Transportation Research Part A 120 (2019) 58–70

Contents lists available at ScienceDirect

Transportation Research Part A

journal homepage: www.elsevier.com/locate/tra

Implications of traffic signal cybersecurity on potential deliberate traffic disruptions

Kenneth A. Perrine[a,*], Michael W. Levin[b], Cesar N. Yahia[c], Melissa Duell[d], Stephen D. Boyles[c]

[a] Network Modeling Center, Center for Transportation Research, Cockrell School of Engineering, The University of Texas at Austin, 3925 W. Braker Ln., Stop D9300, Austin, TX 78759, United States
[b] Department of Civil, Environmental, and Geo-Engineering, University of Minnesota, 140 Civil Engineering Building, 500 Pillsbury Drive SE, Minneapolis, MN 55455, United States
[c] Department of Civil, Architectural and Environmental Engineering, The University of Texas at Austin, 301 E. Dean Keeton St. Stop C1761, Austin, TX 78712-1172, United States
[d] School of Civil and Environmental Engineering, University of New South Wales, Sydney, NSW 2052, Australia

ARTICLE INFO

Keywords:
Cyber security
Traffic operations
Hacking
Risk
Vulnerability
Attack

ABSTRACT

Traffic control systems, including signal controllers, sensors, and centralized coordination software, all have the capacity to be vulnerable to malicious attacks. Although several studies on outages, attacks, and cybersecurity have been conducted in the literature, the effects of district-wide attacks on signals have not been specifically studied in-depth. There is a need for risk assessments to be conducted to establish resilient policies within traffic operations agencies. A key factor in assessing risk is in gaining an idea of the hypothetical impact of an outage. In this preliminary study, a dynamic traffic assignment network is used to model a central business district, where traffic signal-controlled intersections are cyberattacked and selectively disabled (effectively replaced with four-way stops). In one scenario, total delay is multiplied 4.3 times when 26 signals are chosen and disabled according to maximum, decreasing intersection traffic volume. In scenarios where the attacker prioritizes the selection of signals by maximizing the number of travelers affected, 7 signals are needed to exert the same impact.

1. Introduction

As traffic signal control technologies are improved and field equipment is updated, overall infrastructure connectivity increases. States and municipalities continue to install and maintain regional wired and wireless networks to enhance traffic management. One consequence of enhanced connectivity is increased risk of malicious attacks. Known breaches in industrial control systems has risen as these systems have become more technologically developed (Byres and Lowe, 2004), and has influenced cybersecurity policies within respective organizations.

Traffic systems, including signal controllers, sensors, centralized coordination software, variable message signs, and networking devices all have the capacity to be vulnerable to attacks because of misconfigurations, lack of security features, and system failures. Attacks on signal operation (Ghena et al., 2014), sensors (Infosec Institute, 2014), and variable message signs (Sawin, 2010) have

* Corresponding author.
E-mail addresses: kperrine@utexas.edu (K.A. Perrine), mlevin@umn.edu (M.W. Levin), cesaryahia@utexas.edu (C.N. Yahia), sboyles@mail.utexas.edu (S.D. Boyles).

https://doi.org/10.1016/j.tra.2018.12.009
Received 2 December 2015; Received in revised form 28 November 2018; Accepted 4 December 2018
Available online 20 December 2018
0965-8564/

already happened, and many other potential vulnerabilities such as signal controller system access due to default passwords (Econolite, 2014), susceptibilities to denial of service attacks (Hoo, 2000), computer virus and malware infestations (Byres and Lowe, 2004), etc. exist.

Study of traffic systems vulnerability—a largely untapped area—falls within a broader field of industrial control system vulnerability that includes other infrastructure areas as power distribution, water supply, etc. Traffic systems are unique in that performance largely depends upon individual people that directly affect the system operation (Wang et al., 2015). Each analyzed type of disruption and mode of transport further carries its own unique challenges and implications (Faturechi and Miller-Hooks, 2014). One of the most important analyses is in the route choice behavior of travelers in the face of disruptions (Sullivan et al., 2009). Previous literature on network vulnerability has focused on link- or road-specific impacts (Bell and Cassir, 2002; Sullivan et al., 2010). Because traffic signals have been assumed to be secure, the effects of a malicious attack on signals have not yet been fully studied except in small-scale cases (Feng et al., 2018). Consequently, the district-wide study within this paper is preliminary in nature and motivated by an immediate need for increased awareness, better tools, and improved policies for traffic operations personnel and computer network engineers.

