

MIT Open Access Articles

Network Inspection for Detecting Strategic Attacks

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Dahan, Mathieu, Sela, Lina and Amin, Saurabh. 2022. "Network Inspection for Detecting Strategic Attacks." *Operations Research*, 70 (2).

As Published: 10.1287/OPRE.2021.2180

Publisher: Institute for Operations Research and the Management Sciences (INFORMS)

Persistent URL: <https://hdl.handle.net/1721.1/148394>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Network Inspection for Detecting Strategic Attacks

Mathieu Dahan

School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332,
mathieu.dahan@isye.gatech.edu

Lina Sela

Department of Civil, Architectural and Environmental Engineering, University of Texas, Austin, TX 78712,
linasela@utexas.edu

Saurabh Amin

Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139,
amins@mit.edu

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}). Firstly, we construct an approximate solution by utilizing 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. Secondly, 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.

Subject classifications: Games/group decisions: noncooperative; Search and surveillance; Programming: integer: applications.

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 employ 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 due to 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 as well as 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 utilized 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, while the attacker seeks to avoid detection (i.e., maximize 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.

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 and water networks	gas and 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

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, while the defender aims to minimize the number of detectors required to ensure that the expected detection rate under worst-case attacks is above a pre-specified threshold.

Our approach to solve the problem (\mathcal{P}) involves analyzing the equilibrium properties of a zero-sum game, denoted Γ , where the defender (resp. attacker) seeks to minimize (resp. 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}) . Firstly, 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 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. Secondly, a consequence of the equilibrium properties is that a column generation-based procedure 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 which 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 to 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 electronic companion of this article.

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. 2008b, 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. (2008a) 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*, i.e., 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), 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), Zhuang et al. (2010) investigate equilibria in security games with an asymmetric information structure. However, the approaches presented in these papers cannot be applied to solve problem (\mathcal{P}).

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).

2. Problem Description

In this section, we introduce a generic formulation of 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 risk of adversarial disruptions that can compromise the operational functionality of its set of components, denoted \mathcal{E} . All components in the set \mathcal{E} are critical in the sense that they are prone to targeted attacks by an attacker (malicious entity). 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 network 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 disruptions such as damage to pipelines, backflow at fire hydrants, or sabotage of valves (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 defender (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 her inspection strategy, i.e., 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: Firstly, the sensing unit collects relevant state measurements from node i . These measurements capture the state of a subset of components $\mathcal{C}_i \in 2^{\mathcal{E}}$, i.e., the detector at node i monitors the components in \mathcal{C}_i . Secondly, 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 . Thirdly, 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 to 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 $\mathcal{A}_1 := \{S \in 2^{\mathcal{V}} \mid |S| \leq b_1\}$. 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* $F : 2^{\mathcal{V}} \times 2^{\mathcal{E}} \rightarrow \mathbb{N}$. For a detector positioning $S \in 2^{\mathcal{V}}$ and a subset of components $T \in 2^{\mathcal{E}}$, the value of $F(S, T)$ is the number of components of T that are monitored by at least one detector positioned in S , i.e.:

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

Note that under our detection model, if the components of T face attack-induced disruptions, the number of attacks detected by the detector positioning S is $F(S, T)$. The detection function satisfies two natural properties:

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

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

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 $S \in 2^{\mathcal{V}}$, $F(S, \cdot)$ is finitely additive (a direct consequence of (1)):

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

Similar to the defender, the attacker is also resource-constrained, in that she can attack a subset of components $T \in 2^{\mathcal{E}}$ of the network 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 $\mathcal{A}_2 := \{T \in 2^{\mathcal{E}} \mid |T| \leq b_2\}$.

In fact, our solution approach and results on its guarantees 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 $\sum_{e \in T} (1 - (1 - \lambda)^{|\{i \in S \mid e \in \mathcal{C}_i\}|})$, 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 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 (resp. attack plans) are realized from a chosen probability distribution on the set \mathcal{A}_1 (resp. \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 $\Delta(\mathcal{A}_1) := \{\sigma^1 \in [0, 1]^{|\mathcal{A}_1|} \mid \sum_{S \in \mathcal{A}_1} \sigma_S^1 = 1\}$ and $\Delta(\mathcal{A}_2) := \{\sigma^2 \in [0, 1]^{|\mathcal{A}_2|} \mid \sum_{T \in \mathcal{A}_2} \sigma_T^2 = 1\}$ denote the mixed strategy sets. Here, σ_S^1 (resp. σ_T^2) represents the probability assigned to the detector positioning S (resp. attack plan T) by the defender's strategy σ^1 (resp. the attacker's strategy σ^2).

For ease of exposition, we denote $F(i, e) := F(\{i\}, \{e\})$ 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 $\sigma^1 \in \Delta(\mathcal{A}_1)$ (resp. $\sigma^2 \in \Delta(\mathcal{A}_2)$) is defined as $\text{supp}(\sigma^1) = \{S \in \mathcal{A}_1 \mid \sigma_S^1 > 0\}$ (resp. $\text{supp}(\sigma^2) = \{T \in \mathcal{A}_2 \mid \sigma_T^2 > 0\}$). 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 non-zero 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.

Inner problem: In our model, the attacker responds to the defender's inspection strategy by choosing an attack strategy, with the objective of maximizing the expected number of attacks that

remain undetected by the defender. Formally, given an inspection strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$, the attacker aims to find an attack strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$ that maximizes the following objective:

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

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 .

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:

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

Thus, the defender aims to find a minimum-resource inspection strategy while ensuring that the expected detection rate is no less than a pre-specified threshold level $\alpha \in [0, 1]$ against worst-case attack plans. This can be written as the following network inspection problem:

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

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 (for example, due to a regulatory imposition). We denote b_1^* the optimal value of (P).

More generally, the problem (P) captures some of the key features of network inspection in strategic settings; see Table 1 for a comparison of various applications. Firstly, the detection model \mathcal{G} is generic in that it represents the detection capability of the defender, without making further modeling assumptions about how the monitoring sets \mathcal{C}_i ($i \in \mathcal{V}$) depend on specific aspects such as the sensing technology employed by detectors, the different means that the attacker may use in targeting a component, and the network's topological structure. Secondly, 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 $\mathbf{P2}$'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, $\mathbf{P1}$ (resp. $\mathbf{P2}$) selects an inspection strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$ (resp. an attack strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$) and seeks to minimize (resp. maximize) the expected number of undetected attacks (see (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:

$$\forall (\sigma^1, \sigma^2) \in \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2), \quad U(\sigma^{1*}, \sigma^2) \leq U(\sigma^{1*}, \sigma^{2*}) \leq U(\sigma^1, \sigma^{2*}). \quad (5)$$

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.

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

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

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, due to 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.

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 utilizes 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 , i.e., $F(S, e) = 1$, 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, i.e., 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:

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

Solving $(\mathcal{I}_{\text{MSC}})$ amounts to determining the minimum number of detectors and their positioning to monitor all network components. We denote the set (resp. the size) of MSCs by \mathcal{S} (resp. n^*). Since we assumed that each component can be monitored from at least one node in the network (Section 2.1), $(\mathcal{I}_{\text{MSC}})$ 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 , i.e., $F(i, T) \leq 1$, 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:

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

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

Although, $(\mathcal{I}_{\text{MSC}})$ and $(\mathcal{I}_{\text{MSP}})$ 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 $m^* \leq n^*$.

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 $b_1 \geq n^*$, 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, i.e., $b_2 < m^*$. 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 $b_2 \geq m^*$, 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 $b_2 \geq m^*$. Henceforth, our analysis primarily focuses on the case when $b_1 < n^*$ and $b_2 < m^*$ (see Figure 1). We discuss the other cases whenever relevant.

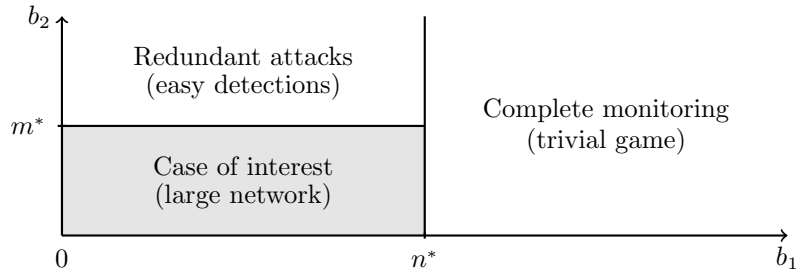


Figure 1 Three cases based on the magnitude of b_1 and b_2 relative to n^* and m^* .

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

We proceed in three steps: Firstly, we derive bounds on the value of the game Γ based on exact or approximate solutions to the MSC and MSP problems (Proposition 1). Secondly, 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).

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 utilize the following construction:

LEMMA 1. Consider a set of nodes $S \in 2^{\mathcal{V}}$ of size $n \geq b_1$, and a set of components $T \in 2^{\mathcal{E}}$ of size $m \geq b_2$. 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:

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

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

For details on the construction of $(\sigma^1(S, b_1), \sigma^2(T, b_2))$, we refer to Lemma EC.3. 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); similarly for 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 $b_2 \left(1 - \frac{b_1}{|S'|}\right)$ for every minimal set cover $S' \in 2^{\mathcal{V}}$, and is lower-bounded by $\max\left\{0, b_2 \left(1 - \frac{b_1}{|T'|}\right)\right\}$ 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:

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

Recall from Section 3.2 that if $\mathbf{P1}$ had at least n^* detectors (i.e., $b_1 \geq n^*$), an equilibrium inspection strategy would be to position n^* detectors on an MSC. Proposition 1 shows that, even for the case when $\mathbf{P1}$ 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 $\mathbf{P2}$ targets components that are spread apart, then it will be difficult for $\mathbf{P1}$ to detect many of these attacks using the available detectors. Thus, by targeting a set packing, $\mathbf{P2}$ 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 (resp. upper) bound on $U^*(b_1, b_2)$ is $\max\left\{0, b_2 \left(1 - \frac{b_1}{m^*}\right)\right\}$ (resp. $b_2 \left(1 - \frac{b_1}{n^*}\right)$).

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 $b_2 < m^*$, $\mathbf{P1}$ must choose an inspection strategy that randomizes over detector positionings of size exactly b_1 , and $\mathbf{P2}$ must randomize her attacks over sets of b_2 components.

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

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

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

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

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 utilized 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 (since $b_2 < m^*$). In addition, showing (11) involves using the features of the detection function F , Proposition 1, and the properties of $(\mathcal{I}_{\text{MSC}})$ and $(\mathcal{I}_{\text{MSF}})$. Theorem 1 also holds when $b_1 < n^*$ and $b_2 = m^*$. However, counterexamples can be found when $b_1 \geq n^*$ or $b_2 > m^*$; see Section EC.4.

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{LP}_1) ; similar reduction applies between (LP_2) and (\overline{LP}_2) . Although (\overline{LP}_1) and (\overline{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”, i.e., 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 :*

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

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 .

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

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 for an example.

Proposition 3 provides an important insight for planning a network inspection operation. 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^* .

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 on the expected detection rate in equilibrium:

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

Firstly, we note that the lower and upper bounds on the expected detection rate in equilibrium are nondecreasing with respect to b_1 . The intuition is that the more detectors **P1** has, the more attacks she will be able to detect. Secondly, these bounds are also 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:

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

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 $\frac{b_1}{n^*}$, 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 $1 - \frac{\max\{b_1, m^*\}}{n^*}$; 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}_{MSC}) and (\mathcal{I}_{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:

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

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

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 when both players have multiple resources.

Table 2 Comparison of security games.

	Detection model (monitoring locations, components, monitoring sets)			Resources def./att.	Combinatorial objects used for equilibrium characterization	
	\mathcal{V}	\mathcal{E}	$\mathcal{C}_i, i \in \mathcal{V}$		MSC	MSP
$\Gamma(b_1, b_2)$				$b_1 \geq 1$ $b_2 \geq 1$	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)	walk	node and start time	(node, time) tuples belonging to walk i	$b_1 = 1$ $b_2 = 1$	minimum covering set	maximum independent set
Garnaev (2000)	cable location	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 edge	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 and time step	walk	walks containing the (node, time) tuple i	$b_1 \geq 1$ $b_2 = 1$	set of covering calendars	set of wait-and-run walks

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 $\mathbf{P1}$'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 $r_{b_1}^*$.*

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

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 $\mathbf{P2}$'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 $\mathbf{P2}$'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:

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

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 to $\mathbf{P2}$'s resources, i.e., $b_2 < m^*$. While Theorem 2 also holds when $b_2 = m^*$, Section EC.4 illustrates a counterexample when $b_2 > m^*$.

