



Detection and mitigation of attacks on transportation networks as a multi-stage security game



Aron Laszka^{a,*}, Waseem Abbas^b, Yevgeniy Vorobeychik^c, Xenofon Koutsoukos^d

^a University of Houston, Houston, USA

^b Information Technology University, Lahore, Pakistan

^c Washington University in St. Louis, St. Louis, USA

^d Vanderbilt University, Nashville, USA

ARTICLE INFO

Article history:

Received 11 March 2019

Revised 18 June 2019

Accepted 17 July 2019

Available online 2 August 2019

Keywords:

Security

Game theory

Transportation network

Cyber-physical system

Intrusion detection

Computational complexity

ABSTRACT

In recent years, state-of-the-art traffic-control devices have evolved from standalone hardware to networked smart devices. Smart traffic control enables operators to decrease traffic congestion and environmental impact by acquiring real-time traffic data and changing traffic signals from fixed to adaptive schedules. However, these capabilities have inadvertently exposed traffic control to a wide range of cyber-attacks, which adversaries can easily mount through wireless networks or even through the Internet. Indeed, recent studies have found that a large number of traffic signals that are deployed in practice suffer from exploitable vulnerabilities, which adversaries may use to take control of the devices. Thanks to the hardware-based failsafes that most devices employ, adversaries cannot cause traffic accidents directly by setting compromised signals to dangerous configurations. Nonetheless, an adversary could cause disastrous traffic congestion by changing the schedule of compromised traffic signals, thereby effectively crippling the transportation network. To provide theoretical foundations for the protection of transportation networks from these attacks, we introduce a game-theoretic model of launching, detecting, and mitigating attacks that tamper with traffic-signal schedules. We show that finding optimal strategies is a computationally challenging problem, and we propose efficient heuristic algorithms for finding near optimal strategies. We also introduce a Gaussian-process based anomaly detector, which can alert operators to ongoing attacks. Finally, we evaluate our algorithms and the proposed detector using numerical experiments based on the SUMO traffic simulator.

© 2019 Elsevier Ltd. All rights reserved.

1. Introduction

The evolution of traffic signals from standalone hardware devices to complex networked systems has provided society with many benefits, such as reducing wasted time and environmental impact. However, it has also exposed traffic signals to a variety of cyber-attacks. While traditional hardware systems were susceptible only to attacks based on direct physical access, modern systems are vulnerable to attacks through wireless interfaces or even to remote attacks through the Internet. To assess the severity of these threats in practice, Ghena et al. recently analyzed the security of real-world traffic infrastructure in cooperation with a road agency located in Michigan (Ghena et al., 2014). This agency operates around a hundred traffic signals, which are all part of the same wireless network, but the signals at every intersection operate independently of the other intersections. The study found three

major weaknesses in the traffic infrastructure: lack of encryption for the wireless network, lack of secure authentication due to the use of default usernames and passwords on the devices, and the presence of exploitable software vulnerabilities.

While all of these known weaknesses could be eliminated, it is extremely difficult to ensure that devices are free of any unknown weaknesses. In general, it is virtually impossible—or prohibitively expensive—to ensure that a system is perfectly secure. In addition to the general difficulty of attaining perfect security, traffic-control devices pose further challenges. Similar to other distributed cyber-physical systems, traffic-control systems have large attack surfaces, and they often have long system lifetime and complicated software-upgrade procedures, which makes fixing vulnerabilities difficult. Consequently, operators cannot hope to stop all cyber-attacks since a determined and sophisticated attacker might always find a way to compromise some of the devices. Therefore, instead of focusing solely on the first line of defense, operators must also consider minimizing the impact of successful cyber-attacks.

* Corresponding author.

E-mail address: alaszka@uh.edu (A. Laszka).

Due to hardware-based failsafes, compromising a traffic signal does not allow an attacker to set the signal to an unsafe configuration that would lead to traffic accidents, such as giving green light to two intersecting directions. However, compromising a signal does enable tampering with its schedule, which allows the attacker to cause disastrous traffic congestion. Such malicious cyber-attacks may be launched by any adversary whose interest is to cause disruption and damage, ranging from cyber-terrorists to disgruntled ex-employees. For instance, during the Los Angeles traffic engineers' 2006 strike, two disgruntled employees allegedly penetrated the traffic-control system of the city, and reprogrammed the traffic lights of four intersections to cause congestion: "[t]he red signal would be on too long for the critical approach and the green signal would be on too long for the noncritical approach, thus resulting in long backups into the airport and other key intersections around the city" (Bernstein and Blankstein, 2007). Furthermore, terrorists could also mount these attacks in conjunction with physical attacks, thereby increasing their impact (e.g., delaying ambulances and firefighters).

To minimize the impact of attacks tampering with traffic signals, operators must be able to detect and mitigate them promptly and effectively. In practice, the detection of novel cyber-attacks poses multiple challenges. Since signature-based detectors are ineffective against novel attack, operators must employ anomaly-based detectors. However, these detectors are prone to raising false alarms, which must be investigated manually, resulting in a waste of manpower and resources. Considering the relative scarcity of attacks, the cost of these investigations may exceed the benefit of early attack detection and mitigation. Therefore, when configuring the sensitivity of detectors, operators must carefully balance the cost of false alarms and the risk from delayed detection. Moreover, sophisticated attackers can act strategically by mounting stealthy attacks, which delay detection but still cause significant impact. In light of this, operators must also plan their defense strategically, by anticipating the attackers' responses.

1.1. Contributions

In Laszka et al. (2016b), we introduced an approach for evaluating the vulnerability of transportation networks to cyber-attacks that tamper with traffic-control devices. In this paper, we extend this approach by considering detectors and countermeasures that operators can implement to mitigate these attacks. In particular, we introduce a game-theoretic model, in which an operator can setup anomaly-based detectors and mitigate ongoing attacks by reconfiguring traffic control. Similar to Laszka et al. (2016b), we build on the cell-transmission model introduced by Daganzo (1994, 1995). To the best of our knowledge, our work is the first to consider the problem of designing and deploying systems based on traffic-sensors measurements to detect tampering attacks against traffic control. Our main contributions in this paper are the following:

- We formulate a multi-stage security game that models the detection and mitigation of cyber-attacks against transportation networks.
- We propose an efficient metaheuristic search algorithm for finding detector configurations that minimize losses in the face of strategic attacks.
- We introduce an anomaly-based detector for attacks against traffic control, which is built on a Gaussian-process based model of normal traffic.
- We evaluate our detector and algorithms based on detailed simulations of traffic using SUMO.

Table 1
List of symbols.

Symbol	Description
Traffic model	
Q_i	Maximum number of vehicles that can flow into or out of cell i
δ_i	Ratio between free-flow speed and backward propagation speed of cell i
N_i	Maximum number of vehicles in cell i
d_i	Demand (inflow) at source cell i
$\Gamma^{-1}(i)$	Set of predecessor cells to cell i
P_{ki}^*	Inflow proportion from cell k to signalized intersection i
\mathcal{I}	Set of signalized intersections
Game constants and functions	
\mathcal{I}_D	Set of signalized intersections with attack detectors
B	Attacker's budget for compromising traffic signals
$\Delta_D(\mathbf{D}, \mathbf{A})$	Detection delay for detector configuration \mathbf{D} and attack \mathbf{A}
Δ_M	Length of mitigation (before returning to normal operation)
T	Congestion with default traffic control
$T_A(\mathbf{A})$	Congestion as a result of attack \mathbf{A} (before mitigation)
C	Cost of investigating a false alarm
$T_M(\mathbf{A}, \mathbf{M})$	Congestion after mitigation \mathbf{M} of attack \mathbf{A}
$\mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M})$	Attacker's gain for actions $(\mathbf{D}, \mathbf{A}, \mathbf{M})$
$\mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{M})$	Defender's loss for actions $(\mathbf{D}, \mathbf{A}, \mathbf{M})$
Game variables	
\mathbf{D}	Defender's detector configuration action (Stage I)
D_i	False-positive rate of detector at intersection $i \in \mathcal{I}_D$
$\mathbf{A} = (\mathcal{I}_A, \hat{\mathbf{p}})$	Attacker's attack action (Stage II)
\mathcal{I}_A	Set of signals compromised by the attacker
\hat{p}_{ki}	Modified (by attack or mitigation) inflow proportion from cell k to intersection i
\mathbf{M}	Defender's mitigation action (Stage III)

1.2. Outline

The remainder of this paper is organized as follows. In Section 2, we introduce our game-theoretic model of detecting and mitigating attacks against transportation networks. In Section 3, we present computational results on our model and propose efficient heuristic algorithms. In Section 4, we introduce a Gaussian-process based detector for attacks against traffic control. In Section 5, we use detailed simulations of transportation networks to evaluate our detector and the heuristic algorithms. In Section 6, we discuss related work on the vulnerability of transportation networks, configuration of attack detectors, and game theory for security of cyber-physical systems. Finally, in Section 7, we offer concluding remarks.

2. Game-theoretic model of attacks on traffic signals

In this section, we introduce our model of launching, detecting, and mitigating cyber-attacks against traffic control in transportation networks. Our model includes two agents: an attacker who can launch cyber-attacks and a defender who attempts to detect and mitigate them. Since these agents may anticipate and react to each other's actions, we formulate our model using game theory, which enables us to capture the strategic interactions between the two agents. For a list of symbols used in our model, see Table 1.

2.1. Traffic model

First, we introduce Daganzo's cell transmission model, the traffic model on which our game-theoretic model, our analysis, and our numerical evaluation are built. Here, we provide only a very brief summary of this traffic model, focusing on the notation that will be used throughout the paper. For a detailed description of the

model, we refer the reader to [Daganzo \(1994\)](#), [Daganzo \(1995\)](#) and [Ziliaskopoulos \(2000\)](#).¹

The cell transmission model divides a road network into *cells*, which represent homogeneous road segments, and divides time into uniform *intervals*. The length of a road segment corresponding to a cell is equal to the distance traveled in light traffic by a typical vehicle in one time interval. Each cell i has three sets of parameters: N_i^t is the maximum number of vehicles that can be present in cell i at time t ; Q_i^t is the maximum number of vehicles that can flow into or out of cell i during time interval t ; and δ_i^t is the ratio between the free-flow speed and the backward propagation speed of cell i at time t (see [Ziliaskopoulos, 2000](#) for a detailed explanation).² Cells that model road segments where vehicles can enter traffic are called *source cells*, and each source cell i has a traffic demand parameter d_i^t , which is the number of vehicles entering traffic at cell i in time interval t . Cells where vehicles may exit traffic are called *sink cells*.