Because the notion of ultimate, total security comes at a conceptually prohibitive cost (Byres and Lowe, 2004), the response within a traffic management organization in addressing known and unknown vulnerabilities is often limited. In facilitating better cybersecurity, the organization must assess the risks of threats, prioritize known vulnerabilities, and choose how much to invest given limited resources (Hoo, 2000). A useful practice in assessing risk is to estimate the severity of consequences that result from probable attacks.

The contributions of this paper are as follows. First, in Section 2 we review demonstrated and potential cybersecurity vulnerabilities that affect or could affect traffic signal controllers. Next, in Section 3 we present a methodology for estimating the impact of traffic signal controller outages within a city, under the simplifying assumption that travelers continue to attempt to carry out daily activities while disruptions are underway. We test this methodology on the downtown Austin, Texas city network on several likely hypothetical scenarios. After discussing further possible cybersecurity mitigations and policy in Section 4, the paper concludes in Section 5.

2. Background

We now review possible traffic signal control system vulnerabilities found in the field, and briefly describe the role of risk assessment for new and pre-existing systems.

2.1. Potential vulnerabilities

A variety of potential vulnerabilities and attack surfaces in traffic signal control equipment exist, each exploitable through malicious intents and hacking techniques. For example, one of the most common and low-level vulnerabilities is the presence of default usernames and passwords on standards-abiding, legacy traffic signal controllers. A possible hacking technique may involve gaining physical access to the regional wired, optical, or radio frequency (RF) computer network, and executing a cyberattack. Certain vulnerabilities may allow one or more controllers to be rendered inoperable. Even though the conflict monitors at each affected intersection would put each signal into flashing red operation (and possibly yellow depending on conflict monitor configuration) and maintain minimally safe traffic operations, the widespread impacts of this flashing operation may be severe until field technicians can restore operations.

A similarly possible attack could involve keeping traffic controllers operational, but forcing a green movement for an infinite amount of time (Ghena et al., 2014). In this case, traffic safety may be worse because of the possibility of drivers disregarding red indications after waiting for a long time, and dangerously entering conflicting traffic. In this scenario, after a certain amount of response time, city officials may choose to deploy police traffic directors and physically revert to flashing red operation to address the infinite red/green light problem.

When considering these examples, many traffic operations personnel are confident about the physical barriers imposed by traffic signal control cabinets in thwarting such attacks. Apart from this physical defense, few other barriers may stand in the way of hackers. Importantly, manufacturers of signal control equipment often justify the presence of vulnerabilities by stating that customers and standards desire such vulnerabilities because of ease of use or other reasons (Cerrudo, 2014). Such manufactures may openly leave the responsibility of implementing cybersecurity to the end user (Bayless et al., 2014). Even so, many jurisdictions historically do not dedicate attention to building up additional defenses (Institute of Transportation Engineers, 2003).

Although there is some awareness of threat possibilities, the relevance of the threats are often underestimated by untrained personnel (Fletcher, 2014). As a result, organizations often set up security schemes to address a set of bare minimum requirements, potentially leaving open many other security holes (Lampson, 2004). As Hoo (2000) argues, “Ignorance is a self-reinforcing problem since organizations are reluctant to act on security concerns unless a real problem has been proven.” Likewise, proper attention and funding for improved cybersecurity may not be present until a major “security catastrophe” happens (Lampson, 2004).

Another source of vulnerability lies within the processes of businesses that create traffic control products (Bayless et al., 2014). It has been observed that rushed project schedules are a severe reality for companies that must expedite new products to market in order to gain a competitive advantage. This is often paired with a “fix-it later mentality” when concerning cybersecurity, where “later” may in fact mean “never”.

In the examples given above, physical access to the regional computer network can be achieved from a compromised traffic signal

control cabinet, or along places where regional network fiber or cable are routed. Alternatively, RF transmitters can be used on vulnerable wireless networks (Cerrudo, 2014). Signal controllers have already experienced an attack by insiders with privileged access as in the 2007 Los Angeles incident (Bernstein and Blackstein, 2007). Two disgruntled city employees disabled signals at four busy intersections for several days. However, attention should also be given to threats originating outside of an organization, as most attacks on industrial control systems within the new millennium have historically come from outsiders (Byres and Lowe, 2004).