Note that for the special case when $n^* = m^*$, $r_{b_1}^* = \frac{b_1}{n^*}$ (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 *undetected* 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 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 $m^* \leq n^*$. 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 utilize 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 $m^* \leq n^*$, we make the following observations: Firstly, 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. Secondly, 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 $\lceil \alpha m^* \rceil \leq b_1^* \leq \lceil \alpha n^* \rceil$. 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. Firstly, 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.

Secondly, 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 .

Thirdly, the MSC-based inspection strategy derived from our approach is desirable from a practical viewpoint. Since $\sigma^1(S^{\min}, b_1)$ 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 while the above-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 utilizing 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:

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

Let $(\sigma^{1*}, z^*) \in \mathbb{R}_+^{|\overline{\mathcal{A}}_1|} \times \mathbb{R}$ (resp. $(\rho^*, z^*) \in \mathbb{R}_+^{|\mathcal{E}|} \times \mathbb{R}$) denote the optimal primal (resp. dual) solution of $(M_{\text{CG}}(\mathcal{I}))$. The reduced cost associated with each $S \in \overline{\mathcal{A}}_1$ is given by $-\sum_{e \in \mathcal{E}} F(S, e) \rho_e^* - z^*$. 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:

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

If the optimal value of $(S_{\text{CG}}(\rho^*))$ is no more than $-z^*$, then this proves that the optimal primal solution of $(M_{\text{CG}}(\mathcal{I}))$, (σ^{1*}, z^*) , is also 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 finally arrive at the following computational procedure to *exactly* solve problem (\mathcal{P}) :

Refinement Procedure: 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 $(M_{CG}(\mathcal{I}))$
- A6: $(x^*, y^*) \leftarrow$ optimal solution of $(S_{CG}(\rho^*))$
- 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 $1 - r'$ respectively denote the current inspection strategy and value of the objective function; 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}_{MSP}) , similarly to (14). Indeed, given m^* , an upper bound on the relative loss in detection performance is given by $\ell' = 1 - \frac{\max\{b_1, m^*\}}{b_1} r'$. 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 $r' = \frac{b_1}{n^*}$, 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. Note that 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 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 this refinement procedure 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 \(2014\)](#), [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 ([US 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. \(2013\)](#), [Sela Perelman et al. \(2016\)](#).

We then apply our solution approach for each network (see Section EC.3.1 for an illustrative example): We solve (\mathcal{I}_{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 $\sigma^1(S^{\text{min}}, \lceil \alpha n^* \rceil)$. Next, we solve (\mathcal{I}_{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 $1 - \frac{\max\{\lceil \alpha n^* \rceil, m^*\}}{n^*}$ derived from (14). The computational results are summarized in Table 3.

We note that the sizes of MSCs and MSPs are equal for 6 out of the 13 networks. Thus, for these 6 networks, our MSC/MSP-based solution is optimal for (\mathcal{P}). For the remaining 7 networks, we note that the relative difference between n^* and m^* is small, which implies that our estimate

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]			m* n* $\lceil \alpha n^* \rceil$	Optimality gap	Relative loss of performance
				(\mathcal{I}_{MSP})	(\mathcal{I}_{MSC})				
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	1124	811	0.39	0.41	18 19 15	1 (7.14%)	5.3%	
ky4	260.24	1156	959	0.03	0.05	62 64 48	1 (2.13%)	3.1%	
ky8	247.34	1614	1325	0.14	0.22	45 45 34	0%	0%	
dover	779.86	16000	14965	4.34	8.36	119 121 91	1 (1.11%)	1.7%	
bswn2	1,844.04	14822	12523	0.77	4.06	352 361 271	7 (2.65%)	2.5%	
mnsr	476.67	25484	24681	58.89	68.67	50 52 39	1 (2.63%)	3.8%	

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 to the optimal performance is also small (2.7% on average over all networks).

Furthermore, the time to solve (\mathcal{I}_{MSC}) and (\mathcal{I}_{MSP}) is fairly small. For networks with less than 1500 nodes and components, Gurobi computes an optimal solution in less than half a second, which directly enables us to construct an approximate solution to (\mathcal{P}). For larger networks, we can obtain n^* and m^* in about a minute. Thus, our MSC/MSP-based solution is scalable to large-scale networks.

Next, we run the refinement procedure presented in Section 4.2 to improve our solution for the 7 networks for which $m^* < n^*$. Table 4 summarizes the computational results for the 4 networks for which the procedure terminated in a reasonable time.

Table 4 Computational results of the refinement procedure, $\alpha = 0.75$.

Network	Running time [s]	b_1^*	$r_{b_1^*}^*$	Improvement		$ \text{supp}(\sigma^{1^*}) $	$ \mathcal{V}_{\sigma^{1^*}} $
				# detectors	detection performance		
				$\lceil \alpha n^* \rceil - b_1^*$	$r_{b_1^*}^* - b_1^*/n^*$		
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	2901.90	47	0.7510	1	0.0165 (2.26%)	73	311

For instance, we find that for network ky13, an optimal solution of (\mathcal{P}) requires 1 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 4 out of the remaining 7 networks, the refinement procedure optimally solves (\mathcal{P}) . First, we observe that the optimal solutions require only one fewer detector than the MSC-based solution. Secondly, given b_1^* detectors, the improvement between the optimal and MSC-based inspection strategies is between 1.79% and 3.79%, and is achieved under 50 minutes. We note that solving (\mathcal{I}_{MSP}) significantly reduces the runtime of the refinement procedure by limiting the binary search to the interval $[[\alpha m^*], [\alpha n^*]]$: (\overline{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 3 larger networks (dover, bsw2, mnsr), the refinement procedure did not terminate after 72 hours of runtime. One of the main reasons is that the restricted master problem $(M_{CG}(\mathcal{I}))$ faces degeneracy issues when the subset of variables \mathcal{I} is small: Even if the subproblem $(S_{CG}(\rho^*))$ finds a variable with negative reduced cost, adding that variable to \mathcal{I} does not change the new optimal solution of $(M_{CG}(\mathcal{I}))$. 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{LP}_1) for the ky4 network with and without the MSC-based warm-start. Additional figures are presented in Section EC.3.2.




Figure 2 Results of column generation applied to (\overline{LP}_1) for the ky4 network ($b_1 = 47$).

Figure 2 shows that 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, i.e.,

the number of distinct monitored locations, of the solution σ^1 to the restricted master problem ($M_{CG}(\mathcal{I})$). In the first iterations, $\mathbf{P2}$'s best response (given by the dual variables of ($M_{CG}(\mathcal{I})$)) 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, $\mathbf{P2}$ 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 the ky4 network, the MSC-based strategy randomizes over $n^* = 64$ different locations, while the optimal strategy randomizes over 311 different locations. Similarly, for the ky2 network, the MSC-based strategy has a support of size $n^* = 19$, while 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. On the other hand, 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.

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 utilizes minimum number of detectors, while ensuring that the expected detection performance against worst-case attack plans 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 imperfect detection of attacks (and other types of compromises). Typically, the diagnostic ability of 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 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 infrastructure 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 critical for reducing risks and for building greater resilience in critical infrastructure systems.

Appendix. Column Generation and Nash Equilibria 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 $\frac{1}{m^*}$. Then, given $b_2 < m^*$, Lemma EC.4 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 $\sum_{\{T \in \overline{\mathcal{A}}_2 \mid e \in T\}} \sigma_T^2 = b_2 \rho_e^*$ for every $e \in \mathcal{E}$. This can be done by considering the following auxiliary linear problem:

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

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 guarantees that the optimal value of $(\mathcal{F}(b_2 \rho^*))$ is 0, 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 $m^* \leq n^*$:

Procedure: 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 $(\sigma^{1*}, \sigma^{2*}) \in \Delta(\overline{\mathcal{A}}_1) \times \Delta(\overline{\mathcal{A}}_2)$.

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

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

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

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

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

Acknowledgments

This work was supported by NSF grant CNS 1239054, NSF CAREER award CNS 1453126, the UT Austin New Faculty Start Up Grant, and MIT Schoettler Fellowship. We are grateful to Özlem Ergun, the Associate Editor, and the reviewers who handled the first round of submission, for useful suggestions. We are also grateful to Ali Jadbabaie, Patrick Jaillet, Asuman E. Ozdaglar, Georgia Perakis, and Zuo-Jun Max Shen for their feedback.

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 efficiency. *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 Res. Plan. Man.* 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, Culpén 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, Salmern J, Wood K (2006) Defending critical infrastructure. *INFORMS J. Applied Analytics* 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* 49(11):3351–3358.
- Fujishige S (2005) *Submodular Functions and Optimization*, volume 58 - 2nd ed. (Elsevier, Amsterdam).

- Gal S, Casas J (2014) Succession of hide–seek and pursuit–evasion at heterogeneous locations. *J. Roy. 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 Heidelberg).
- Garnaev A, Garnaeva G, Goutal P (1997) On the infiltration game. *Int. 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. Hydraul. Engrg.* 134(5):626–635.
- Goyal S, Vigier A (2014) Attack, defense, and contagion in networks. *Rev. Econ. Stud.* 81(4):1518–1542.
- Hansen P, Jaumard B, Savard G (1992) New branch-and-bound rules for linear bilevel programming. *SIAM J. Sci. Stat. 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. Environ. Engrg.* 146(5):03120003.
- 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 Res. Plan. Man.* 140(4):410–416.
- Karlin A, Peres Y (2016) *Game Theory, Alive* (AMS, Providence, RI).
- Krause A, McMahan B, Guestrin C, Gupta A (2008a) Robust submodular observation selection. *J. Mach. Learn. Res.* 9:2761–2801.
- Krause A, Singh A, Guestrin C (2008b) Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.* 9:235–284.
- Lipton RJ, Markakis E, Mehta A (2003) Playing large games using simple strategies. *Proc. 4th ACM Conf. Elec. Commerce* (ACM, New York, NY), 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. *Reliab. Engrg. Sys. Safe.* 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, Udvani 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 Res. Plan. Man.* 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, Zhao K, Jackson RB (2013) Mapping urban pipeline leaks: Methane leaks across Boston. *Environ. Pollut.* 173:1–4.
- Pita J, Jain M, Marecki J, Ordóñez F, Portway C, Tambe M, Western C, Paruchuri P, Kraus S (2008) Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles International Airport. *Proc. 7th Internat. Conf. Autonomous Agents Multiagent Sys.* (IFAAMAS, Richland, SC), 125–132.
- Powell R (2007) Allocating defensive resources with private information about vulnerability. *Amer. Polit. 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* 72:166–176.
- Smith JC, Lim C (2008) Algorithms for network interdiction and fortification games. Chinchuluun A, Pardalos PM, Migdalas A, Pitsoulis L, eds., *Pareto Optimality, Game Theory And Equilibria* (Springer, New York, NY), 609–644.
- 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 Process. Sys.* 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. 56th Conf. Dec. 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-resilience-pareto-fronts/data-files/>.
- US 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 Heidelberg).
- Von Neumann J (1953) A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contr. 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, Pappas P, Stoianov I (2015) Control of water distribution networks with dynamic DMA topology using strictly feasible sequential convex programming. *Water Resour. 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.
- Yuhan A (2016) Pipeline erupts in fiery explosion in Mexico, killing many. *The 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 attackerdefender signaling game. *Eur. J. Oper. Res.* 203(2):409–418.