Every cell is connected to one or more other cells: cells that correspond to consecutive road segments or road segments that are joined by an intersection are connected. The set of cells from which vehicles can move into cell i is called the set of *predecessor cells*, denoted by $\Gamma^{-1}(i)$. Cell that have multiple predecessors are called *merging cells*, while cells that are the predecessors of multiple cells are called *sink cells*. To model signal control at intersections, we follow Daganzo's proposition ([Daganzo, 1995](#)) and introduce the time-dependent parameters p_{ki}^t controlling the inflow proportions of merging cell i . We let \mathcal{I} denote the set of merging cells that model signalized intersections.

To solve the traffic model (i.e., to determine the traffic flow for a given network and set of parameters), we use Ziliaskopoulos's linear programming approach ([Ziliaskopoulos, 2000](#)). The objective of this linear program is the sum of the number of vehicles traveling (i.e., number of vehicles on the road) over time, which is clearly equal to the *total travel time* of all the vehicles. As a consequence, we can use the value of the linear program—which can be computed efficiently for a given instance—as a measure of network congestion.

Finally, we consider relatively short-term attack scenarios, in which the parameters of the cells and the default (i.e., unattacked) schedules of the traffic signals are constant. Hence, in our game-theoretic model, we will omit the superscript t from Q_i^t , N_i^t , δ_i^t , and p_{ki}^t .

2.2. Multi-stage security game

We model defensive countermeasures and attacks in a transportation network as a two-player multi-stage security game between a defender and an attacker ([Laszka et al., 2019](#)). The defender represents the operator of the transportation network, who can configure traffic-control devices and aims to minimize congestion in the network. The attacker represents any strategic adversary that can compromise and tamper with traffic signals and aims to maximize congestion.

In a nutshell, our game consists of the following three stages.

- I. **Detector Configuration:** In the first stage, the defender configures detectors, which are deployed in the transportation network, to detect cyber-attacks against traffic control. The detectors may be traffic-anomaly based detectors or conventional cyber-security intrusion detection systems (IDS). When configuring these detectors, the defender should anticipate the attacker's possible adversarial actions in the second stage.

- II. **Attack on Traffic Control:** In the second stage, the attacker mounts a cyber-attack against the transportation network by compromising traffic signals and tampering with their schedule to cause congestion. When choosing its attack, the attacker must take into account both the detector configuration chosen by the defender in the first stage as well as the defender's possible mitigation actions in the third stage.
- III. **Mitigation of Attack:** In the third stage, the defender attempts to mitigate the attack by changing the configuration of uncompromised traffic-control devices to minimize congestion. As this is the final stage, mitigation simply needs to respond to the particular attack that was launched in the second stage. Note that once the defender detects an ongoing attack, it should also try to regain control of the compromised devices as soon as possible. However, since the devices may be physically scattered throughout the transportation network, regaining control of them can take a long time. For instance, an attacker could have changed remote login passwords, severed communication-network connections, etc., forcing the defender to physically reset or reinstall compromised devices. Meanwhile, disastrous traffic congestions may form in the transportation network, which the defender must mitigate immediately.

2.2.1. Stages and strategic choices

Next, we provide a detailed description of the three stages of the game and the players' action spaces.

Stage I: Detector Configuration To detect stealthy cyber-attacks, detectors are deployed on the traffic-control devices at a subset \mathcal{I}_D of the signalized intersections \mathcal{I} . These detectors can be either traffic-anomaly based detectors, such as the one that we will introduce in [Section 4](#), or conventional cyber-security intrusion detection systems. We assume that detectors are imperfect, which means that they may raise false alarms (when there is no attack in progress) and they may detect actual attacks with some delay. The rate of false-positive errors and detection delay both depend on how sensitive a detector is: a more sensitive detector is more likely to raise false alarms but detects actual attacks earlier, and vice versa. We assume that the operator can configure the sensitivity of every one of the $|\mathcal{I}_D|$ detectors individually.

Specifically, in the first stage of the game, the defender chooses a sensitivity configuration \mathbf{D} for the detectors. For ease of presentation, we let the sensitivity of the detector at each intersection $i \in \mathcal{I}_D$ be represented by the false-positive rate of the detector. Formally, a detector configuration \mathbf{D} is an $|\mathcal{I}_D|$ -dimensional non-negative vector, where D_i is the rate of false alarms raised by the detector at intersection $i \in \mathcal{I}_D$.

Stage II: Cyber-Attack on Traffic Control In the second stage of the game, the attacker compromises a subset of the traffic signals and changes their schedule. We let $\mathcal{I}_A \subseteq \mathcal{I}$ denote the set of traffic signals that the attacker chooses to compromise. We assume that the attacker is resource bounded, which means that it can compromise signals in at most $B \leq |\mathcal{I}|$ intersections at the same time. Hence, the attacker's choice \mathcal{I}_A must satisfy $|\mathcal{I}_A| \leq B$.

Once the attacker has compromised a set of traffic signals \mathcal{I}_A , it can reconfigure every one of them. However, most traffic control devices have hardware-level safety mechanisms in practice (see, e.g., [Chena et al., 2014](#)), which constrain the configurations that may be set by an adversary. In particular, traffic signals typically use malfunction management units as a safety feature against controller faults (e.g., overriding a faulty controller that would give green lights to two intersecting directions). These hardware-based failsafes also limit the impact of cyber attacks (e.g., preventing the attacker from causing traffic accidents) by overriding invalid configurations. In our traffic model, the attacker's reconfiguration corresponds to setting new inflow proportions \hat{p}_{ki} for the cells in \mathcal{I}_A . Therefore, we can model hardware-level failsafes by requiring the

¹ For readers who are familiar with the cell-transmission model, we recommend to continue with [Section 2.2](#).

² This constant is used to quantify how the speed of traffic decreases as the cell becomes congested, and can model traffic phenomena such as shockwaves.

inflow proportions chosen by the attacker to constitute a valid configuration. Specifically, we assume that the inflow proportions set by a feasible attack must sum up to 1 for each compromised intersection:

$$\forall i \in \mathcal{I}_A : \sum_{k \in \Gamma^{-1}(i)} \hat{p}_{ki} = 1. \quad (1)$$

In sum, we can represent a feasible attack action \mathbf{A} as a pair $\mathbf{A} = (\mathcal{I}_A, \hat{\mathbf{p}})$ that satisfies $|\mathcal{I}_A| \leq B$ and $\sum_{k \in \Gamma^{-1}(i)} \hat{p}_{ki} = 1$ for every $i \in \mathcal{I}_A$, where \mathcal{I}_A is the set of compromised signals, and $\hat{\mathbf{p}}$ are the tampered signal schedules.

Between Stages II and III: Detection Once the attacker has perpetrated its attack in the second stage, the compromised signals begin to operate with tampered schedules, which results in increased congestion in the transportation network. Eventually, the detectors deployed by the defender will detect the attack (based on either traffic or cyber anomalies). We let Δ_D denote the detection delay, that is, the amount of time between the launch and detection of the attack. The detection delay depends on both the configuration \mathbf{D} chosen by the defender and the attack \mathbf{A} chosen by the attacker, which we express by representing delay as a function $\Delta_D(\mathbf{D}, \mathbf{A})$ of \mathbf{D} and \mathbf{A} . Once the attack is detected, the game progresses to the third stage.

Stage III: Mitigation of Attack. In the third stage, the defender mitigates the detected attack by reconfiguring traffic-control devices to alleviate congestion. We assume that the defender can reconfigure any device that is still under its control, that is, any traffic signal that is not compromised by the attacker. Since the attacker has compromised signals \mathcal{I}_A , the defender can set new inflow proportions \hat{p}_{ki} for the cells $i \in \mathcal{I} \setminus \mathcal{I}_A$. We again require the new configuration to be valid, which means that the inflow proportions must sum up to 1 for each reconfigured intersection:

$$\forall i \in \mathcal{I} \setminus \mathcal{I}_A : \sum_{k \in \Gamma^{-1}(i)} \hat{p}_{ki} = 1. \quad (2)$$

Finally, we let $\mathbf{M} = \{\hat{p}_{ki} \mid i \in \mathcal{I} \setminus \mathcal{I}_A\}$ denote the defender's mitigation (i.e., reconfiguration) action.

2.2.2. Player's utilities

We now define the defender's and the attacker's utilities resulting from the various strategic choices that they can make in the game. First, we let T denote the level of congestion in the transportation network with the default configuration of traffic-control devices (i.e., with inflow proportions \mathbf{p}). In practice, we can quantify congestion T as, e.g., the average or total travel time of the vehicles in the transportation network between their origin and destination. Recall from Section 2.1 that we can efficiently compute travel time with default proportions \mathbf{p} in our traffic model using a linear program.

Next, we let $T_A(\mathbf{A})$ denote the level of congestion after the attack but before the mitigation, which depends on the attacker's action \mathbf{A} chosen in the second stage. Similar to T , we can compute $T_A(\mathbf{A})$ using our traffic model with the default proportion p_{ki} for every cell $i \in \mathcal{I} \setminus \mathcal{I}_A$ but with the adversarial proportion \hat{p}_{ki} for every cell $i \in \mathcal{I}_A$. Finally, we let $T_M(\mathbf{A}, \mathbf{M})$ denote the level of congestion after the attack has been mitigated, which depends on both attacker's action \mathbf{A} and the defender's mitigation action \mathbf{M} .

We let the attacker's gain \mathcal{G} (i.e., positive utility) for actions $(\mathbf{D}, \mathbf{A}, \mathbf{M})$ be the total impact of the attack in terms of increased congestion level:

$$\mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}) = (T_A(\mathbf{A}) - T) \cdot \Delta_D(\mathbf{D}, \mathbf{A}) + (T_M(\mathbf{A}, \mathbf{M}) - T) \cdot \Delta_M, \quad (3)$$

where Δ_M is the amount of time between mitigation and the transportation network returning to normal operation (e.g., manually resetting compromised devices). The first term quantifies the impact of the attack before mitigation, while the second term

quantifies impact after mitigation but before returning to normal operation.

Next, we define the defender's loss (i.e., negative utility) resulting from actions $(\mathbf{D}, \mathbf{A}, \mathbf{M})$. Recall that the detectors deployed in the transportation network are imperfect, and each detector $i \in \mathcal{I}_D$ is continuously generating false alerts (i.e., false-positive errors) at rate D_i . Since the defender cannot tell which alerts are false, it has to investigate every single alert, which costs manpower and resources. Hence, we define the defender's loss considering both the total impact of attacks and the cost of investigating false alerts:

$$\mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{M}) = \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}) + \sum_{i \in \mathcal{I}_D} D_i \cdot C, \quad (4)$$

where C is the cost of investigating an alert. The first term quantifies the total impact of the attack, while the second term captures the cost of investigating false alerts.