The question of how much damage one attacker can inflict may be affected by how much knowledge the attacker has on the inner workings of a system. Although it is possible for security practices to involve the protection of sensitive information, experience from industrial control fields has shown that the effectiveness of "security through obscurity" is diminishing (Byres and Lowe, 2004). Many avenues have emerged online for exploitable information to be shared. The proprietary network control protocol for a popular traffic signal controller has been successfully reverse-engineered (Goodspeed, 2008). Vulnerabilities in variable message signs have been shared online and exploited (Sawin, 2010). RF communications to wireless detectors have been compromised and documented in an online blog (Infosec Institute, 2014). Other wireless exploits have allowed for success in tampering with traffic signal operations (Ghena et al., 2014).

When assessing security vulnerabilities and determining potential fixes, the overall system can be divided among three levels (Lampson, 2004). First, *network security* includes physical barriers to hardware and software barriers such as firewalls or virtual private network (VPN) devices. Second, *operating system (OS) security* pertains to system-level access to individual controllers or central computers, including user authentication. Third, *application-level* security pertains to security features that are specific to a software solution, such as a central software application that uses domain-specific communication protocols to control and monitor traffic controllers in the field. Traditional traffic operations practices and solutions have tended to focus on a minimal degree of security at the network level, with less flexible availability of operating system security options, and even fewer security features at the application level.

Another facet of security does not involve malicious attacks, but instead relates to operator and equipment failure (Sawin, 2010). Intuitively, the effects of operator error or equipment failure are minimized when adequate security features and practices are in place. An example of inadequate application-level guarding against system failure was observed in the 2009 Montgomery County, Maryland incident, where coordination among 750 signals was lost for over a day, severely impacting the commutes of thousands (Halsey, 2009).

2.2. Risk assessment

We now briefly introduce the role of risk assessment for planning new and improving already-existing traffic signal systems. Our subsequent design of the experimental scenarios is an initial effort in facilitating an approximate, but useful risk assessment that is intended to be applicable to a variety of major urban locations.

Several cybersecurity guides indicate that risk assessment is a primary goal that should be accomplished before other steps are executed, such as drafting incident response plans (U.S. Department of Homeland Security, 2013; Frazier et al., 2009). A simple model of risk is (Frazier et al., 2009):

$$\text{Risk} = \text{Impact} \times \text{Likelihood}$$

In some literature, "impact" may also be described as "consequence". Often impact can be based upon an estimate of financial loss, but it can also include other things such as health effects or environmental consequences (Byres and Lowe, 2004). Estimating likelihood of a successful attack is said to be far more difficult as this can be a function of perceived threat, known vulnerabilities, and target attractiveness. Little historic data are available for assisting in making reasonable estimates on these factors. Furthermore, "most organizations are highly reluctant to report security incidents as they are viewed as potential embarrassments." (Byres and Lowe, 2004) Despite complications, risk assessments help in answering questions on how much risk is acceptable in a given traffic operations system. In the cybersecurity guides, this also informs how response plans are made and prioritized. For significant work to be done, estimates need not necessarily be precise, but should be reasonably bounded.

Despite efforts in risk estimation, prioritization of response plans, etc. there is a fundamental tension between usability and security (Hoo, 2000). In considering one extreme limit, "we can't afford the infinite cost of perfect security." Byres and Lowe (2004). On the other hand, some degree of planning should happen as reinforced by the saying, "if you fail to plan, then plan to fail" (Frazier et al., 2009). There should be a balance between the impacts of active security practices and the level of security that is perceived as needed (Lampson, 2004). Security that limits legitimate access may receive more attention than "security that keeps the bad guys out", since the latter can go undetected for long periods of time.

The analysis of risk is challenged by lack of good statistics on computer security crimes (Hoo, 2000), and at least as limited in the transportation field. When trying to acquire managerial and political buy-in for cybersecurity, it becomes necessary to quantify estimates of hypothetical losses. In our work, the traffic models on which these estimates are made may not represent all phenomena that occur in reality. Rather, balance must be found between simplicity and accurate portrayal of the real traffic system (Hoo, 2000).

In attempts to assess the risk of vulnerabilities that, say, force affected intersections into flashing red operation, much insight can be gained by creating a set of hypothetical scenarios that characterize the effects of possible attacks, as seen in this research. The measured severity of problems in each scenario can then inform the best types of mitigations. For example, scenario results can highlight the positive and negative aspects of broadly safeguarding selected traffic corridors or urban regions versus finely limiting possible damage to individual intersections. In the end, a security policy that is drafted with the help of the risk assessment defines what "cybersecurity" really means within a given system (Bishop, 2003).