Supplementary Material

EC.1. Preliminary Results

First, we define the following quantities: For a strategy $\sigma^1 \in \Delta(\mathcal{A}_1)$ of **P1**, the *inspection probability* of node $i \in \mathcal{V}$, denoted $\rho_{\sigma^1}(i)$, is the probability with which i is inspected, i.e.:

$$\forall \sigma^1 \in \Delta(\mathcal{A}_1), \forall i \in \mathcal{V}, \quad \rho_{\sigma^1}(i) := \mathbb{E}_{\sigma^1} [\mathbb{1}_{\{i \in S\}}] = \sum_{\{S \in \mathcal{A}_1 \mid i \in S\}} \sigma_S^1. \quad (\text{EC.1})$$

Given a strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$, the *attack probability* of component $e \in \mathcal{E}$, denoted $\rho_{\sigma^2}(e)$, is the probability with which e is targeted by σ^2 , i.e.:

$$\forall \sigma^2 \in \Delta(\mathcal{A}_2), \forall e \in \mathcal{E}, \quad \rho_{\sigma^2}(e) := \mathbb{E}_{\sigma^2} [\mathbb{1}_{\{e \in T\}}] = \sum_{\{T \in \mathcal{A}_2 \mid e \in T\}} \sigma_T^2. \quad (\text{EC.2})$$

LEMMA EC.1. *The detection function defined in (1) satisfies the following properties:*

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

$$\forall T \in 2^{\mathcal{E}}, \forall (S, S') \in (2^{\mathcal{V}})^2, \quad F(S \cup S', T) + F(S \cap S', T) \leq F(S, T) + F(S', T), \quad (\text{EC.3})$$

$$S \subseteq S' \implies F(S, T) \leq F(S', T). \quad (\text{EC.4})$$

(ii) *For any detector positioning $S \in 2^{\mathcal{V}}$, $F(S, \cdot)$ is finitely additive:*

$$\forall S \in 2^{\mathcal{V}}, \forall (T, T') \in (2^{\mathcal{E}})^2 \mid T \cap T' = \emptyset, \quad F(S, T \cup T') = F(S, T) + F(S, T'). \quad (\text{EC.5})$$

Proof of Lemma EC.1.

(i) Consider a subset of components $T \in 2^{\mathcal{E}}$, and a pair of detector positionings $(S, S') \in (2^{\mathcal{V}})^2$.

Then, $\mathcal{C}_{S \cup S'} = \mathcal{C}_S \cup \mathcal{C}_{S'}$ and $\mathcal{C}_{S \cap S'} \subseteq \mathcal{C}_S \cap \mathcal{C}_{S'}$, and we obtain:

$$\begin{aligned} F(S \cup S', T) + F(S \cap S', T) &\stackrel{(1)}{=} |\mathcal{C}_{S \cup S'} \cap T| + |\mathcal{C}_{S \cap S'} \cap T| = |(\mathcal{C}_S \cap T) \cup (\mathcal{C}_{S'} \cap T)| + |\mathcal{C}_{S \cap S'} \cap T| \\ &= |\mathcal{C}_S \cap T| + |\mathcal{C}_{S'} \cap T| - |\mathcal{C}_S \cap \mathcal{C}_{S'} \cap T| + |\mathcal{C}_{S \cap S'} \cap T| \leq |\mathcal{C}_S \cap T| + |\mathcal{C}_{S'} \cap T| \stackrel{(1)}{=} F(S, T) + F(S', T). \end{aligned}$$

Furthermore, if $S \subseteq S'$, then: $F(S, T) \stackrel{(1)}{=} |\mathcal{C}_S \cap T| \leq |\mathcal{C}_{S'} \cap T| \stackrel{(1)}{=} F(S', T)$.

(ii) Consider a detector positioning $S \in 2^{\mathcal{V}}$. Then, for every $(T, T') \in (2^{\mathcal{E}})^2$ such that $T \cap T' = \emptyset$, $F(S, T \cup T') \stackrel{(1)}{=} |\mathcal{C}_S \cap (T \cup T')| = |\mathcal{C}_S \cap T| + |\mathcal{C}_S \cap T'| - |\mathcal{C}_S \cap T \cap T'| \stackrel{(1)}{=} F(S, T) + F(S, T')$. \square

COROLLARY EC.1. *The detection function defined in (1) satisfies the following properties:*

$$\forall (S, T) \in 2^{\mathcal{V}} \times 2^{\mathcal{E}}, \quad F(S, T) \leq \sum_{i \in S} F(i, T), \quad (\text{EC.6})$$

$$\forall (S, T) \in 2^{\mathcal{V}} \times 2^{\mathcal{E}}, \quad F(S, T) = \sum_{e \in T} F(S, e). \quad (\text{EC.7})$$

Proof of Corollary EC.1.

(i) Consider $T \in 2^\mathcal{E}$. Since $F(\cdot, T)$ is a submodular and nonnegative function, then $F(\cdot, T)$ is subadditive, i.e., for all $(S, S') \in (2^\mathcal{V})^2$, $F(S \cup S', T) \leq F(S, T) + F(S', T)$.

Therefore, by induction, we obtain that for all $S \in 2^\mathcal{V}$, $F(S, T) = F(\cup_{i \in S} \{i\}, T) \leq \sum_{i \in S} F(i, T)$.

(ii) Consider $S \in 2^\mathcal{V}$. Since $F(S, \cdot)$ is finitely additive (Lemma EC.1), we obtain by induction that for every $T \in 2^\mathcal{E}$, $F(S, T) = F(S, \cup_{e \in T} \{e\}) \stackrel{\text{(EC.5)}}{=} \sum_{e \in T} F(S, e)$. \square

LEMMA EC.2. *The size of MSPs is no greater than the size of MSCs, i.e., $m^* \leq n^*$.*

Proof of Lemma EC.2. Consider an MSP $T^{max} = \{e_1, \dots, e_{m^*}\} \in \mathcal{M}$ and an MSC $S^{min} = \{i_1, \dots, i_{n^*}\} \in \mathcal{S}$. Then, we have the desired inequality:

$$m^* \stackrel{\text{(6)}}{=} \sum_{l=1}^{m^*} F(S^{min}, e_l) \stackrel{\text{(EC.6)}}{\leq} \sum_{l=1}^{m^*} \sum_{k=1}^{n^*} F(i_k, e_l) = \sum_{k=1}^{n^*} \sum_{l=1}^{m^*} F(i_k, e_l) \stackrel{\text{(EC.7)}}{=} \sum_{k=1}^{n^*} F(i_k, T^{max}) \stackrel{\text{(7)}}{\leq} n^*. \quad \square$$

EC.2. Proofs of Section 3

LEMMA EC.3. *Consider a set of nodes $S = \{i_1, \dots, i_n\} \in 2^\mathcal{V}$ of size $n \geq b_1$, and a set of components $T = \{e_1, \dots, e_m\} \in 2^\mathcal{E}$ of size $m \geq b_2$. We define the following pure actions:*

$$\forall k \in \llbracket 1, n \rrbracket, S^k = \begin{cases} \{i_k, \dots, i_{k+b_1-1}\} & \text{if } k \leq n - b_1 + 1, \\ \{i_k, \dots, i_n, i_1, \dots, i_{k+b_1-n-1}\} & \text{if } k \geq n - b_1 + 2, \end{cases} \quad \text{(EC.8)}$$

$$\forall l \in \llbracket 1, m \rrbracket, T^l = \begin{cases} \{e_l, \dots, e_{l+b_2-1}\} & \text{if } l \leq m - b_2 + 1, \\ \{e_l, \dots, e_m, e_1, \dots, e_{l+b_2-m-1}\} & \text{if } l \geq m - b_2 + 2, \end{cases} \quad \text{(EC.9)}$$

and a strategy profile $(\sigma^1(S, b_1), \sigma^2(T, b_2)) \in \Delta(\mathcal{A}_1) \times \Delta(\mathcal{A}_2)$ supported over $\{S^1, \dots, S^n\}$ and $\{T^1, \dots, T^m\}$, where:

$$\forall k \in \llbracket 1, n \rrbracket, \sigma^1(S, b_1)_{S^k} := \frac{1}{n}, \quad \text{(EC.10)}$$

$$\forall l \in \llbracket 1, m \rrbracket, \sigma^2(T, b_2)_{T^l} := \frac{1}{m}. \quad \text{(EC.11)}$$

Then, the strategy profile $(\sigma^1(S, b_1), \sigma^2(T, b_2))$ satisfies the following properties:

(i) *Each node in S (resp. each component in T) is inspected (resp. targeted) with an identical probability given by:*

$$\forall i \in S, \rho_{\sigma^1(S, b_1)}(i) = \frac{b_1}{n}, \quad \text{(EC.12)}$$

$$\forall e \in T, \rho_{\sigma^2(T, b_2)}(e) = \frac{b_2}{m}. \quad \text{(EC.13)}$$

(ii) *Each node in S (resp. each component in T) belongs to b_1 actions (resp. b_2 actions) in the support of $\sigma^1(S, b_1)$ (resp. $\sigma^2(T, b_2)$):*

$$\forall i \in S, |\{k \in \llbracket 1, n \rrbracket \mid i \in S^k\}| = b_1, \quad \text{(EC.14)}$$

$$\forall e \in T, |\{l \in \llbracket 1, m \rrbracket \mid e \in T^l\}| = b_2. \quad \text{(EC.15)}$$

(iii) The following inequality is satisfied:

$$\forall e \in \mathcal{E}, F(S, e) \leq \frac{1}{b_1} \sum_{k=1}^n F(S^k, e). \quad (\text{EC.16})$$

Proof of Lemma EC.3. We show the result for a set of nodes $S \in 2^{\mathcal{V}}$ of size $n \geq b_1$. First we note that, by construction, each node $i \in S$ belongs to the same number of detector positionings S^k , $k \in \llbracket 1, n \rrbracket$. Thus, (EC.14) follows from the following calculations:

$$nb_1 \stackrel{(\text{EC.8})}{=} \sum_{k=1}^n \sum_{i' \in \mathcal{V}} \mathbb{1}_{\{i' \in S^k\}} = \sum_{i' \in \mathcal{V}} |\{k \in \llbracket 1, n \rrbracket \mid i' \in S^k\}| = n \times |\{k \in \llbracket 1, n \rrbracket \mid i \in S^k\}|, \quad \forall i \in S.$$

Then, we can show (EC.12). For every node $i \in S$, we have:

$$\rho_{\sigma^1(S, b_1)}(i) \stackrel{(\text{EC.1})}{=} \sum_{k=1}^n \sigma^1(S, b_1)_{S^k} \mathbb{1}_{\{i \in S^k\}} \stackrel{(\text{EC.10})}{=} \frac{1}{n} \sum_{k=1}^n \mathbb{1}_{\{i \in S^k\}} = \frac{1}{n} |\{k \in \llbracket 1, n \rrbracket \mid i \in S^k\}| \stackrel{(\text{EC.14})}{=} \frac{b_1}{n}.$$

An analogous proof can be applied to $T \in 2^{\mathcal{E}}$ of size $m \geq b_2$ to show (EC.13) and (EC.15).

Finally, let us show (EC.16). Consider $e \in \mathcal{E}$. If $F(S, e) = 1$, then there exists $i_0 \in S$ such that $F(i_0, e) = 1$. Since there are b_1 detector positionings in $\{S^k, k \in \llbracket 1, n \rrbracket\}$ that contain i_0 , then $\frac{1}{b_1} \sum_{k=1}^n F(S^k, e) \geq 1 = F(S, e)$. If $F(S, e) = 0$, then $\frac{1}{b_1} \sum_{k=1}^n F(S^k, e) = 0 = F(S, e)$. \square

We illustrate this construction with an example:

EXAMPLE EC.1. Consider a set of three nodes $S = \{i_1, i_2, i_3\}$ and suppose that **P1** has two detectors ($b_1 = 2$). First, we define three pure actions $S^1 = \{i_1, i_2\}$, $S^2 = \{i_2, i_3\}$, and $S^3 = \{i_3, i_1\}$; see Figure EC.1. The strategy $\sigma^1(S, b_1)$ is then obtained by assigning uniform probability (i.e., $\frac{1}{3}$) to each pure action. One can check that each node in S is inspected with probability $\frac{2}{3} = \frac{b_1}{|S|}$. \triangle




Figure EC.1 Support of $\sigma^1(S, b_1)$ when S is composed of three nodes and $b_1 = 2$.

Proof of Proposition 1.