2.3. Solution concept and problem formulation

We assume that both players have perfect information: in the second stage, the attacker knows the detector configuration \mathbf{D} chosen by the defender in the first stage; and in the third stage, the defender knows the attack action \mathbf{A} chosen by the attacker in the second stage. We assume that the attacker has perfect information because we are considering a sophisticated, worst-case attacker, who has extensive knowledge of its target (i.e., Kerckhoffs's principle) and may know the algorithms or techniques employed by the defender for configuring the detectors. On the other hand, we assume that the defender has perfect information because we are considering a smart transportation network with monitoring capabilities; however, this assumption could be relaxed.

Under this assumption, we can model the players' optimal choices most naturally using the solution concept of *subgame perfect equilibrium*. Our goal is to find an optimal strategy for the defender by solving the game. We can do so by finding the players' best-response actions for every stage using backward induction, that is, by solving each stage starting with the third and finishing with the first, in each stage building on the solutions for the subgame formed by the subsequent stages. Given detector configuration \mathbf{D} and attack \mathbf{A} , a best-response mitigation is

$$\underset{\mathbf{M}}{\operatorname{argmin}} \mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{M}). \quad (5)$$

Note that the best-response mitigation does not actually depend on the detector configuration \mathbf{D} . To prove this, observe that the only term of $\mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{M})$ that depends on mitigation action \mathbf{M} is $(T_M(\mathbf{A}, \mathbf{M}) - T) \cdot \Delta_M$. Since this term does not depend on detector configuration \mathbf{D} , neither does the best-response mitigation. Intuitively, the explanation for this is that once the defender has detected the attack, it does not matter how it was detected (and all costs associated with detection are sunk).

Given detector configuration \mathbf{D} , a best-response attack is

$$\underset{\mathbf{A}}{\operatorname{argmax}} \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}) \Big|_{\mathbf{M} \in \operatorname{argmin}_{\mathbf{M}'} \mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{M}')}$$

$$= \underset{\mathbf{A}}{\operatorname{argmax}} \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}) \Big|_{\mathbf{M} \in \operatorname{argmin}_{\mathbf{M}'} \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}') + \sum_{i \in \mathcal{I}_D} D_i \cdot C} \quad (6)$$

$$= \underset{\mathbf{A}}{\operatorname{argmax}} \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}) \Big|_{\mathbf{M} \in \operatorname{argmin}_{\mathbf{M}'} \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}')} \quad (7)$$

$$= \underset{\mathbf{A}}{\operatorname{argmax}} \min_{\mathbf{M}} \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}). \quad (8)$$

Note that the attacker must anticipate the defender's mitigation action in the next stage; however, since the game is strategically equivalent to a zero-sum game, the attacker's problem simplifies to a maximin optimization. Our threat model assumes a worst-case attacker, whose goal is to minimize the defender's utility. This is

a safe assumption since the defender's utility can only be higher if the attacker behaves differently. In contrast, if our threat model assumed a particular attacker behavior, then the defender's strategy would be vulnerable to deviations from that assumed behavior.

Finally, an optimal (i.e., equilibrium) detector configuration is

$$\begin{aligned} & \operatorname{argmin}_{\mathbf{D}} \min_{\mathbf{M}} \mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{M}) \Big|_{\mathbf{A} \in \operatorname{argmax}_{\mathbf{A}'} \min_{\mathbf{M}'} \mathcal{G}(\mathbf{D}, \mathbf{A}', \mathbf{M}')} \\ & = \operatorname{argmin}_{\mathbf{D}} \max_{\mathbf{A}} \min_{\mathbf{M}} \mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{M}). \end{aligned} \quad (9)$$

Note that the defender needs to anticipate the attacker's attack action in the next stage; however, since the game is strategically equivalent to a zero-sum game, the defender's problem simplifies to a minimax optimization (i.e., minimaximin if we consider the mitigation choice as well).

3. Analysis

Even though we can express the players' optimal strategies as relatively simple maximin and minimax optimization problems, actually finding optimal strategies is computationally challenging due to the sizes of the players' strategy spaces. Hence, we focus our analysis on the computational aspects of solving the transportation security game. First, in Section 3.1, we show that finding an optimal action for the attacker is a computationally hard problem. Although we study the complexity of only the attacker's problem in this paper, our computational-complexity argument could be easily extended to the defender's problem. Due to the complexity of solving the game, we focus the remainder of this section on providing efficient heuristic algorithms: we introduce a greedy heuristic for the attacker in Section 3.2.1, and a metaheuristic search algorithm for the defender in Section 3.2.2.

3.1. Computational complexity

We begin our analysis by showing that the attacker's problem (i.e., finding a worst-case attack) is computationally hard. Given a detector configuration \mathbf{D} , the attacker's problem is to find an optimal attack \mathbf{A}^* that maximizes $\min_{\mathbf{M}} \mathcal{G}(\mathbf{D}, \mathbf{A}^*, \mathbf{M})$. Following the backward-induction approach, we assume that we have an oracle that finds the optimal mitigation action \mathbf{M} for any attack \mathbf{A} , and we study the attacker's problem by building on this oracle. In practice, the oracle can be replaced with, for example, a linear program for finding optimal traffic control. Further, for ease of presentation, we overload the notation \mathcal{G} as follows

$$\mathcal{G}(\mathbf{D}, \mathbf{A}) = \min_{\mathbf{M}} \mathcal{G}(\mathbf{D}, \mathbf{A}, \mathbf{M}). \quad (10)$$

First, we formulate a decision version of the attacker's problem as follows.

Definition 1. Attacker's Decision Problem: Given a transportation network, a budget B , a detector configuration \mathbf{D} , and a threshold gain \mathcal{G}^* , determine if there exists an attack \mathbf{A}^* satisfying the budget constraint such that $\mathcal{G}(\mathbf{D}, \mathbf{A}^*) > \mathcal{G}^*$.

We show that the above problem is computationally hard by reducing a well-known NP-hard problem, the Set Cover Problem, to the above problem.

Definition 2. Set Cover Problem: Given a base set U , a collection \mathcal{C} of subsets of U , and a number k , determine if there exists a sub-collection $\mathcal{C}' \subseteq \mathcal{C}$ of at most k subsets such that every element of U is contained by at least one subset in \mathcal{C}' .

The following theorem establishes the computational complexity of the attacker's problem.

Theorem 1. The Attacker's Decision Problem is NP-hard.




Fig. 1. Illustration for the proof of Theorem 1.

Proof. Given an instance of the Set Cover Problem (i.e., a set U , a collection \mathcal{C} of subsets, and a number k), we construct an instance of the Attacker's Decision Problem as follows:

- let the transportation network be the following (see Fig. 1 for an illustration):
 - there is one source cell r , with $Q_r = k + 1$, $d_r^1 = k + 1$, and $d_r^t = 0$ for $t > 1$;
 - there is one sink cell s ;
 - for every element $u \in U$, there is a merging cell u ;
 - for every subset $C \in \mathcal{C}$, there is a diverging cell C ;
 - each diverging cell C is connected to every merging cell $u \in C$;
 - for every cell i , $N_i = k + 1$ and $\delta_i = 1$;
 - for every merging cell u , $Q_u = k + 1$;
 - for every diverging cell C , $Q_C = 1$;
- let the attacker's budget be $B = |U|$;
- let the detector configuration be such that $\forall \mathbf{A}: \Delta_D(\mathbf{D}, \mathbf{A}) \equiv 1$
- let the default congestion be $T = 0$, let the congestion after the attack $T_A(\mathbf{A})$ be equal to the total travel time of the vehicles, and let the mitigation time be $\Delta_M = 0$;
- let the threshold gain be $\mathcal{G}^* = 3(k + 1)$.

Clearly, the above reduction can be carried out in time that is polynomial in the size of the Set Cover Problem instance.

It remains to show that the above instance of the Attacker's Decision Problem has a solution \mathbf{A}^* if and only if the given instance of the Set Cover Problem has a solution \mathcal{C}' . Before we proceed to prove this equivalence, notice that the values Q_r , N_i and δ_i for every cell i , and Q_u for every merging cell u will not play any role, since they are high enough to allow any traffic to pass through. Furthermore, since $B = |U|$ and $\Delta_D \equiv 1$, the attacker will be able to reconfigure every traffic signal without decreasing detection time. Hence, the attacker's problem is simply to pick the values \hat{p}_{Cu} for every $u \in C$ so that the total travel time is at least $\mathcal{G}^* = 3(k + 1)$.

First, suppose that there exists a set cover \mathcal{C}' of size at most k . Then, we construct an attack as follows: for every merging cell u , choose one diverging cell C from \mathcal{C}' that is connected to u (if there are multiple, then choose an arbitrary one), and let $\hat{p}_{Cu} = 1$. We have to show that the total travel time in the transportation network is greater than $3(k + 1)$ after the attack. Since the distance between the source cell and the sink cell is 3 hops and there are $k + 1$ vehicles, all the vehicles must move one step closer to the sink in every time interval in order for the total travel time to be at most $3(k + 1)$. However, from the source cell, the vehicles may only move to the cells in \mathcal{C}' ; otherwise, they would get "stuck" at one of the diverging cells that are not in \mathcal{C}' . Consequently, in the second time interval, at most k of the $k + 1$ vehicles may move on, which means that the total travel time has to be greater than $3(k + 1)$.

Second, suppose that there does not exist a set cover \mathcal{C}' of size at most k . Then, we have to prove that there cannot

exist an attack which increases the total travel time to more than $3(k+1)$. Firstly, we show that there exists an optimal attack which assigns either 0 or 1 to every \hat{p}_{Cu} . To prove this, consider an attack in which there is a merging cell v with a \hat{p}_{Cv} value other than 0 or 1. If none of its predecessor cells C has a positive \hat{p}_{Cw} value for some other merging cell w , then the assignment for v can clearly be changed to 0 and 1 values without changing the total travel time. Next, suppose that one (or more) of the predecessor cells C of the merging cell has a positive \hat{p}_{Cw} value for some other merging cell w . Then, the total travel time maximizing assignment is clearly one which assigns $\hat{p}_{Cv} = 1$ to a predecessor cell C for which $\sum_{u \in C} \hat{p}_{Cu}$ is maximal, since this “wastes” the most “merging capacity.” Thus, for the remainder of the proof, it suffices to consider only attacks where every \hat{p}_{Cu} value is either 0 or 1.