To the best of the authors' knowledge, no immediate examples exist in the literature that attempt to quantify the effects of signal controller attacks and failures on a *district-wide* scale for the purpose of risk assessment. However, related work had been accomplished in analyzing the effects of falsified data on four types of attack surfaces within a six-intersection corridor (Feng et al., 2018). Impact on total travel time is measured as one to four attacks are simulated in various scenarios. Furthermore, (Reilly et al., 2015) hypothesizes possible outcomes of unauthorized ramp meter tampering. In simulation, scenarios are devised at the expense of many simulated travelers that recreate for one vehicle a "VIP lane" (a path of travel on an expressway that is clear of congested traffic) and also a scheme that assists a getaway vehicle in fleeing from a crime scene.

3. Experiment

To motivate greater attention to signal controller security and to provide an example of an approximate model that can assist in assessing risk, we quantify the potential impacts of hacking signals on a dynamic traffic assignment (DTA) model of the downtown Austin city network. Although we focus upon impacts in terms of travel time, other measures such as impact per unit time, safety and environmental impacts may be addressed in future experimental designs. Because of potential severity of traffic signal outages on a city, we look only at signal outages in this study. Other kinds of attacks such as tampering with sensors and changing timing plans deserve future study, even though their effects may be less severe. However, some types of attacks such as conflicting greens or shortened yellow intervals may be sufficiently infeasible over a signal control network because of "hard-wired" safety features (e.g. conflict monitors) found in all signal control systems.

We first describe a model of stop sign-controlled intersections for the cell transmission model (CTM) developed by Daganzo (1994, 1995) based on a previously-calibrated DTA model of reservation-based intersection control (Levin and Boyles, 2015). We use this model to demonstrate the effects of attacks on traffic signals, turning them into stop signs. A DTA model is well-suited to the analysis in this paper because it can simulate network-wide impacts of intersection failures with computational efficiency (Chiu et al., 2011).

3.1. Stop sign model

Our DTA traffic signal model cycles through its phases, assigning saturation flows at each time step proportional to the green time from active phases. This results in capturing both average intersection flow as well as average delays due to traffic signals. The design goals for the stop sign model are similar. We develop a model that attempts to predict both the average intersection capacity as well as the minimum delays due to stopping at the intersection.

Stop signs in DTA are modeled by adapting the reservation-based intersection control developed for autonomous vehicles by Dresner and Stone (2004). Reservations are essentially an evolution of stop signs that use digital communications and intersection agents to reduce safety margins necessary for human drivers. With reservations, the intersection is divided into a grid of tiles. Vehicles reserve the use of tiles at specific times, and the intersection agent prevents two vehicles from using the same tile at the same time. Stop signs are more constrained than reservations because they require vehicles to stop before entering the intersection and provide exclusive access to the entire intersection rather than a limited set of tiles. We model stop signs by adding these constraints to a DTA model of reservations.

Reservations have previously been modeled in DTA through the conflict region model of Levin and Boyles (2015), which was shown to be compatible with the general intersection model requirements of Tampère et al. (2011). For the purposes of this paper, we adapt the conflict region model for stop signs by adding safety margins and stopping delay. This creates two types of additional constraints on intersection flow: (1) reduced capacity across the intersection reflecting that all vehicles start moving from a complete stop; and (2) minimum delay in the last cell of the link due to the vehicle coming to a stop before entering the intersection. We use a single conflict region for the entire intersection to model exclusive access to the intersection. We define a *turning movement* to be a pair of links $(i, j) \in \Gamma^{-1} \times \Gamma$, where Γ^{-1} is the set of incoming links and Γ is the set of outgoing links for the intersection.

3.1.1. Turning movement capacity

Previous work on macroscopic models of stop signs (Wu, 2000; Li et al., 2011; Wu, 2002) used intersection travel time required for each turning movement to estimate capacity for each turning movement. We estimate these times by geometrically estimating the distance traveled for each turning movement, and using the driver acceleration models of Wang et al. (2004) to determine travel time. Wang et al. developed regression models of driver acceleration in the form

$$a = \alpha + \beta v \tag{1}$$

where a is acceleration, v is speed, and α and β are constants. For vehicles going straight, $\alpha = 1.883 \quad \mathrm{m/s^2}$ and $\beta = -0.021 \quad \mathrm{s^{-1}}$. For vehicles making turning maneuvers, $\alpha = 1.646 \quad \mathrm{m/s^2}$ and $\beta = -0.017 \quad \mathrm{s^{-1}}$.