(i) Consider a minimal set cover $S' = \{i_1, \dots, i_n\} \in 2^{\mathcal{V}}$ of size n . Necessarily, S' is such that:

$$\forall k \in \llbracket 1, n \rrbracket, \exists e_k \in \mathcal{E} \mid F(i_k, e_k) = 1 \text{ and } F(i_j, e_k) = 0, \quad \forall j \neq k. \quad (\text{EC.17})$$

– First, let us show that $U^*(b_1, b_2) \leq b_2 \left(1 - \frac{b_1}{n}\right)$. Consider $\sigma^1(S', b_1) \in \Delta(\mathcal{A}_1)$ defined in (EC.10). Recall that we are in the case when $b_1 < n^*$, implying that $b_1 < n^* \leq n$. Since S' is a set cover, then for every $e \in \mathcal{E}$, there exists $k_e \in \llbracket 1, n \rrbracket$ such that $F(i_{k_e}, e) = 1$. Furthermore, F is a nonnegative function. Therefore:

$$\forall e \in \mathcal{E}, \sum_{S \in \mathcal{A}_1} \sigma^1(S', b_1)_S F(S, e) \stackrel{\text{(EC.4)}}{\geq} \sum_{\{S \in \mathcal{A}_1 \mid i_{k_e} \in S\}} \sigma^1(S', b_1)_S \stackrel{\text{(EC.1)}}{=} \rho_{\sigma^1(S', b_1)}(i_{k_e}) \stackrel{\text{(EC.12)}}{=} \frac{b_1}{n}. \quad (\text{EC.18})$$

Thus, we obtain:

$$\begin{aligned} \forall T \in \mathcal{A}_2, U(\sigma^1(S', b_1), T) &\stackrel{\text{(2)}}{=} |T| - \sum_{S \in \mathcal{A}_1} \sigma^1(S', b_1)_S F(S, T) \stackrel{\text{(EC.7)}}{=} |T| - \sum_{e \in T} \sum_{S \in \mathcal{A}_1} \sigma^1(S', b_1)_S F(S, e) \\ &\stackrel{\text{(EC.18)}}{\leq} |T| - \sum_{e \in T} \frac{b_1}{n} = \underbrace{\left(1 - \frac{b_1}{n}\right)}_{>0} |T| \stackrel{\leq b_2}{\leq} b_2 \left(1 - \frac{b_1}{n}\right). \end{aligned} \quad (\text{EC.19})$$

Therefore:

$$U^*(b_1, b_2) = \min_{\sigma^1 \in \Delta(\mathcal{A}_1)} \max_{T \in \mathcal{A}_2} U(\sigma^1, T) \leq \max_{T \in \mathcal{A}_2} U(\sigma^1(S', b_1), T) \leq b_2 \left(1 - \frac{b_1}{n}\right).$$

– Now, let us show that the upper bound on $\max_{T \in \mathcal{A}_2} U(\sigma^1(S', b_1), T)$ is tight. Consider $T' = \{e_1, \dots, e_{b_2}\}$ (where the e_k 's are defined in (EC.17)). Then, for every $k \in \llbracket 1, b_2 \rrbracket$:

$$\begin{aligned} \mathbb{E}_{\sigma^1(S', b_1)}[F(S, e_k)] &\stackrel{\text{(EC.17)}}{=} \sum_{\{S \in \mathcal{A}_1 \mid i_k \in S\}} \sigma^1(S', b_1)_S \underbrace{F(S, e_k)}_{=1} + \sum_{\{S \in \mathcal{A}_1 \mid i_k \notin S\}} \underbrace{\sigma^1(S', b_1)_S F(S, e_k)}_{=0} \\ &\stackrel{\text{(EC.1)}}{=} \rho_{\sigma^1(S', b_1)}(i_k), \end{aligned} \quad (\text{EC.20})$$

where we combined the fact that the node basis of $\sigma^1(S', b_1)$ is S' and that i_k is the only node from S' that monitors component e_k (by construction). This implies that:

$$\max_{T \in \mathcal{A}_2} U(\sigma^1(S', b_1), T) \geq U(\sigma^1(S', b_1), T') \stackrel{\text{(2), (EC.7), (EC.20)}}{=} b_2 - \sum_{k=1}^{b_2} \rho_{\sigma^1(S', b_1)}(i_k) \stackrel{\text{(EC.12)}}{=} b_2 \left(1 - \frac{b_1}{n}\right).$$

Therefore:

$$\max_{T \in \mathcal{A}_2} U(\sigma^1(S', b_1), T) = b_2 \left(1 - \frac{b_1}{n}\right). \quad (\text{EC.21})$$

(ii) Consider a set packing $T' = \{e_1, \dots, e_m\} \in 2^{\mathcal{E}}$ of size $m \geq b_2$.

– First, we show that $U^*(b_1, b_2) \geq \max\{0, b_2 \left(1 - \frac{b_1}{m}\right)\}$. Consider $\sigma^2(T', b_2) \in \Delta(\mathcal{A}_2)$ defined in (EC.11). Since T' is a set packing, then for all $i \in \mathcal{V}$, $F(i, T') \leq 1$. This implies that:

$$\begin{aligned} \forall S \in \mathcal{A}_1, U(S, \sigma^2(T', b_2)) &\stackrel{\text{(2)}}{=} \mathbb{E}_{\sigma^2(T', b_2)}[|T| - F(S, T)] \stackrel{\text{(EC.6), (EC.9)}}{\geq} b_2 - \sum_{i \in S} \mathbb{E}_{\sigma^2(T', b_2)}[F(i, T)] \\ &\stackrel{\text{(EC.2), (EC.7)}}{=} b_2 - \sum_{i \in S} \sum_{e \in \mathcal{E}} F(i, e) \rho_{\sigma^2(T', b_2)}(e) \stackrel{\text{(EC.7), (EC.13)}}{=} b_2 - \frac{b_2}{m} \sum_{i \in S} \underbrace{F(i, T')}_{\leq 1} \geq b_2 \left(1 - \frac{b_1}{m}\right). \end{aligned}$$

Furthermore, for all $S \in \mathcal{A}_1$, $U(S, \sigma^2(T', b_2)) \geq 0$. Therefore:

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

– Now, let us show that the lower bound on $\min_{S \in \mathcal{A}_1} U(S, \sigma^2(T', b_2))$ is tight. In Section 2.1, we assumed that each component can be monitored from at least one node. Therefore for all $l \in \llbracket 1, m \rrbracket$, there exists $i_l \in \mathcal{V}$ such that $F(i_l, e_l) = 1$ (note that the i_l 's are distinct since T' is a set packing). Now, consider the detector positioning $S' = \{i_1, \dots, i_{\min\{b_1, m\}}\}$. S' monitors $\{e_1, \dots, e_{\min\{b_1, m\}}\}$, which enables us to show:

$$\begin{aligned} \min_{S \in \mathcal{A}_1} U(S, \sigma^2(T', b_2)) &\leq U(S', \sigma^2(T', b_2)) \stackrel{(2), (\text{EC.7})}{=} \mathbb{E}_{\sigma^2(T', b_2)} [|T|] - \mathbb{E}_{\sigma^2(T', b_2)} \left[\sum_{e \in T} F(S', e) \right] \\ &\stackrel{(\text{EC.2})}{=} b_2 - \sum_{e \in \mathcal{E}} F(S', e) \rho_{\sigma^2(T', b_2)}(e) = b_2 - \sum_{l=1}^{\min\{b_1, m\}} \rho_{\sigma^2(T', b_2)}(e_l) \stackrel{(\text{EC.13})}{=} \max \left\{ 0, b_2 \left(1 - \frac{b_1}{m} \right) \right\}. \end{aligned}$$

Therefore, $\min_{S \in \mathcal{A}_1} U(S, \sigma^2(T', b_2)) = \max \left\{ 0, b_2 \left(1 - \frac{b_1}{m} \right) \right\}$. \square

Proof of Theorem 1.

(i.a) First, let us show by contradiction that **P1** uses all her resources in equilibrium. Suppose that there exist $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$ and $S^0 \in \text{supp}(\sigma^{1*})$ such that $|S^0| < b_1$.

– The first step is to show that **P2**'s strategy σ^{2*} necessarily targets with positive probability at least one component that is not monitored by S^0 . On the contrary, assume that for every $e \in \mathcal{E}$: $\rho_{\sigma^{2*}}(e) > 0 \implies F(S^0, e) = 1$. Then, **P1** can detect all the attacks of σ^{2*} with the detector positioning S^0 . Since S^0 is a best response to σ^{2*} , the value of the game is 0. However, since $b_1 < n^*$, there exists $e' \in \mathcal{E}$ such that $F(S^0, e') = 0$, i.e., e' is not monitored by S^0 . Since $S^0 \in \text{supp}(\sigma^{1*})$, then e' is not monitored with positive probability. Therefore, if **P2** targets e' , she will get a positive payoff, which contradicts the equilibrium condition (5) for $(\sigma^{1*}, \sigma^{2*})$. Thus, there exists $e_0 \in \mathcal{E}$ such that $\rho_{\sigma^{2*}}(e_0) > 0$ and $F(S^0, e_0) = 0$.

– Now, we can show that **P1** can improve her payoff by positioning one more detector. Let us denote $i_0 \in \mathcal{V} \setminus S^0$ that satisfies $F(i_0, e_0) = 1$. Then, by considering the detector positioning $S' = S^0 \cup \{i_0\} \in \mathcal{A}_1$, we obtain that $U(S', \sigma^{2*}) \leq U(S^0, \sigma^{2*}) - \rho_{\sigma^{2*}}(e_0) < U(S^0, \sigma^{2*})$, which violates the equilibrium condition (5) for $(\sigma^{1*}, \sigma^{2*})$. Therefore, for all $S \in \text{supp}(\sigma^{1*})$, $|S| = b_1$.

(i.b) Now, let us show that **P2** uses all her resources in equilibrium. By contradiction, suppose that there exist $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$ and $T^0 \in \text{supp}(\sigma^{2*})$ such that $|T^0| < b_2$.

– The first step is to show that there exists a component e' not in T^0 that is not monitored by every detector positioning in the support of σ^{1*} . Let us assume the contrary, i.e., that for every $S \in \text{supp}(\sigma^{1*})$ and every $e \notin T^0$, $F(S, e) = 1$. First, let us denote $T^1 \subseteq T^0$ the subset of components of T^0 that are unmonitored by at least one detector positioning $S \in \text{supp}(\sigma^{1*})$, i.e., $T^1 := \bigcup_{S \in \text{supp}(\sigma^{1*})} (\mathcal{E} \setminus \mathcal{C}_S) \subseteq T^0$.

For all $S \in \text{supp}(\sigma^{1*})$, let k_S denote the number of components of T^1 that are not monitored by S . Since every component outside of T^1 is monitored by every detector positioning in the support of σ^{1*} , then **P2**'s best response to σ^{1*} is any attack strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$ such that for all $T \in \text{supp}(\sigma^2)$, $T^1 \subseteq T$ (note that $|T^1| \leq b_2$). This implies that:

$$\forall \sigma^2 \in B_2(\sigma^{1*}), U(\sigma^{1*}, \sigma^2) = \mathbb{E}_{\sigma^{1*}}[k_S] =: k^*. \quad (\text{EC.22})$$

Consider $S \in \text{supp}(\sigma^{1*})$. We know that S leaves k_S network components unmonitored, that we denote e_1, \dots, e_{k_S} . For $l \in \llbracket 1, k_S \rrbracket$, let i_l be a node from where a detector can monitor component e_l . Then, $S \cup \{i_1, \dots, i_{k_S}\}$ is a set cover (of size at most $b_1 + k_S$). By definition of n^* , we deduce that $b_1 + k_S \geq n^*$. Therefore, for all $S \in \text{supp}(\sigma^{1*})$, $k_S \geq n^* - b_1$. Since we are in the case when $b_2 < m^*$, and we have $m^* \leq n^*$ (Lemma [EC.2](#)), then we obtain:

$$k^* = \mathbb{E}_{\sigma^{1*}}[k_S] \geq n^* - b_1 \geq \frac{m^*}{n^*}(n^* - b_1) > \frac{b_2}{n^*}(n^* - b_1). \quad (\text{EC.23})$$

Since σ^{1*} is an equilibrium inspection strategy, the value of the game is $U^*(b_1, b_2) \stackrel{(\text{EC.22})}{=} k^*$. However, from Proposition [1](#), we know that $U^*(b_1, b_2) \leq b_2 \left(1 - \frac{b_1}{n^*}\right)$, which contradicts [\(EC.23\)](#).