Now, consider an optimal attack \mathbf{A}^* against the transportation network, and let C^* be the set of diverging cells C for which there exists a merging cell u such that $\hat{p}_{Cu} = 1$. Clearly, C^* forms a set cover of U since for every element u , there is a subset $C \in C^*$ such that $u \in C$ (i.e., C is connected to u). From our initial supposition, it follows readily that the cardinality of set C^* must be at least $k+1$. However, this also implies that the total travel time after the attack is equal to $3(k+1)$: in the second time interval, all $k+1$ vehicles may move forward to the diverging cells in set C^* ; in the third time interval, all the vehicles may again move forward to the merging cells (since every cell in C has at least one “enabled” connection); and all the vehicles may leave the network by the next interval through the sink cell. Since the total travel time after an optimal attack \mathbf{A}^* is equal to $T^* = 3(k+1)$, the attacker’s problem does not have a solution. Therefore, the constructed instance of the Attacker’s Decision Problem has a solution if and only if the given instance of the Set Cover Problem has one, which concludes our proof. \square

3.2. Algorithms

Mitigation—in our model—means adapting the schedule of uncompromised traffic signals given the schedule of compromised signals, which is equivalent to optimizing traffic control in a non-adversarial setting, with the compromised signals acting as fixed-schedule signals. Since optimizing traffic control in non-adversarial settings has been studied in prior work, we focus on providing efficient algorithms for solving the first two stages of the game.

3.2.1. Greedy algorithm for attacks

Since the attacker’s problem is NP-hard, we cannot hope for a polynomial-time algorithm that always finds a worst-case attack (unless $P = NP$). Hence, to provide an alternative to computationally infeasible exhaustive search, we turn our attention to designing an efficient heuristic algorithm.

The attacker’s problem can be viewed as the composition of two problems: finding a subset \mathcal{I}_A of at most B signalized intersections and finding new inflow proportions \hat{p}_{ki} for the cells $i \in \mathcal{I}_A$. For finding a subset \mathcal{I}_A , we propose to use a greedy heuristic, which starts with an empty set and adds new cells to it one-by-one, always picking the one that leads to the greatest increase in the attacker’s gain. Finding new inflow proportions \hat{p}_{ki} is particularly challenging, since the set of possible choices is continuous. However, we observe that in most networks, the worst-case configuration is an “extreme” one, which assigns proportion $\hat{p}_{ki} = 1$ to one predecessor cell k and proportion $\hat{p}_{ji} = 0$ to every other predecessor cell j (Laszka et al., 2016b). In fact, we have tested this property on hundreds of networks that we generated according to the Grid model with Random Edges (see Section 5.2), resembling road networks from the U.S. and Europe, and we have not found a single counterexample. Based on these numerical results,

we conjecture that this property holds in general, that is, for any network.³ Hence, for every new cell i added to the set of attacked intersections, we propose to search over the possible extreme configurations by iterating over the predecessors of cell i . Based on the above ideas, we formulate Algorithm 1.

Algorithm 1: Polynomial-Time Greedy Heuristic for Finding an Attack.

Data: transportation network security game, detector configuration \mathbf{D}

Result: attack \mathbf{A}^*

```

 $\mathcal{I}_A^* \leftarrow \emptyset, \hat{\mathbf{p}}^* \leftarrow \mathbf{p};$ 
for  $b = 1, \dots, B$  do
   $\mathcal{I}'_A \leftarrow \mathcal{I}_A^*, \hat{\mathbf{p}}' \leftarrow \hat{\mathbf{p}}^*;$ 
  for  $i \in \mathcal{I}$  do
     $\mathcal{I}_A \leftarrow \mathcal{I}_A^* \cup \{i\};$ 
    for  $k \in \Gamma^{-1}(i)$  do
       $\forall l, j: \hat{p}_{lj} \leftarrow \begin{cases} 1 & \text{if } j = i \wedge l = k \\ 0 & \text{if } j = i \wedge l \in \Gamma^{-1}(i) \setminus \{k\}; \\ \hat{p}_{lj} & \text{otherwise.} \end{cases}$ 
      if  $\mathcal{G}(\mathbf{D}, (\mathcal{I}_A, \hat{\mathbf{p}})) \geq \mathcal{G}(\mathbf{D}, (\mathcal{I}'_A, \hat{\mathbf{p}}'))$  then
         $\mathcal{I}'_A \leftarrow \mathcal{I}_A, \hat{\mathbf{p}}' \leftarrow \hat{\mathbf{p}};$ 
      end
    end
  end
   $\mathcal{I}_A^* \leftarrow \mathcal{I}'_A, \hat{\mathbf{p}}^* \leftarrow \hat{\mathbf{p}}';$ 
end
output  $\mathbf{A}^* = (\mathcal{I}_A^*, \hat{\mathbf{p}}^*)$ 

```

It is fairly easy to see that we can implement Algorithm 1 as a polynomial-time algorithm. Due to the three nested iterations, the running time of Algorithm 1 is $O(B \cdot |\mathcal{I}| \cdot (\max_{i \in \mathcal{I}} |\Gamma^{-1}(i)|))$ times the running time of computing \mathcal{G} . Since we can compute $\mathcal{G}(\mathbf{D}, \mathbf{A})$ for any attack \mathbf{A} using a linear program, it follows readily that the running time of the algorithm can be upper bounded by a polynomial function of the input size (i.e., size of the transportation network and budget B). We can formally state this observation as the following proposition.

Proposition 1. *With a polynomial-time oracle for computing \mathcal{G} , the running time of Algorithm 1 is a polynomial function of the input size.*

3.2.2. Metaheuristic search algorithm for detector configuration

Next, we present an algorithm for finding a detector configuration (i.e., false-positive rates) based on a metaheuristic approach. In particular, we use *simulated annealing* to find a near-optimal detector configuration \mathbf{D} . The basic idea of this approach is to start with an arbitrary configuration \mathbf{D} , which we then improve iteratively. In each iteration, we generate a new solution \mathbf{D}' in the neighborhood of \mathbf{D} . If the new configuration \mathbf{D}' is better in terms of minimizing the defender’s loss (against an attacker playing a best response), then the current configuration \mathbf{D} is replaced with the new one. On the other hand, if the new configuration \mathbf{D}' increases the defender’s loss, then new configuration replaces the current one with only a small probability. This probability depends on the difference between the two solutions in terms of loss as well as a parameter commonly referred to as “temperature,” which is a decreasing function of the number of iterations. These random replacements prevent the search from “getting stuck” in a local minimum. The algorithm is presented below as Algorithm 2.

In Algorithm 2, $\text{Perturb}(\mathbf{D})$ picks a random configuration \mathbf{D}' from the neighborhood of \mathbf{D} . In particular, we implement

³ We leave the theoretical proof of this claim for future work.

Algorithm 2: Polynomial-Time Metaheuristic for Finding a Detector Configuration.

Data: transportation network security game, iterations k_{\max} , initial temperature T_0 , cooling parameter β

Result: detector configuration D^*

$D \leftarrow \mathbf{1}$;

$L \leftarrow \max_{\mathbf{A}} \mathcal{G}(D, \mathbf{A}) + \sum_{i \in \mathcal{I}_D} D_i \cdot C$;

for $k = 1, \dots, k_{\max}$ **do**

$D' \leftarrow \text{Perturb}(D)$;

$L' \leftarrow \max_{\mathbf{A}} \mathcal{G}(D', \mathbf{A}) + \sum_{i \in \mathcal{I}_{D'}} D'_i \cdot C$;

$T \leftarrow T_0 \cdot e^{-\beta k}$;

$pr \leftarrow e^{(L'-L)/T}$;

if $(L' < L) \vee (\text{rand}(0, 1) \leq pr)$ **then**

$D \leftarrow D', L \leftarrow L'$

end

end

output D

$\text{Perturb}(D)$ as choosing a value for each D'_i uniformly at random from $[D_i \cdot (1 - \varepsilon), D_i \cdot (1 + \varepsilon)]$, where ε is a small constant (e.g., 0.1). For solving $\max_{\mathbf{A}} \mathcal{G}(D, \mathbf{A})$ in practice, we can use the greedy heuristic (Algorithm 1). The temperature T is decreasing exponentially with iteration number k , and the rate of the decrease is controlled by the “cooling” parameter β . Finally, we note that a simpler algorithm could also be obtained, in which D is updated with D' in each iteration if and only if D' is strictly better than D . This heuristic search, commonly known as *hill climbing*, also works well for our problem; however, Algorithm 2 gives better results.

4. Anomaly-based detector

Now, we introduce a traffic-anomaly based detector against stealthy attacks that tamper with traffic control. The core idea of anomaly-based detection is to build a probabilistic model of normal traffic conditions, which can then be used to estimate the likelihood that observed traffic conditions are normal. Note that we must employ a probabilistic model to account for the uncertainty in parameter values since many parameters (e.g., traffic demand) can only be estimated in practice. We can estimate the likelihood that the observed traffic is normal as the probability that our model of normal traffic would generate the observed traffic. We can then compare the likelihood value to a threshold, and if the likelihood is lower, we can raise an alarm. In our detector, the model of normal traffic is based on Gaussian processes, which have been successfully used in prior work for traffic volume forecasting (Chen et al., 2012; Xie et al., 2010). Note that we cannot use macroscopic traffic models, such as the cell transmission model, to detect attacks because these models abstract away details for the sake of tractability (e.g., inflow proportions instead of actual traffic light schedules); however, such details can be crucial for the detection of stealthier attacks that alter traffic control only slightly.

4.1. Gaussian processes

We begin giving a very brief overview of Gaussian processes. For a comprehensive discussion of Gaussian processes in machine learning, we refer the reader to Rasmussen and Williams (2006).

In principle, Gaussian processes are an extension of multivariate Gaussian distributions to infinite collections of random variables. Formally, a Gaussian process is a stochastic process such that any finite collection of variables (X_1, \dots, X_n) follows a multivariate Gaussian distribution. A Gaussian process is typically described us-

ing a *mean function*

$$\mu(X) = \mathbb{E}(X) \quad (11)$$

and a *covariance function*

$$k(X_1, X_2) = \mathbb{E}[(X_1 - m(X_1))(X_2 - m(X_2))]. \quad (12)$$

The covariance function is often chosen to be some well-known kernel function, such as squared exponential, whose parameters can be estimated from a training dataset (x_1, \dots, x_n) . A common application of Gaussian processes is regression: given the values of a set of training variables (x_1, \dots, x_n) , we can easily compute the expected value and variance of a target variable Y using the mean and covariance functions. Gaussian-process based regression models have been successfully applied to a wide range of problem, such as traffic volume forecasting (Chen et al., 2012; Xie et al., 2010), spatial modeling of extreme snow depth (Blanchet et al., 2011), wind power forecasting (Kou et al., 2013), estimation of water chlorophyll concentration (Bazi et al., 2012), and spectrum sensing (Nevat et al., 2012).