For estimating distance, we distinguish between three types of turning movements: straight, right turns, and left turns. We assume that U-turns are not used in this model because in our DTA model, vehicle route choice is completely determined before vehicles depart and shortest paths are acyclic. Therefore, we do not code U-turns, which simplifies the analysis. (We acknowledge that in reality, frustrated drivers may make U-turns and exhibit other impromptu behaviors in high-delay situations). Since the study network has 546 intersections (with 164 signalized), we use an automatic procedure based on the change in direction a vehicle makes along its turning movement. Let θ_i be the direction of link i. Then the change in direction for turning movement (i, j) is $\Delta\theta_{ij} = \theta_j - \theta_i$. Without loss of generality, let $\Delta\theta_{ij} \in [0, 2\pi]$. If $\Delta\theta_{ij} \leq \frac{\pi}{4}$ or $\Delta\theta_{ij} \geq \frac{7\pi}{4}$, then (i, j) is labeled as a straight movement. If $\frac{\pi}{4} < \Delta\theta_{ij} < \pi$, (i, j)

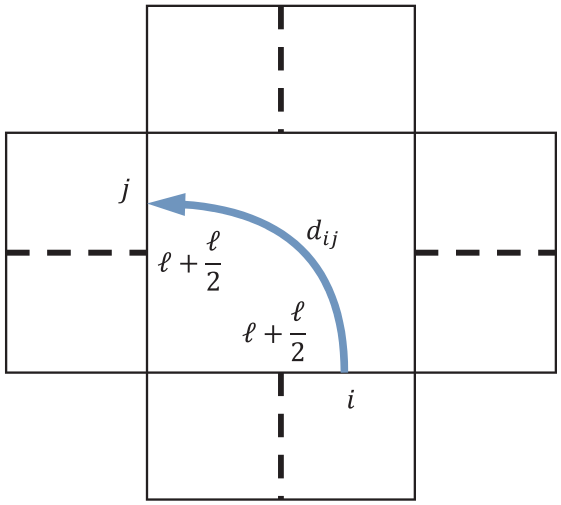

Fig. 1. Stop sign displacement calculation.

is labeled a left turn, and if $\pi < \Delta\theta_{ij} < \frac{7\pi}{4}$, (i, j) is labeled a right turn. Link directions are primarily determined from node coordinates, which implicitly assumes that links are straight. Since right angle turns are most common, we assume that left and right turns are right angles instead of using the estimation of link direction to determine the angle.

Vehicles going straight must cross some l_{ij} lanes, resulting in a distance of ℓl_{ij}, where ℓ is the lane width. For right turns, we assume that vehicles turn from the right-most lane of i into the right-most lane of j, traversing a quarter of the circumference of a circle with radius ℓ, resulting in distance $\frac{\pi}{2}\ell$. Vehicles making a left turn traverse a quarter of the perimeter of an ellipse with axes depending on the number of lanes crossed. Let $\hat{l}_i$ and $\hat{l}_j$ be the numbers of lanes crossed, then the axes are $\hat{d}_i = \left(\hat{l}_i + \frac{1}{2}\right)\ell$ and $\hat{d}_j\left(\hat{l}_j + \frac{1}{2}\right)\ell$. We approximate the distance as $\frac{\pi}{4}(3\left(\hat{d}_i + \hat{d}_j - \sqrt{(3\hat{d}_i + \hat{d}_j)(\hat{d}_i + 3\hat{d}_j)}\right)$ (Ramanujan's approximation). Without more specific data, we assume lane widths of 3.7 m (12 feet), which is a typical width for arterial roads in the USA (Stein and Neuman, 2007).

Integrating the acceleration function of equation (1) results in the following equation for displacement:

$$d_{ij} = \frac{\alpha(e^{-\beta t_{ij}} - 1)}{\beta^2} - \frac{\alpha t_{ij}}{\beta} \tag{2}$$

Eq. (2) is illustrated in Fig. 1. We approximate the solution to Eq. (2) to find the travel time t_{ij} for turning movement (i, j). The maximum number of vehicles that can make the turning movement (i, j) in unit time, assuming no conflicting traffic, is then the inverse of t_{ij}, the time required per vehicle. This is used to determine intersection capacity.