Thus, there exists $(e', S') \in \mathcal{E} \setminus T^0 \times \text{supp}(\sigma^{1*})$ such that $F(S', e') = 0$.

– Now, we can show that **P2** can increase her payoff by targeting component e' and the components in T^0 . Let $T' = T^0 \cup \{e'\} \in \mathcal{A}_2$. Then, $U(\sigma^{1*}, T') \geq \sigma_{S'}^{1*} + U(\sigma^{1*}, T^0) > U(\sigma^{1*}, T^0)$, which is a contradiction. Therefore, for every $T \in \text{supp}(\sigma^{2*})$, $|T| = b_2$.

(ii) Finally, we show that **P1**'s strategies in equilibrium of Γ are the optimal solutions of $(\overline{\text{LP}}_1)$. First, from (i), we can deduce that the set of optimal solutions of (LP_1) is a subset of $\Delta(\overline{\mathcal{A}}_1)$. Therefore, the equilibrium inspection strategies are the optimal solutions of $\min_{\sigma^1 \in \Delta(\overline{\mathcal{A}}_1)} \max_{T \in \mathcal{A}_2} U(\sigma^1, T)$.

Now, consider an inspection strategy $\sigma^1 \in \Delta(\overline{\mathcal{A}}_1)$. Since $\overline{\mathcal{A}}_2 \subseteq \mathcal{A}_2$, then we trivially have $\max_{T \in \mathcal{A}_2} U(\sigma^1, T) \geq \max_{T \in \overline{\mathcal{A}}_2} U(\sigma^1, T)$. To obtain the reverse inequality, let $T^0 \in \mathcal{A}_2$ be an attack plan that satisfies $T^0 \in \arg \max_{T \in \mathcal{A}_2} U(\sigma^1, T)$. Then, consider $T' \in \overline{\mathcal{A}}_2$ such that $T^0 \subseteq T'$. We can deduce that $\max_{T \in \mathcal{A}_2} U(\sigma^1, T) = U(\sigma^1, T^0) \leq U(\sigma^1, T') \leq \max_{T \in \overline{\mathcal{A}}_2} U(\sigma^1, T)$.

Therefore, for all $\sigma^1 \in \Delta(\overline{\mathcal{A}}_1)$, $\max_{T \in \mathcal{A}_2} U(\sigma^1, T) = \max_{T \in \overline{\mathcal{A}}_2} U(\sigma^1, T)$, which implies that $\min_{\sigma^1 \in \Delta(\overline{\mathcal{A}}_1)} \max_{T \in \mathcal{A}_2} U(\sigma^1, T) = \min_{\sigma^1 \in \Delta(\overline{\mathcal{A}}_1)} \max_{T \in \overline{\mathcal{A}}_2} U(\sigma^1, T)$. Thus, the equilibrium inspection strategies are the optimal solutions of $(\overline{\text{LP}}_1)$. An analogous proof can be applied to show that the attack strategies in equilibrium of Γ are the optimal solutions of $(\overline{\text{LP}}_2)$. \square

Proof of Proposition [2](#).

(i) Let σ^{1*} be an optimal solution of $\max_{\sigma^1 \in \Delta(\mathcal{A}_1)} \min_{\sigma^2 \in B_2(\sigma^1)} r(\sigma^1, \sigma^2)$. We now show that any best response of **P2** to σ^{1*} uses all resources. Instead, suppose that there exist $\sigma^2 \in B_2(\sigma^{1*})$ and

$T^0 \in \text{supp}(\sigma^2)$ such that $|T^0| < b_2$. Analogously to the proof of Theorem 1, showing that there exists a component not in T^0 that is not monitored by every detector positioning in the support of σ^{1*} leads to a contradiction. By definition of σ^{1*} , we have the following lower bound:

$$\min_{\sigma^2 \in B_2(\sigma^{1*})} r(\sigma^{1*}, \sigma^2) \geq \min_{\sigma^2 \in B_2(\sigma^1(S^{\min}, b_1))} r(\sigma^1(S^{\min}, b_1), \sigma^2). \quad (\text{EC.24})$$

Note that a consequence of Proposition 1 and (EC.19) is that for every $\sigma^2 \in B_2(\sigma^1(S^{\min}, b_1))$ and every $T \in \text{supp}(\sigma^2)$, $|T| = b_2$. Therefore, we have:

$$\begin{aligned} & \min_{\sigma^2 \in B_2(\sigma^1(S^{\min}, b_1))} r(\sigma^1(S^{\min}, b_1), \sigma^2) \stackrel{(2),(3)}{=} \min_{\sigma^2 \in B_2(\sigma^1(S^{\min}, b_1))} \mathbb{E}_{(\sigma^1(S^{\min}, b_1), \sigma^2)} \left[1 - \frac{U(S, T)}{|T|} \right] \\ & = 1 - \frac{1}{b_2} \max_{\sigma^2 \in B_2(\sigma^1(S^{\min}, b_1))} U(\sigma^1(S^{\min}, b_1), \sigma^2) \stackrel{(\text{EC.21})}{=} \frac{b_1}{n^*}. \end{aligned} \quad (\text{EC.25})$$

Now, if we assume that for every $S \in \text{supp}(\sigma^{1*})$ and every $e \notin T^0$, $F(S, e) = 1$, we obtain the following upper bound from the calculations in the proof of Theorem 1:

$$\min_{\sigma^2 \in B_2(\sigma^{1*})} r(\sigma^{1*}, \sigma^2) \stackrel{(2),(3)}{\leq} \min_{\sigma^2 \in B_2(\sigma^{1*})} \mathbb{E}_{(\sigma^{1*}, \sigma^2)} \left[1 - \frac{1}{b_2} U(S, T) \right] \stackrel{(\text{EC.22})}{=} 1 - \frac{1}{b_2} k^* \stackrel{(\text{EC.23})}{<} \frac{b_1}{n^*},$$

which contradicts the lower bound obtained from (EC.24) and (EC.25). Thus, there exists $(e', S') \in \mathcal{E} \setminus T^0 \times \text{supp}(\sigma^{1*})$ such that $F(S', e') = 0$, which contradicts T^0 being in the support of a best response to σ^{1*} . Therefore, for every $\sigma^2 \in B_2(\sigma^{1*})$ and every $T \in \text{supp}(\sigma^2)$, $|T| = b_2$.

(ii) Consider $b_1 \in \mathbb{N}$. First, we note that:

$$\exists \sigma^1 \in \Delta(\mathcal{A}_1) \mid r(\sigma^1, \sigma^2) \geq \alpha, \forall \sigma^2 \in B_2(\sigma^1) \iff \max_{\sigma^1 \in \Delta(\mathcal{A}_1)} \min_{\sigma^2 \in B_2(\sigma^1)} r(\sigma^1, \sigma^2) \geq \alpha.$$

Then, from Theorem 1, and by definition of the best response function B_2 , we have:

$$\begin{aligned} \max_{\sigma^1 \in \Delta(\mathcal{A}_1)} \min_{\sigma^2 \in B_2(\sigma^1)} r(\sigma^1, \sigma^2) \geq \alpha & \stackrel{(2),(3),(12)}{\iff} 1 - \frac{1}{b_2} \min_{\sigma^1 \in \Delta(\mathcal{A}_1)} \max_{\sigma^2 \in \Delta(\mathcal{A}_2)} U(\sigma^1, \sigma^2) \geq \alpha \\ & \iff 1 - \frac{1}{b_2} U(\sigma^{1*}, \sigma^{2*}) \geq \alpha, \forall (\sigma^{1*}, \sigma^{2*}) \in \Sigma(b_1, b_2) \\ & \stackrel{(2),(3),(11)}{\iff} r(\sigma^{1*}, \sigma^{2*}) \geq \alpha, \forall (\sigma^{1*}, \sigma^{2*}) \in \Sigma(b_1, b_2). \end{aligned}$$

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

(iii) Let σ^{1*} be an optimal solution of (LP₁). Then, we have:

$$\begin{aligned} & \max_{\sigma^2 \in B_2(\sigma^{1*})} U(\sigma^{1*}, \sigma^2) \leq \max_{\sigma^2 \in B_2(\sigma^1)} U(\sigma^1, \sigma^2), \forall \sigma^1 \in \Delta(\mathcal{A}_1) \\ & \stackrel{(2),(3),(11)}{\iff} \max_{\sigma^2 \in B_2(\sigma^{1*})} (1 - r(\sigma^{1*}, \sigma^2)) b_2 \leq \max_{\sigma^2 \in B_2(\sigma^1)} \mathbb{E}_{(\sigma^1, \sigma^2)} [(1 - r(S, T)) \underbrace{|T|}_{\leq b_2}], \forall \sigma^1 \in \Delta(\mathcal{A}_1) \\ & \implies \min_{\sigma^2 \in B_2(\sigma^{1*})} r(\sigma^{1*}, \sigma^2) \geq \min_{\sigma^2 \in B_2(\sigma^1)} r(\sigma^1, \sigma^2), \forall \sigma^1 \in \Delta(\mathcal{A}_1). \end{aligned}$$

Therefore, $\sigma^{1*} \in \arg \max_{\sigma^1 \in \Delta(\mathcal{A}_1)} \min_{\sigma^2 \in B_2(\sigma^1)} r(\sigma^1, \sigma^2)$. Thus, an equilibrium inspection strategy of $\Gamma(b_1^*, b_2)$ is an optimal inspection strategy of (\mathcal{P}). \square

Proof of Proposition 3.

(i) We show the result by contradiction, that is, suppose that there exists $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$ such that $\mathcal{V}_{\sigma^{1*}}$ is not a set cover. For simplicity, we introduce the following notation:

$$\forall e \in \mathcal{E}, \eta_{\sigma^{1*}}(e) := \mathbb{E}_{\sigma^{1*}}[\mathbf{F}(S, e)], \quad (\text{EC.26})$$

which is the probability with which component e is monitored by σ^{1*} . Let us sort the components in nondecreasing order of monitoring probability: $\eta_{\sigma^{1*}}(e_1) \leq \eta_{\sigma^{1*}}(e_2) \leq \dots \leq \eta_{\sigma^{1*}}(e_{|\mathcal{E}|})$. Then, $T' = \{e_1, \dots, e_{b_2}\}$ is a best response to σ^{1*} for **P2** (recall that **P2** uses all her resources; see Theorem 1). Let $T^0 = \{e_1, \dots, e_k\} \in 2^{\mathcal{E}}$ denote the components that are not monitored by any detector positioning $S \in \text{supp}(\sigma^{1*})$. Then, $\eta_{\sigma^{1*}}(e_1) = \dots = \eta_{\sigma^{1*}}(e_k) = 0$, and the value of the game Γ is:

$$U^*(b_1, b_2) = U(\sigma^{1*}, T') \stackrel{(2), (\text{EC.7}), (\text{EC.26})}{=} b_2 - \sum_{e \in T'} \eta_{\sigma^{1*}}(e) = b_2 - \sum_{i=k+1}^{b_2} \eta_{\sigma^{1*}}(e_i). \quad (\text{EC.27})$$

Now, to show the contradiction, we construct another strategy $\hat{\sigma}^1$ that will provide a better payoff than σ^{1*} to **P1**.

– **Case 1:** $k \geq b_2$. Then the value of the game is b_2 (it corresponds to zero detections). However, **P1** has an incentive to switch her strategy, and by randomizing over the nodes that can monitor the k components in T^0 , she will improve her payoff; which is a contradiction.

– **Case 2:** $k < b_2$. Then **P2** will randomize over attack plans that contain T^0 .

- **Case 2.1:** $k \geq b_1$. Then, $U^*(b_1, b_2) \leq b_2 - b_1$ (since **P1** can monitor b_1 components in T^0 that are always targeted). However, from Proposition 1, we have $U^*(b_1, b_2) \geq b_2 - \frac{b_1 b_2}{m^*} > b_2 - b_1$ (since $b_2 < m^*$). Therefore there is a contradiction.

- **Case 2.2:** $k < b_1$. Then, the idea is to construct another strategy that positions k detectors to monitor the components in T^0 (that were previously unmonitored), and that randomizes the positioning of the remaining $b_1 - k$ detectors over the node basis of σ^{1*} .