4.2. Model

We assume that traffic sensors, such as induction loop sensors, have been deployed for monitoring the transportation network. Since modeling an entire network would be computationally challenging—and would certainly not scale well—we divide sensors into subsets, and we build a separate model and detector for each one of these subsets. For example, traffic sensors that are deployed next to the same intersection $i \in \mathcal{I}$ may be grouped together and provide traffic data for one detector (see, e.g., Fig. 2 in Section 5.1). The outputs of all the detectors deployed in a transportation network can then be combined together to form a single detector for the entire network.

We assume that sensors measure and report traffic values, such as traffic flow or occupancy, in fixed-length intervals (e.g., report one measurement value for every 15-minute intervals).⁴ In our Gaussian-process model, we model each one of these measurements as a random variable. Formally, for every sensor s and time interval t , there exists a random variable X_s^t whose value is equal to the traffic value measured by sensor s in interval t . Hence, our model has a discrete but—due to the time dimension—potentially infinite set of variables.

A key part of modeling is establishing mean and covariance functions for these variables. For each sensor s , we model the mean values $\mu(X_s^t)$ of variables X_s^t using a periodic function:

$$\mu(X_s^t) \equiv \mu(X_s^{t+P}), \quad (13)$$

where P is the length of the period. We call this time period, which is measured in number of discrete time intervals, the *detector window* W . Similarly, for sensors s_1 and s_2 , we model the covariance values $k(X_{s_1}^{t_1}, X_{s_2}^{t_2})$ between variables $X_{s_1}^{t_1}$ and $X_{s_2}^{t_2}$ using a periodic function:

$$k(X_{s_1}^{t_1}, X_{s_2}^{t_2}) \equiv k(X_{s_1}^{t_1+P}, X_{s_2}^{t_2+P}). \quad (14)$$

Further, we assume that $k(X_{s_1}^{t_1}, X_{s_2}^{t_2}) \equiv 0$ if $|t_1 - t_2| > P$. To train the model, the actual values of these functions must be learned from traffic values observed during normal operation.

4.3. Training

Before training, our model has $S \cdot P + \frac{S \cdot (S-1) \cdot P \cdot (2P+1)}{2} + S \cdot P \cdot (2P+1)$ unknown values, where S is the number of sensors:

⁴ Note that the length of these measurement intervals is independent of the time intervals of the cell-transmission model. However, for ease of presentation, we will reuse notation t to identify measurement intervals.




Fig. 2. Intersection used for evaluating the detector. Yellow rectangles represent induction-loop sensors.

- $S \cdot P$ mean values: for each sensor s , the mean function μ can be described by P values since its period length is P ;
- $\frac{S \cdot (S-1) \cdot P \cdot (2P+1)}{2}$ covariance values: for each distinct pair of sensors s_1 and s_2 , the covariance function k can be described by $P \cdot (2 \cdot P + 1)$ values since its period length is P , the maximum difference between t_1 and t_2 is P , and covariance values are symmetric;
- and $S \cdot P \cdot (2P + 1)$ variance values: for each sensor s , the covariance function k can be described by $P \cdot (2 \cdot P + 1)$.

Training the model means learning these mean and covariance values for normal traffic. In practice, we can train the model by observing sensor measurements $(x_{s_1}^{t_1}, x_{s_1}^{t_2}, x_{s_2}^{t_1}, \dots)$ of traffic under normal conditions, and then simply estimating the most likely mean and covariance values from these observations (i.e., maximum likelihood estimation).

4.4. Detection

Once we have trained the model, we can use it to detect attacks against traffic control. First, we take sensor measurements (x_s^t, \dots) , which are observed in the network that might be under attack, and we use the Gaussian process to compute the likelihood of these measurement values being generated by our model of normal traffic.⁵ Since the measurement values are continuous, we can use the probability density of the Gaussian distribution (X_s^t, \dots) at (x_s^t, \dots) as the likelihood value. We then interpret this likelihood as the likelihood of the network operating under normal control (i.e., not being under attack). Finally, we com-

⁵ Due computational limitations, we restrict detection to observations from a single detector window (i.e., measurement values from some range $(t, t + P - 1)$). If more observations are available, we can evaluate the detector multiple times.

pare the likelihood to a *detector threshold* τ_i , and raise an alarm if the likelihood is lower than the threshold. Thus, our detector has two parameters, the detector window W and the detector threshold τ_i , which together determine the rate of false alarms D_i and the detection delay.

Note that the detector threshold τ_i and the rate of false alarms D_i are closely related to each other: lower thresholds result in fewer false alarms since more observations are accepted as likely; and vice versa. Hence, we can express $\tau_i(D_i)$ and $D_i(\tau_i)$ as increasing functions, and we may specify the configuration of the detector either as the desired false-positive rate D_i or as the threshold τ_i . We chose to use representation D_i in our game-theoretic model and analysis for ease of presentation. The exact relation $\tau_i(D_i)$ can be determined experimentally by evaluating the detector with various configurations on normal traffic.

5. Numerical results

In this section, we present numerical results on our heuristic algorithms from Section 3 and our anomaly-based detector from Section 4.

5.1. Anomaly-based attack detection

We begin by training and evaluating our detector based on simulated flows of traffic under normal conditions and under attacks. We will then use the results of this evaluation to instantiate our game-theoretic model. In particular, we will use the measured false-alarms rate and detection delay values as numerical parameters for our game-theoretic model.

Table 2
Posterior predictive checking p -values.

Sensor	Test statistic T				
	Mean	Variance	Median	Quantile $Q(0.3)$	Quantile $Q(0.7)$
East in	0.54	0.52	0.54	0.79	0.49
East out	0.51	0.47	0.66	0.78	0.36
South in	0.46	0.53	0.49	0.79	0.26
South out	0.49	0.48	0.66	0.79	0.28
West in	0.54	0.52	0.55	0.80	0.49
West out	0.50	0.47	0.66	0.78	0.36
North in	0.47	0.52	0.64	0.84	0.57
North out	0.49	0.48	0.66	0.79	0.28

5.1.1. Setup

For these experiments, we simulate the signalized four-way intersection shown by Fig. 2 using SUMO (Simulation of Urban Mobility),⁶ a well-known and widely-used micro simulator (Behrisch et al., 2011; Krajzewicz et al., 2002). Signals are deployed on both the incoming and outgoing lanes (represented by yellow rectangles in the figure), and these signals measure traffic flow in 15-second intervals. We generate one month of traffic data with the original signal schedule (45 seconds for both roads, including left-turning and yellow phases) for training the Gaussian-process based model. In these simulations, vehicles enter the intersection from all four directions, and each car turns left, continues straight, or turns right with probability 5.3%, 73.7%, and 21.1%, respectively. From each direction, 0.19 vehicles arrive each second on average. We also generate one month of test traffic-data with the original schedule for measuring the false-positive rate of the detector. Finally, for each of the attacks considered below, we generate one day of traffic data, which includes 1 hour with the original schedule and then 23 hours with the tampered schedule.

Posterior Predictive Check. To confirm that our Gaussian-process based traffic model fits observations well, we perform *posterior predictive checking* (Gelman et al., 2013). The idea of posterior predictive checking is to draw simulated samples from the joint posterior predictive distribution and compare them to the observed sample. If there were significant systematic differences between the simulated and observed samples, that would indicate that our model did not fit well. Note that the applicability of our model is also demonstrated by the low false-positive rate exhibited by our detector.

The significance of difference can be quantified as the classical p -value (Gelman et al., 2013, Chapter 6.3):

$$p = \Pr[T(x^{\text{rep}}) \geq T(x) | \mu, k], \quad (15)$$

where x^{rep} is the replicated data generated according to our model, x is the observed data, μ and k are the model parameters (see Section 4), and T is a test statistic. In our checks, we compute p values for various standard test statistics, including mean, variance (with Bessel's correction), median, and quantiles. Note that a perfectly fitting model will yield p -values around 0.5. To reliably estimate the probability p , we generate and evaluate the test statistics on 10,000 replicated samples (drawn independently according to our trained model).

Table 2 shows the p -values for various statistical tests (p -values around 0.5 indicate a perfect fit). Since our samples are multidimensional (i.e., one value for each sensor around the intersection), we apply the statistical tests to the marginal distributions corresponding to the individual sensors. We list sensors in clockwise order, starting with the sensors on the incoming and outgoing lanes

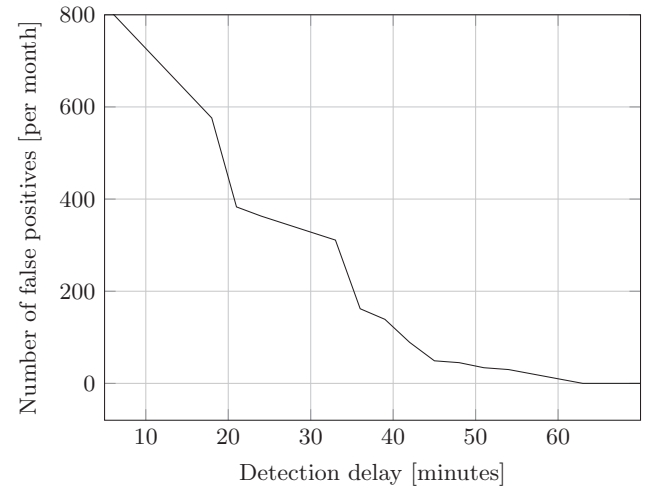


Fig. 3. Trade-off between false-positive rate and detection delay.

of the road eastward of the intersection (see Fig. 2), denoted 'east in' and 'east out.' For mean and variance statistics, we see that our model produces an almost perfect fit for all sensors, which is important since these play key role in determining likelihood, on which our detector is built. Further, we see that our model produces a good fit for the median statistic as well, which indicates that there is no significant asymmetry that could not be captured by our model. Finally, we test quantiles $Q(0.3)$ and $Q(0.7)$, and see a reasonable fit for most sensors. The worst fit is for sensor 'north in' (i.e., inward lane of the road northwards) and $Q(0.3)$; however, we see an almost perfect fit for the same sensor for $Q(0.7)$, indicating a skewed distribution. Next, we study the detection performance of our model, showing that observations under normal conditions and under attacks exhibit significantly different likelihood values.

5.1.2. Detector configuration

We first consider the defender's problem of balancing the number of false-positive errors and the detection delay. For this experiment, we assume an attack which changes 4.4% of the traffic-signal schedule.

Fig. 3 shows the trade-off between the false-positive rate and the detection delay. Each point on the curve is a Pareto optimal point that is attainable with some detector window W and threshold τ . The figure shows that with a negligible false-positive rate, even the stealthy attack considered in this example can be detected in approximately one hour. The configuration of the detector when the false-positive rate reaches zero is detector window being equal to $W = 48$ minutes and log-likelihood threshold being equal to $\ln \tau = -112.58$.