3.1.2. Minimum delay

Since all vehicles must stop at the intersection, we impose a minimum delay in the last cell of the link of $\Delta t + \frac{v_f}{a_d}$, where Δt is the CTM timestep, v_f is the free flow speed of the link and a_d is the braking deceleration, which we assumed to be 4.6 m (15 feet) per second squared for all vehicles. This produces an additional time step spent in the last cell leading to a stop sign, and possibly more for links with a high free flow speed.

3.2. Dynamic traffic assignment model

To study the possible effects of targeted attacks on a traffic signal system, we applied the DTA model described in the previous sections to the downtown Austin city network. The downtown Austin network has 546 intersections, of which 173 have traffic signals as seen in Fig. 2. Intersections include freeway merges and diverges along the right side of the figure. The network has 1247 links and 62,836 AM weekday peak period trips distributed over 64 zones throughout a 2 h, 15 min peak time period. Although the PM peak period expectedly offers worse congestion than the AM peak period due to increased trip chaining behavior, the experimental scenarios in this work assume that cyberattacks are set up and executed during the early morning.

Static travel demand is based on a Year 2010 calibrated model produced by the region's municipal planning organization. Travel demand over the AM peak for each pair of zones is uniformly distributed with respect to time. At the conclusion of the peak time period, no new vehicles are generated during the model's cool-down period, which lasts until the final vehicle arrives at its destination. For simplicity, all trips are made by passenger vehicles; the use of freight vehicles in future work would offer opportunities for further economic analyses. The traffic signal timings are based upon actual configurations in the field. In this preliminary model, disabled (flashing red) signals remain disabled throughout the entire simulation time.