For now, assume that **P1** has $b_1 - k$ detectors. For any detector positioning $S \in \text{supp}(\sigma^{1*})$ (viewed as a set of nodes), we consider $\sigma^1(S, b_1 - k)$ defined in (EC.10) (in this case, we randomize the positioning of $b_1 - k$ detectors over the set S of size b_1). Recall that $\text{supp}(\sigma^1(S, b_1 - k)) = \{S^1, \dots, S^{b_1}\}$, and that for all $l \in \llbracket 1, b_1 \rrbracket$, $\sigma^1(S, b_1 - k)_{S^l} = \frac{1}{b_1}$. Now, let us construct the following inspection strategy:

$$\sigma^{1'} = \sum_{S \in \text{supp}(\sigma^{1*})} \sigma_S^{1*} \sigma^1(S, b_1 - k). \quad (\text{EC.28})$$

One can check that $\sigma^{1'} \in \Delta(\mathcal{A}_1)$, and is such that $\text{supp}(\sigma^{1'}) = \cup_{S \in \text{supp}(\sigma^{1*})} \{S^1, \dots, S^{b_1}\}$. Thus, $\sigma^{1'}$ is a probability distribution that randomizes over detector positionings of size $b_1 - k$. Then,

for every $S \in \text{supp}(\sigma^{1'})$, we augment S by positioning k additional detectors to monitor the subset of components T^0 that was previously unmonitored and that is always targeted in equilibrium by **P2**. We denote $\{i_1, \dots, i_k\}$ the positioning of such additional detectors, and we denote $\widehat{S} = S \cup \{i_1, \dots, i_k\}$ for every $S \in \text{supp}(\sigma^{1'})$ the augmented detector positioning. Then, we consider the probability distribution $\widehat{\sigma}^1$ with support equal to $\cup_{S \in \text{supp}(\sigma^{1*})} \{\widehat{S}^1, \dots, \widehat{S}^{b_1}\}$ and such that:

$$\forall S \in \text{supp}(\sigma^{1'}), \sigma_S^{1'} = \widehat{\sigma}_{\widehat{S}}^1, \quad (\text{EC.29})$$

i.e., $\widehat{\sigma}^1$ is the same probability distribution as $\sigma^{1'}$ except that it randomizes over the augmented detector positionings present in the support of $\sigma^{1'}$.

Then, we can derive the following calculations which combine the previous construction of $\widehat{\sigma}^1$ with a property of the detection function derived in (EC.16), and which will lead to a contradiction:

$$\begin{aligned} & \forall T \in \mathcal{A}_2 \mid T^0 \subset T \text{ and } |T| = b_2, \mathbb{E}_{\widehat{\sigma}^1}[\mathbf{F}(S, T)] \stackrel{(\text{EC.5})}{=} \underbrace{\mathbb{E}_{\widehat{\sigma}^1}[\mathbf{F}(S, T^0)]}_{=k} + \mathbb{E}_{\widehat{\sigma}^1}[\mathbf{F}(S, T \setminus T^0)] \\ & = k + \sum_{S \in \text{supp}(\sigma^{1*})} \sum_{l=1}^{b_1} \widehat{\sigma}_{\widehat{S}^l}^1 \mathbf{F}(\widehat{S}^l, T \setminus T^0) \stackrel{(\text{EC.4}), (\text{EC.29})}{\geq} k + \sum_{S \in \text{supp}(\sigma^{1*})} \sum_{l=1}^{b_1} \sigma_{S^l}^{1'} \mathbf{F}(S^l, T \setminus T^0) \\ & \stackrel{(\text{EC.10}), (\text{EC.28})}{=} k + \sum_{S \in \text{supp}(\sigma^{1*})} \sum_{l=1}^{b_1} \frac{1}{b_1} \sigma_S^{1*} \mathbf{F}(S^l, T \setminus T^0) \stackrel{(\text{EC.16})}{\geq} k + \sum_{S \in \text{supp}(\sigma^{1*})} \frac{b_1 - k}{b_1} \sigma_S^{1*} \mathbf{F}(S, T \setminus T^0) \\ & \stackrel{(\text{EC.7}), (\text{EC.26})}{=} k + \underbrace{\frac{b_1 - k}{b_1} \sum_{e \in T \setminus T^0} \eta_{\sigma^{1*}}(e)}_{>0} \geq k + \frac{b_1 - k}{b_1} \sum_{i=k+1}^{b_2} \eta_{\sigma^{1*}}(e_i) \\ & = \sum_{i=k+1}^{b_2} \eta_{\sigma^{1*}}(e_i) + k \left(1 - \frac{1}{b_1} \sum_{i=k+1}^{b_2} \eta_{\sigma^{1*}}(e_i) \right). \end{aligned} \quad (\text{EC.30})$$

From Proposition 1, we know that $U^*(b_1, b_2) \geq b_2 - \frac{b_1 b_2}{m^*}$. Therefore:

$$U(\sigma^{1*}, \sigma^{2*}) \stackrel{(\text{EC.27})}{=} b_2 - \sum_{i=k+1}^{b_2} \eta_{\sigma^{1*}}(e_i) \geq b_2 - \frac{b_1 b_2}{m^*} \iff \frac{1}{b_1} \sum_{i=k+1}^{b_2} \eta_{\sigma^{1*}}(e_i) \leq \frac{b_2}{m^*} < 1,$$

since $b_2 < m^*$. Then, by combining the previous inequality with (EC.30), we obtain:

$$\forall T \in \mathcal{A}_2 \mid T^0 \subset T \text{ and } |T| = b_2, U(\widehat{\sigma}^1, T) \stackrel{(2), (\text{EC.30})}{<} b_2 - \sum_{i=k+1}^{b_2} \eta_{\sigma^{1*}}(e_i) \stackrel{(\text{EC.27})}{=} U(\sigma^{1*}, \sigma^{2*}).$$

Since each attack plan in the support of an equilibrium strategy uses all the resources (Theorem 1), and must contain T^0 (beginning of Case 2), then: $U(\widehat{\sigma}^1, \sigma^{2*}) = \mathbb{E}_{\sigma^{2*}}[U(\widehat{\sigma}^1, T)] < U(\sigma^{1*}, \sigma^{2*})$, which violates the equilibrium condition (5).

Therefore, for all $(\sigma^{1*}, \sigma^{2*}) \in \Sigma$, $\mathcal{V}_{\sigma^{1*}}$ is a set cover.

(ii) Next, we show that both players must randomize their actions in equilibrium. Since the node basis in equilibrium is a set cover and $b_1 < n^*$, then **P1** must randomize her detector positionings in equilibrium. Now, assume that there exists a NE $(\sigma^{1^*}, T) \in \Sigma$ such that **P2** chooses a pure strategy T (of size b_2 from Theorem 1).

– If $b_1 \geq b_2$, then **P1** can detect all the attacks in T by positioning b_2 detectors at the nodes that can monitor the components of T , and the value of the game is 0. However, we showed in the proof of Theorem 1 that there exists a component outside of T that is not monitored with positive probability by σ^{1^*} . Therefore, **P2** can increase her payoff by targeting that component, thus leading to a contradiction.

– If $b_1 < b_2$, then **P1** can detect at least b_1 attacks in T by positioning detectors on b_1 nodes that can collectively monitor b_1 components of T . Thus, $U^*(b_1, b_2) \leq b_2 - b_1$. However, from Proposition 1, we know that $U^*(b_1, b_2) \geq b_2(1 - \frac{b_1}{m^*}) > b_2 - b_1$ (since $b_2 < m^*$), thus leading to a contradiction.

Therefore, in equilibrium, both players must randomize their actions. \square

PROPOSITION EC.1. *For $b_1 < n^*$ and $b_2 < m^*$, the expected detection rate in equilibrium of $\Gamma(b_1, b_2)$ is constant and bounded as follows:*

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

Furthermore, 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:

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

Proof of Proposition EC.1. Theorem 1 implies that for all $\sigma^* \in \Sigma$, $r(\sigma^*) \stackrel{(2),(3),(11)}{\geq} 1 - \frac{1}{b_2} U^*(b_1, b_2)$. Therefore, the equilibrium expected detection rate is constant. Then, the MSC/MSP-based bounds in Proposition 1 can be translated to bounds on the equilibrium detection rate as follows: For all $\sigma^* \in \Sigma$, $\frac{b_1}{n^*} \leq r(\sigma^*) \leq \min \left\{ \frac{b_1}{m^*}, 1 \right\} = \frac{b_1}{\max\{b_1, m^*\}}$. Finally, we can deduce that for all $\sigma^* \in \Sigma$, $\frac{\max\{b_1, m^*\}}{n^*} r(\sigma^*) \leq \frac{b_1}{n^*} \stackrel{(2),(3),(11),(EC.21)}{=} \min_{\sigma^2 \in \Delta(\mathcal{A}_2)} r(\sigma^1(S^{min}, b_1), \sigma^2)$. \square

LEMMA EC.4. *Given $b_2 < m^*$, let $\rho \in [0, 1]^{|\mathcal{E}|}$ that satisfies $\sum_{e \in \mathcal{E}} \rho_e = b_2$. Then, there exists an attack strategy $\sigma^2 \in \Delta(\mathcal{A}_2)$ that satisfies $\rho_{\sigma^2}(e) = \rho_e$ for all $e \in \mathcal{E}$.*

Proof of Lemma EC.4. Given $b_2 < m^*$, let \mathbf{A} be the $|\mathcal{E}| \times \binom{|\mathcal{E}|}{b_2}$ binary matrix whose rows (resp. columns) are indexed by the components (resp. the size- b_2 subsets) of \mathcal{E} , and which satisfies for all $(e, T) \in \mathcal{E} \times \overline{\mathcal{A}}_2$, $a_{e,T} = \mathbb{1}_{\{e \in T\}}$. Then, given $\rho \in [0, 1]^{|\mathcal{E}|}$ that satisfies $\sum_{e \in \mathcal{E}} \rho_e = b_2$, we must show that the following system of equations has a feasible solution:

$$\begin{aligned}\mathbf{A}\sigma^2 &= \rho \\ \mathbf{1}_{|\overline{\mathcal{A}_2}|}^T \sigma^2 &= 1 \\ \sigma^2 &\geq \mathbf{0}_{|\overline{\mathcal{A}_2}|}.\end{aligned}$$

Since each $T \in \overline{\mathcal{A}_2}$ is of size b_2 , it is easy to see that $\frac{1}{b_2} \mathbf{1}_{|\mathcal{E}|}^T \mathbf{A} = \mathbf{1}_{|\overline{\mathcal{A}_2}|}^T$. Furthermore, since $\frac{1}{b_2} \mathbf{1}_{|\mathcal{E}|}^T \rho = 1$, it implies that if σ^2 satisfies $\mathbf{A}\sigma^2 = \rho$, it also satisfies $\mathbf{1}_{|\overline{\mathcal{A}_2}|}^T \sigma^2 = 1$. Therefore, we only need to show that there exists $\sigma^2 \geq \mathbf{0}_{|\overline{\mathcal{A}_2}|}$ such that $\mathbf{A}\sigma^2 = \rho$. By Farkas' lemma, such a solution exists if and only if there does not exist $w \in \mathbb{R}^{|\mathcal{E}|}$ such that $w^T \mathbf{A} \leq \mathbf{0}_{|\overline{\mathcal{A}_2}|}^T$ and $w^T \rho > 0$.