Fig. 4 shows the likelihood values output by the Gaussian-process model for traffic data resulting from the original and the attacked traffic-signal schedules. For this figure, we set the detector window to be $W = 3$ minutes, which results in highly variable likelihood values. The figure shows that after one hour (i.e., when the attack starts), the likelihood values for the tampered schedule become much lower than for the original one. In other words, the detector correctly estimates that the traffic with tampered schedule is less likely to be normal.

5.1.3. Stealthy attacks

Next, we consider the attacker's problem of balancing stealthiness and impact. If stealthy attacks could avoid detection for extended periods of time while having substantial impact on traffic, they could pose a significant threat to the transportation network. To show that stealthy attacks are not effective against our

⁶ http://sumo.dlr.de/wiki/Main_Page.

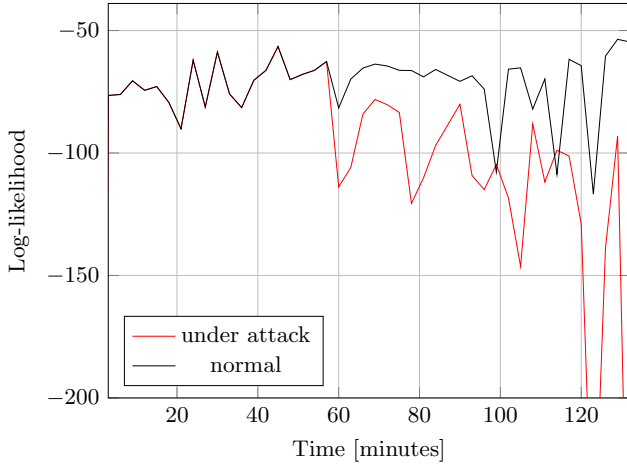


Fig. 4. Likelihood values output by the Gaussian-process based model.

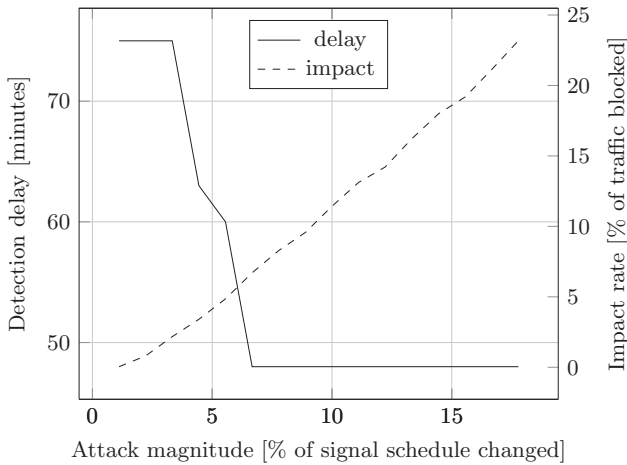


Fig. 5. Impact and detection delay for various attacks.

detector, we compare a wide range of attacks, from stealthy ones that change control only slightly to non-stealthy ones that change control fundamentally. To maximize the advantage of stealthiness, we consider a low detection threshold, which allows attacks to remain undetected for long periods of time. In particular, for this experiment, we assume a detector window of $W = 48$ minutes and a log-likelihood threshold of $\ln \tau = -112.58$, which result in zero false-positive rate for the one-month test interval (see the discussion of Fig. 3).

Fig. 5 shows detection delay and impact for attacks of various magnitudes. The impact of an attack is measured as the fraction of traffic that is “blocked” by the attack, i.e., the decrease in the number of vehicles passing through the intersection compared to the normal traffic-signal schedule; the magnitude of an attack is the fraction of the traffic-signal schedule that is modified by the attacker. The figure shows that attacks with higher magnitude may be less stealthy (i.e., detected earlier), but they cause much more significant impact. In fact, the total impact of attacks, measured as the number of vehicles that could not pass through the intersection due to the attack until its detection, is a strictly increasing function of the attack magnitude. This means that our detector can make stealthy attacks essentially pointless in this example.

5.2. Multi-stage security game strategies

Next, we provide numerical results on the algorithms that we proposed in Section 3.2 for finding strategies in practice.

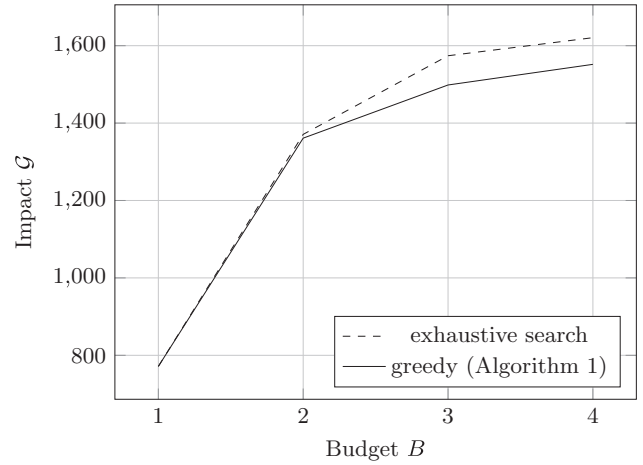


Fig. 6. Impact resulting from attacks found by exhaustive and greedy searches.

We first compare the proposed greedy heuristic for finding attacks (Algorithm 1) to an exhaustive search, and then study the metaheuristic search algorithm for finding detector configurations (Algorithm 2).

5.2.1. Setup

To provide meaningful numerical results, we have to evaluate our algorithms on a large number of transportation networks. Every point plotted in the figures of this subsection represents a mean value computed over a large number of random networks with the same parameters. To obtain these networks, we use the Grid model with Random Edges (GRE) to generate random network topologies (Peng et al., 2012), which closely resemble real-world transportation networks. For a detailed description of this model, we refer the reader to Peng et al. (2014, 2012).

We set both the width and height of the generated grids to be 4, and let the bottom-left corner be a source and the upper-right corner be a sink. For the parameters controlling the randomness of the generation, we use the values from (Peng et al., 2012), which were derived from measurements on actual road networks from the USA. We let the inflow at the source cell be $d^0 = 8$, $d^1 = 12$, $d^2 = 8$, and $d^t = 0$ for $t \geq 3$. For every other cell i , we let the parameters be $Q_i = 6$, $\delta_i = 1.0$, and $N_i = 10$. Finally, we let every merging cell be a signalized intersection, and optimize the inflow proportions for every intersection using a linear program.

We assume that there is an anomaly detector deployed in each intersection (i.e., $\mathcal{I}_D = \mathcal{I}$). We also assume that every one of these detectors exhibits the false-positive rate and detection delay characteristics observed in Section 5.1. In other words, for each intersection, the defender chooses one of the Pareto optimal configurations that were identified in the experiments of Section 5.1 by choosing a false-positive rate D_i ; the delay of this detector is then determined by the magnitude of the attack against the corresponding intersection. Finally, we assume that the attack is detected as soon as one detector raises an alarm, that attack mitigation takes $\Delta_M = 20$ minutes, and that congestion levels T , T_A , and T_M are measured in travel time.

5.2.2. Attacks

We begin by comparing the greedy attack heuristic to an exhaustive search. To perform an exhaustive search, we quantize the space of possible schedules for each intersection, so that we have a finite and discrete search space. For this experiment, we assume that the defender uses a detector configuration that sets the false-positive rate of every intersection $i \in \mathcal{I}_D$ to $D_i = 1$.

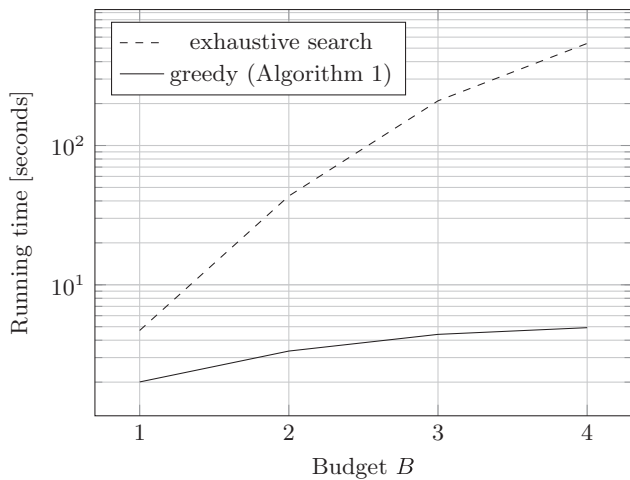


Fig. 7. Running time of exhaustive and greedy searches.

Fig. 6 shows the impact of attacks found by exhaustive and greedy (Algorithm 1) searches for various budget values. The vertical axis shows the total impact \mathcal{G} of the attacks, which includes impact that was caused both before and after detection. The figure shows that the attacks found by the greedy search are very close to the ones found by the exhaustive search in terms of total impact, with the largest difference being 5%.

Fig. 7 compares the greedy heuristic (Algorithm 1) to the exhaustive search in terms of running time. Note that we used fairly small problem instances for our experiments in order to be able to apply the algorithms to a large number of networks. We observe that the running time of the greedy heuristic is much lower than that of the exhaustive search, and it grows slower as the attacker's budget increases.

5.2.3. Detector configuration

Next, we evaluate the metaheuristic search algorithm for finding detector configurations. We compare our strategic configurations to a non-strategic baseline represented by uniform configurations, which assign the same false-positive rate to all detectors. We find quasi-optimal uniform configurations using the same search algorithm, but restricting the search space to a single scalar value, which is used for all detectors. For these experiments, we let the attacker's budget be enough to compromise $B = 2$ intersections; we assume that the attacker always mounts a best-response attack; and we let the unit cost of false positives be equal to $C = 10$.

Fig. 8 shows the defender's total loss—which includes both the cost of investigating false alarms and the total impact of the attack—with strategic and uniform detector configurations. The horizontal axis shows the number of iterations k_{\max} for which the search algorithms was run. We can learn two important lessons from this figure. First, strategic thresholds result in much lower losses than uniform ones, which suggests that game-theoretic optimization can have a significant practical impact. Second, losses decrease rapidly in the first 100 or 500 hundred iterations, but they do not decrease further even after a significant number of additional iterations,⁷ which suggests that the search algorithm is a very efficient practical approach for finding near optimal detector configurations.

Fig. 9 shows the cost of false positives and the total impact of attacks with strategic and uniform detector configurations found by the search algorithm. We observe that strategic detector configurations




Fig. 8. Defender's loss resulting from uniform and strategic detector configurations output by the metaheuristic search algorithm.

