optimality Game Algorithms for Network Interdiction and Fortification Games* (Springer, New York).

Srirangarajan S, Allen M, Preis A, Iqbal M, Lim H, Whittle A (2013) Wavelet-based burst event detection and localization in water distribution systems. *J. Signal Processing Systems* 72(1):1–16.

Tiemann M (2017) Safe drinking water act (SDWA): A summary of the act and its major requirements. Technical report, Congressional Research Service, Washington, DC.

Tzoumas V, Gatsis K, Jadbabaie A, Pappas GJ (2017) Resilient monotone submodular function maximization. *Proc. IEEE 56th Annual Conf. on Decision and Control* (IEEE, Piscataway, NJ), 1362–1367.

University of Exeter (2014) Centre for Water Systems. Accessed October 24, 2014, http://emps.exeter.ac.uk/engineering/research/cws/resources/benchmarks/design-resiliance-pareto-fronts/data-files/.

U.S. Environmental Protection Agency (2007) Factoids: Drinking water and ground water statistics for 2007. Technical report, Office of Water, US EPA, Washington, DC.

Vazirani VV (2001) *Approximation Algorithms* (Springer, Berlin).

Von Neumann J (1953) A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions Theory Games* 2:5–12.

Washburn A, Wood K (1995) Two-person zero-sum games for network interdiction. *Oper. Res.* 43(2):243–251.

Wright R, Abraham E, Parpas P, Stoianov I (2015) Control of water distribution networks with dynamic DMA topology using strictly feasible sequential convex programming. *Water Resources Res.* 51(12):9925–9941.

Xing L, Sela L (2019) Unsteady pressure patterns discovery from high-frequency sensing in water distribution systems. *Water Res.* 158:291–300.

Yuhas A (2016) Pipeline erupts in fiery explosion in Mexico, killing many. *New York Times* (June 18), https://www.nytimes.com/2019/01/18/world/americas/mexico-gas-pipeline-explosion.html.

Zhuang J, Bier VM (2007) Balancing terrorism and natural disasters—Defensive strategy with endogenous attacker effort. *Oper. Res.* 55(5):976–991.

Zhuang J, Bier VM, Alagoz O (2010) Modeling secrecy and deception in a multiple-period attacker–defender signaling game. *Eur. J. Oper. Res.* 203(2):409–418.

Mathieu Dahan is an assistant professor in the School of Industrial and Systems Engineering at Georgia Tech. His research interests are in combinatorial optimization, game theory, and predictive analytics, with applications to service and healthcare operations, humanitarian systems, and logistics and supply chain management. He is a recipient of the MIT Robert Thurber fellowship and the MIT Robert Guenassia award.

Lina Sela is an assistant professor in environmental water resources engineering in the Civil, Architectural and Environment Engineering department at the University of Texas at Austin. Her research focuses on improving the efficiency of water distribution systems facing challenges related to finite water sources, aging infrastructure, and population growth. She is a recipient of the NSF CAREER award.

Saurabh Amin is an associate professor in the Department of Civil and Environmental Engineering at the Massachusetts Institute of Technology. His research focuses on the design of resilient monitoring and control algorithms for infrastructure systems in the face of disruptions, both random and adversarial. He is a recipient of the NSF CAREER award, Google faculty research award, Robert N. Noyce professorship, and Ole Madsen mentoring award.