Let $w \in \mathbb{R}^{|\mathcal{E}|}$ that satisfies $w^T \mathbf{A} \leq \mathbf{0}_{|\overline{\mathcal{A}_2}|}^T$, and let us order the components in \mathcal{E} so that $w_{e_1} \geq \dots \geq w_{e_{|\mathcal{E}|}}$. For notational simplicity, let $w_k := w_{e_k}$ and $\rho_k := \rho_{e_k}$ for all $k \in \llbracket 1, |\mathcal{E}| \rrbracket$. Note that since $T^1 = \{e_1, \dots, e_{b_2}\} \in \overline{\mathcal{A}_2}$, we have $\sum_{k=1}^{b_2} w_k = (w^T \mathbf{A})_{T^1} \leq 0$. Then, we obtain:

$$w^T \rho = \sum_{k=1}^{|\mathcal{E}|} w_k \rho_k = \underbrace{\sum_{k=1}^{b_2} w_k}_{\leq 0} + \sum_{k=b_2+1}^{|\mathcal{E}|} \underbrace{w_k}_{\leq w_{b_2}} \underbrace{\rho_k}_{\geq 0} - \sum_{k=1}^{b_2} \underbrace{w_k}_{\geq w_{b_2}} \underbrace{(1 - \rho_k)}_{\geq 0} \leq w_{b_2} \left(\sum_{k=1}^{|\mathcal{E}|} \rho_k - b_2 \right) = 0.$$

By Farkas' lemma, this implies that there exists an attack strategy $\sigma^2 \in \Delta(\overline{\mathcal{A}_2})$ such that for all $e \in \mathcal{E}$, $\rho_{\sigma^2}(e) = \rho_e$. \square

Proof of Theorem 2. Consider $\mathbf{P1}$'s amount of resources b_1 . If $b_1 \geq n^*$, then for every $b_2 < m^*$ and every $\sigma^* \in \Sigma(b_1, b_2)$, $r(\sigma^*) = 1$. Henceforth, we assume that $b_1 < n^*$. Recall from Theorem 1 that the NE of Γ can be obtained by solving $(\overline{\text{LP}}_1)$ and $(\overline{\text{LP}}_2)$.

(i.a) First, we show that for $b_2 = 1$, there exists an optimal solution of $(\overline{\text{LP}}_2)$, $\sigma^{2*} \in \Delta(\overline{\mathcal{A}_2})$, that satisfies for all $e \in \mathcal{E}$, $\rho_{\sigma^{2*}}(e) \leq \frac{1}{m^*}$. Since $b_2 = 1$, then $\overline{\mathcal{A}_2} = \mathcal{E}$, and for every $\sigma^2 \in \Delta(\mathcal{E})$ and every $e \in \mathcal{E}$, $\sigma_e^2 = \rho_{\sigma^2}(e)$.

– Consider an optimal solution of $(\overline{\text{LP}}_2)$, $\sigma^{2*} \in \Delta(\mathcal{E})$, and assume on the contrary that there exist $e' \in \mathcal{E}$ and $\varepsilon > 0$ such that $\rho_{\sigma^{2*}}(e') = \frac{1}{m^*} + \varepsilon$. Let $S^* \in \arg \min_{S \in \overline{\mathcal{A}_1}} U(S, \sigma^{2*})$. From Proposition 1, we know that:

$$\mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^*, T)] \stackrel{(2)}{=} -U(S^*, \sigma^{2*}) + 1 \leq -\max\left\{1 - \frac{b_1}{m^*}, 0\right\} + 1 \leq \frac{b_1}{m^*}. \quad (\text{EC.31})$$

Therefore, we can show that there exists $i' \in S^*$ such that $\mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^*, T)] - \mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^* \setminus \{i'\}, T)] \leq \frac{1}{m^*}$. Indeed, if for all $i \in S^*$, $\mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^*, T)] - \mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^* \setminus \{i\}, T)] > \frac{1}{m^*}$, then we obtain the following contradiction:

$$\begin{aligned}(b_1 - 1)\mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^*, T)] &\stackrel{(\text{EC.16})}{\leq} \sum_{i \in S^*} \mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^* \setminus \{i\}, T)] < b_1 \mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^*, T)] - \frac{b_1}{m^*} \\ &\stackrel{(\text{EC.31})}{\leq} (b_1 - 1)\mathbb{E}_{\sigma^{2*}}[\mathbf{F}(S^*, T)].\end{aligned}$$

– This implies that $e' \in \mathcal{C}_{S^*}$: If $e' \notin \mathcal{C}_{S^*}$ instead, then repositioning the detector on node i' to a node that can monitor e' improves $\mathbf{P1}$'s payoff by at least ε and contradicts the definition of S^* .

– Then, we show that at least b_1 components are monitored by S^* . If we assume that $|\mathcal{C}_{S^*}| < b_1$, then at least one detector in S^* can be removed without changing $\mathbf{P1}$'s payoff. Let $i_0 \in S^*$ denote the location of that detector. Now, we can show that there exists $e \in \mathcal{E} \setminus \{\mathcal{C}_{S^*}\}$ such that $\sigma_e^{2^*} > 0$: if on the contrary, we had for all $e \in \mathcal{E} \setminus \{\mathcal{C}_{S^*}\}$, $\sigma_e^{2^*} = 0$, then we would obtain $U(S^*, \sigma^{2^*}) = 0$, which is a contradiction: Since $b_1 < n^*$, the value of the game must be strictly positive. Let $e \in \mathcal{E} \setminus \{\mathcal{C}_{S^*}\}$ such that $\sigma_e^{2^*} > 0$ and let $i \in \mathcal{V}$ such that $F(i, e) = 1$. Repositioning the detector from node i_0 to node i will improve $\mathbf{P1}$'s payoff by at least $\sigma_e^{2^*}$, which contradicts the definition of S^* . Thus, $|\mathcal{C}_{S^*}| \geq b_1$.

– Now, we show that we can construct $\sigma^{2'}$ which is the same probability distribution as σ^{2^*} except that it reallocates ε probability from e' to a subset of components T^1 monitored by S^* while ensuring that the attack probability of each component in T^1 is not above $\frac{1}{m^*}$. Let us split \mathcal{C}_{S^*} into $\{e'\}$, $T^1 := \{e \in \mathcal{C}_{S^*} \setminus \{e'\} \mid \sigma_e^{2^*} < \frac{1}{m^*}\}$, and $T^2 := \{e \in \mathcal{C}_{S^*} \setminus \{e'\} \mid \sigma_e^{2^*} \geq \frac{1}{m^*}\}$. Now, we have:

$$\varepsilon = \mathbb{E}_{\sigma^{2^*}}[F(S^*, T)] - \sum_{e \in T^1 \cup T^2} \sigma_e^{2^*} - \frac{1}{m^*} \stackrel{\text{(EC.31)}}{\leq} \frac{|\mathcal{C}_{S^*}| - |T^2| - 1}{m^*} - \sum_{e \in T^1} \sigma_e^{2^*} = \sum_{e \in T^1} \left(\frac{1}{m^*} - \sigma_e^{2^*}\right),$$

which implies that $T^1 \neq \emptyset$, and that it is possible to allocate ε additional probability to components in T^1 so that the attack probability of each component in T^1 is not above $\frac{1}{m^*}$. Thus, $\sigma^{2'}$ can be constructed. It satisfies for all $e \in \mathcal{E} \setminus \{T^1 \cup \{e'\}\}$, $\sigma_e^{2'} = \sigma_e^{2^*}$, $\sigma_{e'}^{2'} = \sigma_{e'}^{2^*} - \varepsilon$, and $\sum_{e \in T^1} (\sigma_e^{2'} - \sigma_e^{2^*}) = \varepsilon$.

– Now, consider $S' \in \arg \min_{S \in \mathcal{A}_1} U(S, \sigma^{2'})$. The goal of this step is to show (by contradiction) that $e' \in \mathcal{C}_{S'}$. First, we derive the following calculations:

$$\begin{aligned} \mathbb{E}_{\sigma^{2'}}[F(S', T)] &= (\sigma_{e'}^{2^*} - \varepsilon)F(S', e') + \sum_{e \in \mathcal{E} \setminus \{T^1 \cup \{e'\}\}} \sigma_e^{2^*} F(S', e) + \sum_{e \in T^1} \underbrace{(\sigma_e^{2'} - \sigma_e^{2^*})}_{\geq 0} F(S', e) + \sum_{e \in T^1} \underbrace{\sigma_e^{2^*}}_{\leq 1} F(S', e) \\ &\leq \sum_{e \in \mathcal{E}} \sigma_e^{2^*} F(S', e) + \sum_{e \in T^1} (\sigma_e^{2'} - \sigma_e^{2^*}) - \varepsilon F(S', e') = \mathbb{E}_{\sigma^{2^*}}[F(S', T)] + \varepsilon(1 - F(S', e')). \end{aligned} \quad \text{(EC.32)}$$

Thus, if $e' \notin \mathcal{C}_{S'}$, then $\mathbb{E}_{\sigma^{2'}}[F(S', T)] \leq \mathbb{E}_{\sigma^{2^*}}[F(S', T)] + \varepsilon$. Let $i^* \in \mathcal{V} \setminus S'$ such that $e' \in \mathcal{C}_{i^*}$. Then, we have:

$$\begin{aligned} \forall i \in S', \mathbb{E}_{\sigma^{2^*}}[F(S^*, T)] &\stackrel{(2)}{=} 1 - U(S^*, \sigma^{2^*}) \geq 1 - U(\{i^*\} \cup S' \setminus \{i\}, \sigma^{2^*}) \\ &\stackrel{(2)}{=} \mathbb{E}_{\sigma^{2^*}}[F(\{i^*\} \cup S' \setminus \{i\}, T)] \stackrel{\text{(EC.4)}}{\geq} \frac{1}{m^*} + \varepsilon + \mathbb{E}_{\sigma^{2^*}}[F(S' \setminus \{i\}, T)]. \end{aligned} \quad \text{(EC.33)}$$

Next, we can derive the following calculations:

$$\mathbb{E}_{\sigma^{2^*}}[F(S', T)] \stackrel{\text{(EC.16)}}{\leq} \frac{1}{b_1 - 1} \sum_{i \in S'} \mathbb{E}_{\sigma^{2^*}}[F(S' \setminus \{i\}, T)] \stackrel{\text{(EC.31), (EC.33)}}{\leq} \mathbb{E}_{\sigma^{2^*}}[F(S^*, T)] - \frac{b_1}{b_1 - 1} \varepsilon. \quad \text{(EC.34)}$$

Combining everything together, we obtain the following contradiction:

$$\min_{S \in \mathcal{A}_1} U(S, \sigma^{2'}) \stackrel{(2), \text{(EC.32), (EC.34)}}{\geq} U(S^*, \sigma^{2^*}) + \frac{1}{b_1 - 1} \varepsilon > U(S^*, \sigma^{2^*}) = \max_{\sigma^2 \in \Delta(\mathcal{E})} \min_{S \in \mathcal{A}_1} U(S, \sigma^2).$$

– Therefore, we showed that $e' \in \mathcal{C}_{S'}$. This implies that $\varepsilon(1 - F(S', e')) = 0$, and we obtain: $\min_{S \in \overline{\mathcal{A}}_1} U(S, \sigma^{2'}) \stackrel{(2), (\text{EC.32})}{\geq} U(S', \sigma^{2*}) \geq \min_{S \in \overline{\mathcal{A}}_1} U(S, \sigma^{2*}) = \max_{\sigma^2 \in \Delta(\mathcal{E})} \min_{S \in \overline{\mathcal{A}}_1} U(S, \sigma^2)$. Thus, $\sigma^{2'}$ is also an optimal solution of $(\overline{\text{LP}}_2)$. Therefore, if an optimal solution of $(\overline{\text{LP}}_2)$ is such that at least one component is targeted with probability more than $\frac{1}{m^*}$, we can create another optimal solution of $(\overline{\text{LP}}_2)$ with one less component targeted with probability more than $\frac{1}{m^*}$. We can then repeat this process until all attack probabilities are no more than $\frac{1}{m^*}$.

(i.b) Given $b_1 < n^*$, let $z^*(b_2)$ denote the optimal value of $(\overline{\text{LP}}_2)$ for any $b_2 < m^*$. Now, consider $b_2 < m^*$, and let $\sigma^{2*} \in \Delta(\mathcal{A}_2)$ be an optimal solution of $(\overline{\text{LP}}_2)$. Since $\sum_{e \in \mathcal{E}} \frac{\rho_{\sigma^{2*}}(e)}{b_2} = 1$, we can construct an attack strategy $\sigma^{2'} \in \Delta(\mathcal{E})$ such that for all $e \in \mathcal{E}$, $\sigma_e^{2'} = \rho_{\sigma^{2'}}(e) = \frac{\rho_{\sigma^{2*}}(e)}{b_2}$. Then, the additivity of F gives: $z^*(b_2) = \min_{S \in \overline{\mathcal{A}}_1} (b_2 - \sum_{e \in \mathcal{E}} F(S, e) \rho_{\sigma^{2*}}(e)) = b_2 \min_{S \in \overline{\mathcal{A}}_1} U(S, \sigma^{2'}) \leq b_2 z^*(1)$.