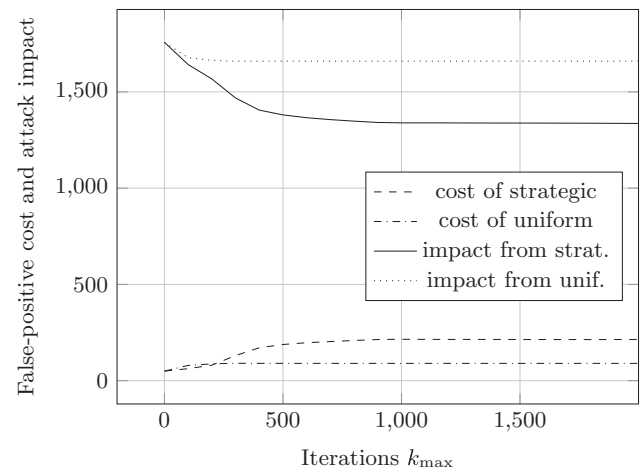


Fig. 9. Defender's false-positive cost and attack impact resulting from uniform and strategic configurations output by the metaheuristic search algorithm.

urations may result in slightly more false-positive errors, but they can significantly decrease the impact of attacks.

6. Related work

In this section, we briefly survey related work on the vulnerability of transportation networks, the optimal configuration of attack detectors, and game theory for security of cyber-physical systems.

6.1. Vulnerability of transportation networks

We first give a brief overview of the related work on the vulnerability of transportation networks. A number of research efforts have studied the vulnerability of transportation networks to natural disasters and attacks. However, to the best of our knowledge, our work is the first one to consider traffic-signal tampering attacks against general transportation networks.

In a closely related work, Reilly et al. consider the vulnerability of freeway control systems to attacks on the sensing and control infrastructure (Reilly et al., 2014). They present an in-depth analysis on the takeover of a series of onramp-metering traffic lights using a methodology based on finite-horizon optimal control techniques and multi-objective optimization.

⁷ We actually run the search with $k_{\max} = 100,000$ iterations, but plot only the first 2000 for clarity, since losses do not decrease significantly after 2000 iterations.

Prior work has studied the impact of other disruptive events as well. Sullivan et al. study short-term disruptive events, such as partial flooding, and propose an approach that employs various link-based capacity-disruption values (Sullivan et al., 2010). The proposed approach can be used to identify and rank the most critical links and to quantify transportation network robustness (i.e., inverse vulnerability). Jenelius and Mattsson introduce an approach for systematically analyzing the robustness of road networks to disruptions affecting extended areas, such as floods and heavy snowfall (Jenelius and Mattsson, 2012). Their methodology is based on covering the area of interest with grids of uniformly shaped and sized cells, where each cell represents the extent of an event. The authors apply their approach to the Swedish road network, and find that the impact of area-covering disruptions are largely determined by the internal, outbound, and inbound travel demands of the affected area itself.

In addition to assessing vulnerability, prior work has also considered the problems of identifying critical links and studying other aspects of vulnerability. Scott et al. propose a comprehensive, system-wide approach for identifying critical links and evaluating network performance (Scott et al., 2006). Using three hypothetical networks, the authors demonstrate that their approach yields different highway planning solutions than traditional approaches, which rely on volume/capacity ratios to identify congested or critical links. Jenelius proposes a methodology for vulnerability analysis of road networks and considers the impact of road-link closures (Jenelius, 2010). The author considers different aspects of vulnerability, and explores the dichotomy between system-wide efficiency and user equity.

Prior work has also considered game-theoretic models of attacks against transportation. Alpcan and Buchegger investigate the resilience aspects of vehicular networks using a game-theoretic model, in which defensive measures are optimized with respect to threats posed by intentional attacks (Alpcan and Buchegger, 2011). The game is formulated in an abstract manner, based on centrality values computed by mapping the centrality values of the car communication network onto the road topology. The authors consider multiple formulations based on varying assumptions on the players' information, and evaluate their models using numerical examples. Bell introduces a two-player non-cooperative game between a network user, who seeks to minimize expected travel cost, and an adversary, who chooses link performance scenarios to maximize the travel cost (Bell, 2000; Bell et al., 2008). The Nash equilibrium of this game can be used to measure network performance when users are pessimistic and, hence, may be used for cautious network design. Wu and Amin study normal-form and sequential attacker-defender games over transportation networks to understand how a defender should prioritize its investment in securing a set of facilities (Wu and Amin, 2018).

6.2. Configuration of detectors

The problem of configuring the sensitivity of intrusion detection systems in the presence of strategic attackers has been studied in a variety of different ways in the academic literature (Laszka et al., 2016a). For example, Alpcan and Basar study distributed intrusion detection in access control systems as a security game between an attacker and an IDS, using a model that captures the imperfect flow of information from the attacker to the IDS through a network (Alpcan and Basar, 2003; Alpcan and Başar, 2004). The authors investigate the existence of a unique Nash equilibrium and best-response strategies under specific cost functions, and analyze long-term interactions using repeated games and a dynamic model. As another example, Dritsoula et al. consider the problem of setting a threshold for classifying an attacker into one of two categories, spammer and spy, based on its intrusion at-

tempts (Dritsoula et al., 2012). They give a characterization of the Nash equilibria in mixed strategies, and show that the equilibria can be computed in polynomial time. More recently, Lisý et al. study randomized detection thresholds using a general model of adversarial classification, which can be applied to e-mail filtering, intrusion detection, steganalysis, etc. (Lisý et al., 2014). The authors analyze both Nash and Stackelberg equilibria based on the true-positive to false-positive curve of the classifier, and find that randomizing the detection threshold may force a strategic attacker to design less efficient attacks. Finally, Zhu and Basar study the problem of optimal signature-based IDS configuration under resource constraints (Zhu and Başar, 2011).

The strategic configuration of the sensitivity of e-mail filtering against spear-phishing and other malicious e-mail is also closely related to the problem considered in this paper. Laszka et al. study a single defender who has to protect multiple users against targeted and non-targeted malicious e-mail (Laszka et al., 2015). The authors focus on characterizing and computing optimal filtering thresholds, and they use numerical results to demonstrate that optimal thresholds can lead to substantially lower losses than naïve ones. Zhao et al. study a variant of the previous model: they assume that the attacker can mount an arbitrary number of costly spear-phishing attacks in order to learn a secret, which is known only by a subset of the users (Zhao et al., 2015; 2016). They also focus on the computational aspects of finding optimal filtering thresholds; however, their variant of the model does not capture non-targeted malicious e-mails, such as spam.

6.3. Game theory for security of cyber-physical systems

Beyond the configuration of detectors, prior efforts have also used game-theory to study a variety of other security problems in cyber-physical systems. For instance, Zhu and Basar introduce a game-theoretic framework for resilient control design and studying the trade-off between robustness, security, and resilience (Zhu and Basar, 2015). They employ a hybrid model, which integrates a discrete-time Markov model that captures the evolution of cyberstates with continuous-time dynamics that capture the underlying controlled physical process. Backhaus et al. consider the problem of designing attack-resilient power grids and control systems (Backhaus et al., 2013). They use game theory to model the conflict between a cyber-physical intruder and a system operator, and use simulation results to assess design options. Pawlick et al. consider Advanced Persistent Threats (APTs) against cloud-controlled cyber-physical systems and design a framework that specifies when a device should trust commands from the cloud (Pawlick et al., 2015). They model this scenario as a three player game between the cloud administrator, the attacker, and the device by combining the FlipIt game (Laszka et al., 2014; Van Dijk et al., 2013) with a signaling game. Li et al. consider jamming attacks against remote state estimation (Li et al., 2015). In particular, they assume that a sensor and an estimate have to communicate over a wireless, formulate a game-theoretic model, and provide Nash equilibrium strategies.

7. Conclusion

As traffic-control devices in practice evolve into complex networks of smart devices, the risks posed to transportation networks by cyber-attacks increases. Thus, it is imperative for traffic-network operators to be prepared to detect and mitigate attacks against traffic control. To provide theoretical foundations for planning and implementing countermeasures, we introduced a game-theoretic model of cyber-attacks against traffic control. Our security game model consists of three stages: defender configuring

detectors, attacker mounting a tampering attack against traffic signals, and defender mitigating the attack. Since mitigation—in our model—means adapting the schedules of some traffic signals given the schedules of other signals (set by the adversary in the previous stage), it is equivalent to optimizing traffic control in a non-adversarial setting, which has been studied in prior work. In light of this, we focused on the computational problem of finding optimal actions in the first two stages. We showed that this is a computationally hard problem, which prompted us to propose efficient heuristic algorithms.

Using numerical results, we demonstrated that the proposed algorithms are practical. In particular, we first showed that the greedy algorithm for attackers is close to optimal and computationally very efficient. Second, we showed that the metaheuristic search algorithm for detector configuration is effective, and it can significantly decrease losses compared to non-strategic detector configuration. We also introduced and studied a Gaussian-process based traffic-anomaly detector, which we showed to be very effective at detecting tampering attacks against traffic signals.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported in part by the National Science Foundation under CAREER Grant IIS-1905558.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.cose.2019.101576