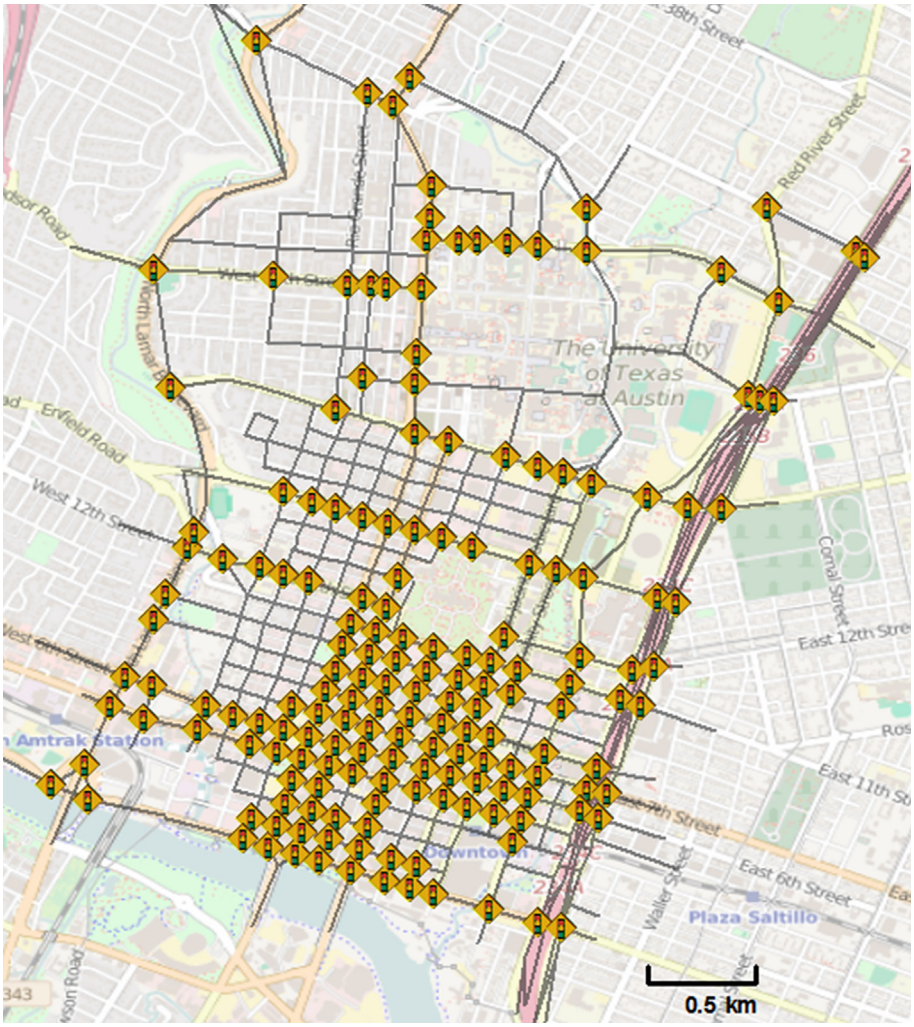

Fig. 2. Signals in the downtown Austin network. (Map imagery© OpenStreetMap contributors).

Using the method of successive averages (MSA), we found dynamic user equilibrium (DUE) for the base case network. In short, DUE involves the assignment of vehicles to travel routes such that all utilized routes between each origin/destination at a specific departure time are equal, and no faster route exists. Due to the fact that hacking signals results in temporary and unexpected intersection control behavior, for the sake of simplicity we assume that drivers are not aware of which signals are operating normally when they make their route choice decisions (and therefore there is no impact on their route choice). A more realistic analysis for future research would include driver behavior modeling, both for "go/no-go" travel decisions, and route choice. For route choice, there are opportunities in future research for looking at changes in congestion caused by increased traffic diverting to side streets, as well as changes in the use of freeway corridors. Key characteristics of Austin during the AM peak are that most of the demand arrives throughout the downtown district, and that the Austin roadway system provides few alternatives for north/south bypass traffic; other models may behave differently in these regards. See Table 1 for an overall spatial characterization of AM peak demand. For the experiment presented in this section, we replace some of the traffic signal controls with stop signs to simulate flashing red signals and

Table 1
Spatial characterization of AM peak demand.

Origin-Destination Type	Percentage
Traffic originating outside arriving to downtown	56.0%
Traffic passing through downtown via local streets	29.7%
Freeway north–south traffic passing through	3.9%
Traffic originating and arriving within downtown	4.0%
Other traffic (including originating downtown and exiting)	6.4%

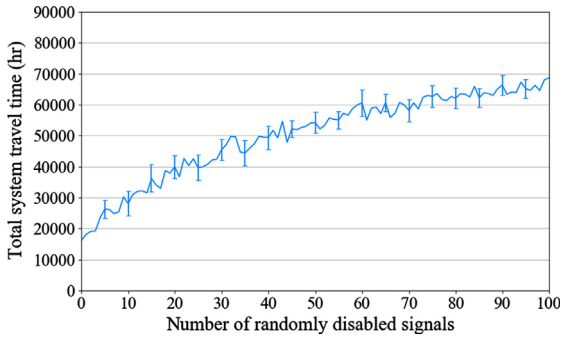

Fig. 3. Total system travel time (h) as a function of the number of randomly chosen disrupted signals, with 95% confidence intervals.

simulate vehicle movement along the routes used for the DUE with normal signals.

This section presents the results for three methods of attacks on the signal system, and two scenarios where measures are taken to protect signal controllers from attacks. We examine the impact of disabling varying numbers of traffic signals from a network-wide perspective using the metric of total system travel time (the sum of travel times for all vehicles modeled) and on individual vehicles. Additionally, we spatially analyze the impact on the average added delay for each travel zone.

3.3. Effects of disabling signals

As expected, the impact of disabling parts of the traffic signal system is significant, especially depending on the number of signals that are disabled. However, identifying which signals to target and in which order presents an additional question. In this experiment, we identify three methods for prioritizing the order in which signals are disabled: disabling intersections at random, prioritizing intersections with maximum flow first, and prioritizing intersections to affect the most vehicles. All methods seek to impact a large number of vehicles.

3.3.1. Random targeting

In the first method, signals are attacked at random. We consider that the hacker will attack a specific number of signals between 0 and 100, and that the signals hacked are chosen at random. To evaluate the impact on the total system travel time, for every number of signals disrupted, we perform 30 simulations and compute the expected total system travel time. We also compute 95% confidence intervals around the expected value and plot the results as shown in Fig. 3. Based on the United States Department of Transportation recommendation of around $14 per hour in traffic (U.S. Department of Transportation, 2016), the worst case with 100 signals disrupted randomly costs a total of $1 Million ($0.79 Million attributed to the outages), or approximately $16 per vehicle. We note that these immediate cost estimates do not account for additional economic damages to city businesses, cost of emissions, safety costs, etc.