Now, consider $\tilde{\sigma}^2 \in \Delta(\mathcal{E})$ which is an optimal solution of $(\overline{\text{LP}}_2)$ (where the number of attack resources is 1) with the additional property that for all $e \in \mathcal{E}$, $\rho_{\tilde{\sigma}^2}(e) \leq \frac{1}{m^*}$. Then, given $b_2 < m^*$, since for all $e \in \mathcal{E}$, $b_2 \rho_{\tilde{\sigma}^2}(e) \leq 1$ and $\sum_{e \in \mathcal{E}} b_2 \rho_{\tilde{\sigma}^2}(e) = b_2$, there exists a probability distribution $\hat{\sigma}^2 \in \Delta(\mathcal{A}_2)$ that satisfies $\rho_{\hat{\sigma}^2}(e) = b_2 \rho_{\tilde{\sigma}^2}(e)$ for all $e \in \mathcal{E}$ (Lemma EC.4). Then, by additivity of F , we obtain: $z^*(1) = \min_{S \in \overline{\mathcal{A}}_1} (1 - \sum_{e \in \mathcal{E}} F(S, e) \rho_{\hat{\sigma}^2}(e)) = \frac{1}{b_2} \min_{S \in \overline{\mathcal{A}}_1} (b_2 - \sum_{e \in \mathcal{E}} F(S, e) \rho_{\hat{\sigma}^2}(e)) \leq \frac{1}{b_2} z^*(b_2)$.

Thus, for every $b_2 < m^*$, $z^*(b_2) = b_2 z^*(1)$. Therefore, we can conclude that for every $\sigma^* \in \Sigma(b_1, b_2)$, $r(\sigma^*) \stackrel{(2), (3), (11)}{=} 1 - \frac{U(\sigma^*)}{b_2} = 1 - \frac{z^*(b_2)}{b_2} = 1 - z^*(1) =: r_{b_1}^*$, which does not depend on b_2 .

(ii) Given $\mathbf{P1}$'s resources $b_1 < n^*$, let $\sigma^{1*} \in \Delta(\overline{\mathcal{A}}_1)$ be an inspection strategy in equilibrium of the game $\Gamma(b_1, 1)$. From Theorem 1, we know that σ^{1*} is an optimal solution of $(\overline{\text{LP}}_1)$ for $b_2 = 1$. Now, consider $b_2 < m^*$. We can derive the following inequality:

$$\forall T \in \overline{\mathcal{A}}_2, U(\sigma^{1*}, T) \stackrel{(\text{EC.7})}{=} \sum_{e \in T} U(\sigma^{1*}, e) \leq \sum_{e \in T} \underbrace{\max_{e' \in \mathcal{E}} U(\sigma^{1*}, e')}_{=1-r_{b_1}^*} = b_2(1 - r_{b_1}^*) = z^*(b_2).$$

Since this inequality is valid for any $T \in \overline{\mathcal{A}}_2$, we deduce that $\max_{T \in \overline{\mathcal{A}}_2} U(\sigma^{1*}, T) \leq z^*(b_2) = \min_{\sigma^1 \in \Delta(\overline{\mathcal{A}}_1)} \max_{T \in \overline{\mathcal{A}}_2} U(\sigma^1, T)$. Therefore, σ^{1*} is an optimal solution of $(\overline{\text{LP}}_1)$ (when the number of attack resources is b_2), and is an inspection strategy in equilibrium of the game $\Gamma(b_1, b_2)$. \square

EC.3. Additional Illustrations

EC.3.1. Illustrative Example

We illustrate the solution approach outlined in Section 4 on a small-sized pipeline network of 23 nodes and 34 pipelines (Giustolisi et al. 2008), as shown in Figure EC.2a. Consider the network inspection problem (\mathcal{P}) , where the defender wants to ensure an expected detection rate of $\alpha = 0.75$ against an adversary with $b_2 \leq m^*$ resources. First, we determine the detection model \mathcal{G} : 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, we determine the monitoring sets \mathcal{C}_i for detecting break events from each location $i \in \mathcal{V}$, as illustrated in Figure EC.2a.

MSC/MSP-based solution: We solve $(\mathcal{I}_{\text{MSC}})$ to compute an MSC $S^{\min} = \{1, 8, 9, 21\}$ of size $n^* = 4$, as shown in Figure EC.2b. From (13), we know that with $\lceil \alpha n^* \rceil = 3$ detectors, the expected detection rate in any equilibrium is at least α . Using Lemma EC.3, we construct the inspection strategy $\sigma^1(S^{\min}, 3)$, which uniformly randomizes over $\text{supp}(\sigma^1(S^{\min}, 3)) = \{\{1, 8, 9\}, \{8, 9, 21\}, \{9, 21, 1\}, \{21, 1, 8\}\}$. Then, we deduce from (14) that $(3, \sigma^1(S^{\min}, 3))$ is a feasible solution of (\mathcal{P}) . Next, we solve $(\mathcal{I}_{\text{MSP}})$ to compute an MSP T^{\max} , which is of size $m^* = 3$, as illustrated in Figure EC.2b. From (13), we deduce that the optimality gap associated with $(3, \sigma^1(S^{\min}, 3))$ is given by $\lceil \alpha n^* \rceil - \lceil \alpha m^* \rceil = \lceil 3 \rceil - \lceil 2.25 \rceil = 0$. Therefore, $(3, \sigma^1(S^{\min}, 3))$ is an optimal solution of (\mathcal{P}) . Although $\sigma^1(S^{\min}, 3)$ satisfies constraints (4), we note that its detection performance is lower than the detection performance in equilibrium: The upper bound on the corresponding loss in detection performance derived from (14) is given by $1 - \frac{\max\{\lceil \alpha n^* \rceil, m^*\}}{n^*} = 25\%$. Figure EC.2c illustrates the bounds on the equilibrium expected detection rate derived in (13).


Refinement Procedure: Next, we improve our MSC-based inspection strategy using the refinement procedure described in Section 4.2. We first select $\mathcal{I} = \{\{1, 8, 9\}, \{8, 9, 21\}, \{9, 21, 1\}, \{21, 1, 8\}\}$ as subset of indices for the restricted master problem $(M_{\text{CG}}(\mathcal{I}))$ to warm-start the column generation algorithm. After 17 iterations of (A5)-(A11), we obtain an equilibrium inspection strategy σ^{1*} that uniformly randomizes over $\text{supp}(\sigma^{1*}) = \{\{1, 8, 9\}, \{8, 9, 21\}, \{21, 1, 8\}, \{5, 12, 17\}, \{1, 10, 17\}, \{6, 9, 21\}\}$. The corresponding node basis is given by $\mathcal{V}_{\sigma^{1*}} = \{1, 5, 6, 8, 9, 10, 12, 17, 21\}$, as illustrated in Figure EC.2d. The equilibrium detection rate is $\frac{5}{6}$, which improves the detection performance of $\sigma^1(S^{\min}, 3)$ by 11.11%. Finally, we find that the optimal dual variables of $(\overline{\text{LP}}_1)$, obtained at termination of the column generation algorithm, imply that an attack strategy in equilibrium of $\Gamma(3, 1)$ uniformly randomizes over the 6 pipes highlighted in Figure EC.2d. Thus, for every $b_2 \leq m^* = 3$, an attack strategy in equilibrium of $\Gamma(3, b_2)$ is such that each of these 6 pipelines is targeted with probability $\frac{b_2}{6}$.

EC.3.2. Additional Computational Results of Column Generation


In Figure EC.3, we illustrate the results of the column generation algorithm (A4) for solving $(\overline{\text{LP}}_1)$ with the optimal number of detectors in networks ky5, ky13, and ky2.

EC.4. Case When $b_2 \geq m^*$


As argued in Section 3.2, the network inspection problem (\mathcal{P}) and the game $\Gamma(b_1, b_2)$ when $b_2 \geq m^*$ are of limited practical interest. However, for the sake of completeness, we now briefly discuss this case. To solve problem (\mathcal{P}) , we derived in Section 3.3 equilibrium properties of the game $\Gamma(b_1, b_2)$




(a) Network layout. The colored region indicates the locations from where a particular break event can be detected.



(b) MSC (blue nodes) and MSP (red pipes).



(c) Equilibrium detection rate with respect to the number of detectors b_1 .



(d) Node basis (blue nodes) and component basis (red crosses) in equilibrium of $\Gamma(3, b_2)$ for $b_2 \leq 3$.

Figure EC.2 Solving (\mathcal{P}) for the Apulian benchmark network facing adversarial pipeline break events.

that hold when $b_1 < n^*$ and $b_2 < m^*$. Note that all these properties, except Proposition 3, also hold when $b_1 < n^*$ and $b_2 = m^*$. This implies that our approach to solve (\mathcal{P}) is still valid when $b_2 = m^*$: the optimality gap and detection guarantees of the MSC/MSP-based solution hold, and can be improved with the refinement procedure.

However, most of the equilibrium properties described in Section 3.3 are not always satisfied by the NE of $\Gamma(b_1, b_2)$ when $b_1 < n^*$ and $b_2 > m^*$, as discussed in the following example:

EXAMPLE EC.2. Consider the detection model $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \{C_i, i \in \mathcal{V}\})$ defined as follows: Let $\mathcal{V} = \{i_1, \dots, i_n\}$, with $n \in \mathbb{N}$, and let $\mathcal{E} = \mathcal{E}_1 \cup \{e_1, \dots, e_n\}$, where \mathcal{E}_1 is a discrete set. Then, the monitoring




Figure EC.3 Results of column generation applied to (\overline{LP}_1) for networks ky5, ky13, and ky2 with and without the MSC-based warm-start.

sets are defined by $\mathcal{C}_{i_k} = \mathcal{E}_1 \cup \{e_k\}$, for every $k \in \llbracket 1, n \rrbracket$. In this example, $S^{min} = \{i_1, \dots, i_n\}$ is an MSC, and $T^{max} = \{e_1, \dots, e_n\}$ is an MSP; so $n^* = m^* = n$. Given any $b_1 < n$ and any $b_2 \in \llbracket n, |\mathcal{E}_1| + n \rrbracket$, one can check that for all $T \in \mathcal{A}_2$ such that $T^{max} \subseteq T$, $(\{i_1, \dots, i_{b_1}\}, T)$ is a pure NE, whose node basis is not a set cover. Therefore, Proposition 3 does not hold anymore.

Now, if we consider $b_2 = |\mathcal{E}_1| + n$, we just showed that we could construct NE where **P2** can use $n, n+1, \dots$, or $n + |\mathcal{E}_1| = b_2$ resources. Thus, (11) in Theorem 1 does not hold anymore. This also implies that the expected detection rate is not constant in equilibrium anymore: We found equilibria where the equilibrium detection rates are equal to $\frac{b_1}{n}, \frac{b_1+1}{n+1}, \dots, \frac{b_1+|\mathcal{E}_1|}{n+|\mathcal{E}_1|}$, which violates Theorem 2. Since $\frac{b_1+|\mathcal{E}_1|}{n+|\mathcal{E}_1|} \xrightarrow{|\mathcal{E}_1| \rightarrow +\infty} 1$, the MSP-based upper bound on the expected detection rates given in (13) is violated. Furthermore, the upper bound on the relative loss of performance derived from (14) is not valid anymore. By choosing $\sigma^1(S^{min}, b_1)$, the expected detection rate may be arbitrarily far from an equilibrium expected detection rate. \triangle

Still, some results remain valid when $b_1 < n^*$ and $b_2 > m^*$: In Theorem 1, property (10) still holds. In (13), the MSC-based lower bound on the equilibrium expected detection rates is still valid. If **P1** selects the strategy $\sigma^1(S^{min}, b_1)$, the expected detection rate is guaranteed to be at least $\frac{b_1}{n^*}$, regardless of **P2**'s strategy. From these remaining results, we can deduce that $(\lceil \alpha n^* \rceil, \sigma^1(S^{min}, \lceil \alpha n^* \rceil))$ is a feasible solution to (\mathcal{P}) when $b_2 > m^*$. However, the optimality gap derived in Section 4.1 does not hold anymore.