References

- Alpcan, T., Basar, T., 2003. A game theoretic approach to decision and analysis in network intrusion detection. In: Proceedings of the 42nd IEEE Conference on Decision and Control (CDC), 3. IEEE, pp. 2595–2600.
- Alpcan, T., Başar, T., 2004. A game theoretic analysis of intrusion detection in access control systems. In: Proceedings of the 43rd IEEE Conference on Decision and Control (CDC), 2. IEEE, pp. 1568–1573.
- Alpcan, T., Buchegger, S., 2011. Security games for vehicular networks. *IEEE Trans. Mob. Comput.* 10 (2), 280–290.
- Backhaus, S., Bent, R., Bono, J., Lee, R., Tracey, B., Wolpert, D., Xie, D., Yildiz, Y., 2013. Cyber-physical security: a game theory model of humans interacting over control systems. *IEEE Trans. Smart Grid* 4 (4), 2320–2327.
- Bazi, Y., Alajlan, N., Melgani, F., 2012. Improved estimation of water chlorophyll concentration with semisupervised Gaussian process regression. *IEEE Trans. Geosci. Remote Sensing* 50 (7), 2733–2743.
- Behrisch, M., Bieker, L., Erdmann, J., Krajzewicz, D., 2011. SUMO – Simulation of Urban MObility: an overview. In: Proceedings of the 3rd International Conference on Advances in System Simulation (SIMUL), pp. 63–68.
- Bell, M.G., 2000. A game theory approach to measuring the performance reliability of transport networks. *Transp. Res. Part B* 34 (6), 533–545.
- Bell, M.G., Kanturska, U., Schmöcker, J.-D., Fonzone, A., 2008. Attacker-defender models and road network vulnerability. *Philos. Trans. R. Soc. Lond. A* 366 (1872), 1893–1906.
- Bernstein, S., Blankstein, A., 2007. Key signals targeted, officials say. *Los Angeles Times*.
- Blanchet, J., Davison, A.C., et al., 2011. Spatial modeling of extreme snow depth. *Ann. Appl. Stat.* 5 (3), 1699–1725.
- Chen, J., Low, K.H., Tan, C.K.-Y., Oran, A., Jaillet, P., Dolan, J.M., Sukhatme, G.S., 2012. Decentralized data fusion and active sensing with mobile sensors for modeling and predicting spatiotemporal traffic phenomena. In: Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI), pp. 163–173.
- Daganzo, C.F., 1994. The cell transmission model: a dynamic representation of highway traffic consistent with the hydrodynamic theory. *Transp. Res. Part B* 28 (4), 269–287.
- Daganzo, C.F., 1995. The cell transmission model, part II: network traffic. *Transportation Research Part B: Methodological* 29 (2), 79–93.
- Dritsoula, L., Loiseau, P., Musacchio, J., 2012. Computing the Nash equilibria of intruder classification games. In: Proceedings of the 3rd International Conference on Decision and Game Theory for Security (GameSec). Springer, pp. 78–97.
- Gelman, A., Stern, H.S., Carlin, J.B., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. *Bayesian Data Analysis*, 3 Chapman and Hall/CRC.
- Ghena, B., Beyer, W., Hillaker, A., Pevarnek, J., Halderman, J.A., 2014. Green lights forever: analyzing the security of traffic infrastructure. In: Proceedings of the 8th USENIX Workshop on Offensive Technologies (WOOT), pp. 1–10.
- Jenelius, E., 2010. Large-Scale Road Network Vulnerability Analysis. KTH Ph.D. thesis.
- Jenelius, E., Mattsson, L.-G., 2012. Road network vulnerability analysis of area-covering disruptions: a grid-based approach with case study. *Transp. Res. Part A* 46 (5), 746–760.
- Kou, P., Gao, F., Guan, X., 2013. Sparse online warped Gaussian process for wind power probabilistic forecasting. *Appl. Energy* 108, 410–428.
- Krajzewicz, D., Hertkorn, G., Rössel, C., Wagner, P., 2002. SUMO (Simulation of Urban MObility): an open-source traffic simulation. In: Proceedings of the 4th Middle East Symposium on Simulation and Modelling (MESM), pp. 183–187.
- Laszka, A., Abbas, W., Sastry, S.S., Vorobeychik, Y., Koutsoukos, X., 2016. Optimal thresholds for intrusion detection systems. In: Proceedings of the 3rd Annual Symposium and Bootcamp on the Science of Security (HotSoS), pp. 72–81.
- Laszka, A., Horvath, G., Felegyhazi, M., Buttyan, L., 2014. FlipThem: modeling targeted attacks with FlipIt for multiple resources. In: Proceedings of the 5th Conference on Decision and Game Theory for Security (GameSec), pp. 175–194.
- Laszka, A., Koutsoukos, X., Vorobeychik, Y., 2019. Towards high-resolution multi-stage security games. In: *Proactive and Dynamic Network Defense*. Springer, pp. 139–161.
- Laszka, A., Potteiger, B., Vorobeychik, Y., Amin, S., Koutsoukos, X., 2016. Vulnerability of transportation networks to traffic-signal tampering. In: Proceedings of the 7th ACM/IEEE International Conference on Cyber-Physical Systems (ICCPs), pp. 1–10.
- Laszka, A., Vorobeychik, Y., Koutsoukos, X., 2015. Optimal personalized filtering against spear-phishing attacks. In: Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), pp. 958–964.
- Li, Y., Shi, L., Cheng, P., Chen, J., Quevedo, D.E., 2015. Jamming attacks on remote state estimation in cyber-physical systems: a game-theoretic approach. *IEEE Trans. Automat. Control* 60 (10), 2831–2836.
- Lisý, V., Kessl, R., Pevný, T., 2014. Randomized operating point selection in adversarial classification. In: Proceedings of the 2014 European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD), Part II. Springer, pp. 240–255.
- Nevat, I., Peters, G.W., Collings, I.B., 2012. Location-aware cooperative spectrum sensing via Gaussian processes. In: Proc. of the 13th Australian Comm. Theory Work. (AusCTW). IEEE, pp. 19–24.
- Pawlick, J., Farhang, S., Zhu, Q., 2015. Flip the cloud: cyber-physical signaling games in the presence of advanced persistent threats. In: Proceedings of the 6th Conference on Decision and Game Theory for Security (GameSec). Springer, pp. 289–308.
- Peng, W., Dong, G., Yang, K., Su, J., 2014. A random road network model and its effects on topological characteristics of mobile delay-tolerant networks. *IEEE Trans. Mob. Comput.* 13 (12), 2706–2718.
- Peng, W., Dong, G., Yang, K., Su, J., Wu, J., 2012. A random road network model for mobility modeling in mobile delay-tolerant networks. In: Proceedings of the 8th International Conference on Mobile Ad-hoc and Sensor Networks (MSN). IEEE, pp. 140–146.
- Rasmussen, C.E., Williams, C.K.I., 2006. *Gaussian Processes for Machine Learning*. MIT Press.
- Reilly, J., Martin, S., Payer, M., Payer, M., 2014. On cybersecurity of freeway control systems: analysis of coordinated ramp metering attacks. *Transp. Res., Part B*.
- Scott, D.M., Novak, D.C., Aultman-Hall, L., Guo, F., 2006. Network robustness index: a new method for identifying critical links and evaluating the performance of transportation networks. *J. Transport Geograph.* 14 (3), 215–227.
- Sullivan, J., Novak, D., Aultman-Hall, L., Scott, D.M., 2010. Identifying critical road segments and measuring system-wide robustness in transportation networks with isolating links: a link-based capacity-reduction approach. *Transp. Res. Part A* 44 (5), 323–336.
- Van Dijk, M., Juels, A., Oprea, A., Rivest, R.L., 2013. FlipIt: the game of stealthy takeover. *J. Cryptol.* 26 (4), 655–713.
- Wu, M., Amin, S., 2018. Security of transportation networks: modeling attacker-defender interaction. <https://arxiv.org/abs/1804.00391v1>.
- Xie, Y., Zhao, K., Sun, Y., Chen, D., 2010. Gaussian processes for short-term traffic volume forecasting. *Transp. Res. Rec.* 2165 (1), 69–78.
- Zhao, M., An, B., Kiekintveld, C., 2015. An initial study on personalized filtering thresholds in defending sequential spear phishing attacks. In: Proceedings of the 2015 IJCAI Workshop on Behavioral, Economic and Computational Intelligence for Security, pp. 1–9.
- Zhao, M., An, B., Kiekintveld, C., 2016. Optimizing personalized email filtering thresholds to mitigate sequential spear phishing attacks. In: Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI), pp. 658–664.
- Zhu, Q., Başar, T., 2011. Indices of power in optimal IDS default configuration: theory and examples. In: Proceedings of the 2nd International Conference on Decision and Game Theory for Security (GameSec). Springer, pp. 7–21.
- Zhu, Q., Basar, T., 2015. Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: games-in-games principle for

optimal cross-layer resilient control systems. *IEEE Control Syst. Mag.* 35 (1), 46–65.

Ziliaskopoulos, A.K., 2000. A linear programming model for the single destination system optimum dynamic traffic assignment problem. *Transp. Sci.* 34 (1), 37–49.

Aron Laszka is an Assistant Professor in the Department of Computer Science at the University of Houston, TX, USA. Previously, he was a Research Assistant Professor in the Department of Electrical Engineering and Computer Science at Vanderbilt University, and a Postdoctoral Scholar at the University of California, Berkeley between 2015 and 2016. He graduated summa cum laude with a Ph.D. in Computer Science from the Budapest University of Technology and Economics in 2014. In 2013, he was a Visiting Research Scholar at the Pennsylvania State University. His research interests broadly revolve around the security and resilience of cyber-physical systems and the Internet-of- Things, the economics of security, and game-theoretic modeling of security problems.

Waseem Abbas is an Assistant Professor in the Electrical Engineering Department at the InformationTechnology University, Lahore, Pakistan. Previously, he was a post-doctoral research scholar in the Department of Electrical Engineering and Computer Science at Vanderbilt University, Nashville, TN. He received Ph.D. (2013) and M.Sc. (2010) degrees, both in Electrical and Computer Engineering, from Georgia Institute of Technology, Atlanta, GA. He was a Fulbright scholar from 2009 till 2013. His research interests are in the area of network control systems, graph-theoretic methods in complex networks, social networks, and resilience of cyber-physical systems.

Yevgeniy Vorobeychik is an Associate Professor of Computer Science and Engineering at Washington University in Saint Louis. Previously, he was an Assistant Profes-

sor of Computer Science at Vanderbilt University. Between 2008 and 2010 he was a post-doctoral research associate at the University of Pennsylvania Computer and Information Science department. He received Ph.D. (2008) and M.S.E. (2004) degrees in Computer Science and Engineering from the University of Michigan, and a B.S. degree in Computer Engineering from Northwestern University. His work focuses on game theoretic modeling of security and privacy, adversarial machine learning, algorithmic and behavioral game theory and incentive design, optimization, agent-based modeling, complex systems, network science, and epidemic control. Dr. Vorobeychik received an NSF CAREER award in 2017, and was invited to give an IJCAI-16 early career spotlight talk. He also received several Best Paper awards, including one of 2017 Best Papers in Health Informatics. He was nominated for the 2008 ACM Doctoral Dissertation Award and received honorable mention for the 2008 IFAAMAS Distinguished Dissertation Award.

Xenofon Koutsoukos received the Ph.D. degree in electrical engineering from the University of NotreDame, Notre Dame, IN, USA, in 2000. He is a Professor with the Department of Electrical Engineering and Computer Science and a Senior Research Scientist with the Institute for Software Integrated Systems (ISIS), Vanderbilt University, Nashville, TN, USA. He was a Member of Research Staff at the Xerox Palo Alto Research Center (PARC) (2000–2002), working in the embedded collaborative computing area. His research work is in the area of cyber-physical systems with emphasis on formal methods, data-driven methods, distributed algorithms, security and resilience, diagnosis and fault tolerance, and adaptive resource management. He has published more than 250 journal and conference papers and he is co-inventor of four US patents. Prof. Koutsoukos was the recipient of the NSF Career Award in 2004, the Excellence in Teaching Award in 2009 from the Vanderbilt University School of Engineering, and the 2011 NASA Aeronautics Research Mission Directorate (ARMD) Associate Administrator (AA) Award in Technology and Innovation. He was named a Fellow of the IEEE for the class of 2018.