3.3.2. Maximum vehicle flow targeting

In the second method, the signals with the greatest amount of vehicle flow are targeted first. This is referred to as the “max-vehicle-flow” method. In the DTA model, the intersections with the greatest flow are identified based on the vehicle paths. Fig. 4(a) presents the results for this method, where the horizontal axis indicates the number of signals that are targeted, i.e., replaced with four-way stops, and the vertical axis presents the total system travel time. In the base case, the total travel time of all vehicles is about

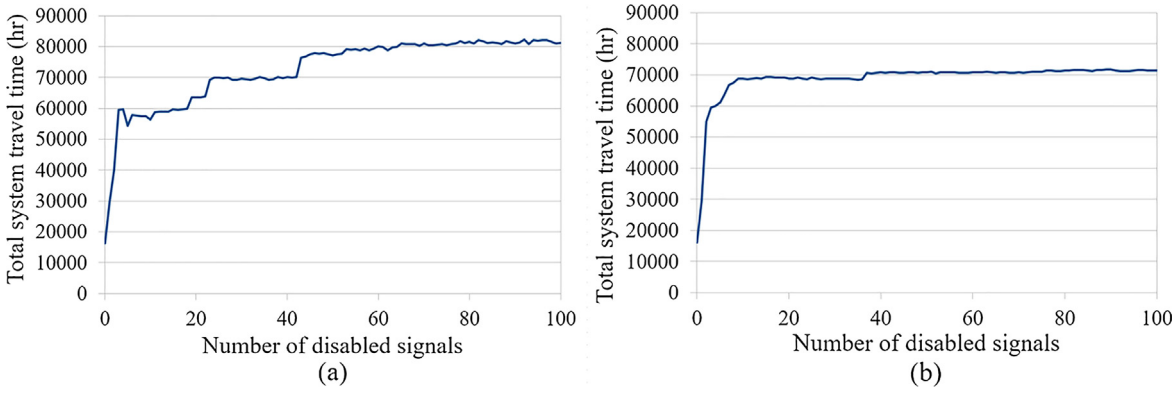

Fig. 4. Total system travel time results for (a) the max-vehicle-flow targeting method, and (b) the max-vehicles-affected method.

16,100 h, while in the worst case the total system travel time is 82,000 h. The case with 100 disabled signals over the course of the simulation duration costs $1.15 Million (of which $0.92 Million is attributed to the outages), or approximately $18 per vehicle. However, significant impact can still happen when fewer signals are disabled. For example, the outage of 26 signals costs $0.98 Million ($0.75 Million due to the outages), or approximately $15.5 per vehicle. The rate at which the delay increases converges to zero as more signals are disrupted. This is expected since as the network conditions become worse, a new disruption would not lead to a significant increase in travel time. However, we observe that when the number of hacked signals is below 20, the rate at which the total system travel time increases is larger than the corresponding rate in the random scenario. This is also an anticipated result since the targeted attacks will impact a greater number of people with fewer disruptions.

3.3.3. Maximum vehicles affected targeting

The third method for prioritizing which traffic signals to disable is based on maximizing the number of vehicles affected. This is referred to as the "max-vehicles-affected" method. This method used a greedy heuristic in which we first identify a subset of unaffected vehicles, disable a signal at the intersection that has the greatest use from the subset of unaffected vehicles, remove those vehicles from the subset, and iterate. This represents attacks that depend upon in-depth knowledge of downtown traffic patterns. Fig. 4(b) presents the results for the max-vehicles-affected method, where again the horizontal axis shows the number of disabled signals and the vertical axis shows the total system travel time in hours. Here, disabling only 7 signals runs a cost of $0.93 Million (of which $0.71 Million is attributed to the outages), or approximately $15 per vehicle. Compared to the max-vehicle-flow method, the max-affected-vehicles method results in a faster increase in total travel time with a fewer number of disabled signals. However, it would be the more difficult of the two methods for an attacker to implement.

3.3.4. Further analysis

Next, we further demonstrate the potential magnitude of the problem by further analyzing the scenario where 100 signals are disabled via the max-vehicle-flow method. Fig. 5 presents a spatial analysis of the average added delay (in minutes) for each origin zone. The average added delay is colored in blue and scaled by size according to the legend. In addition, Fig. 5 shows the total vehicle

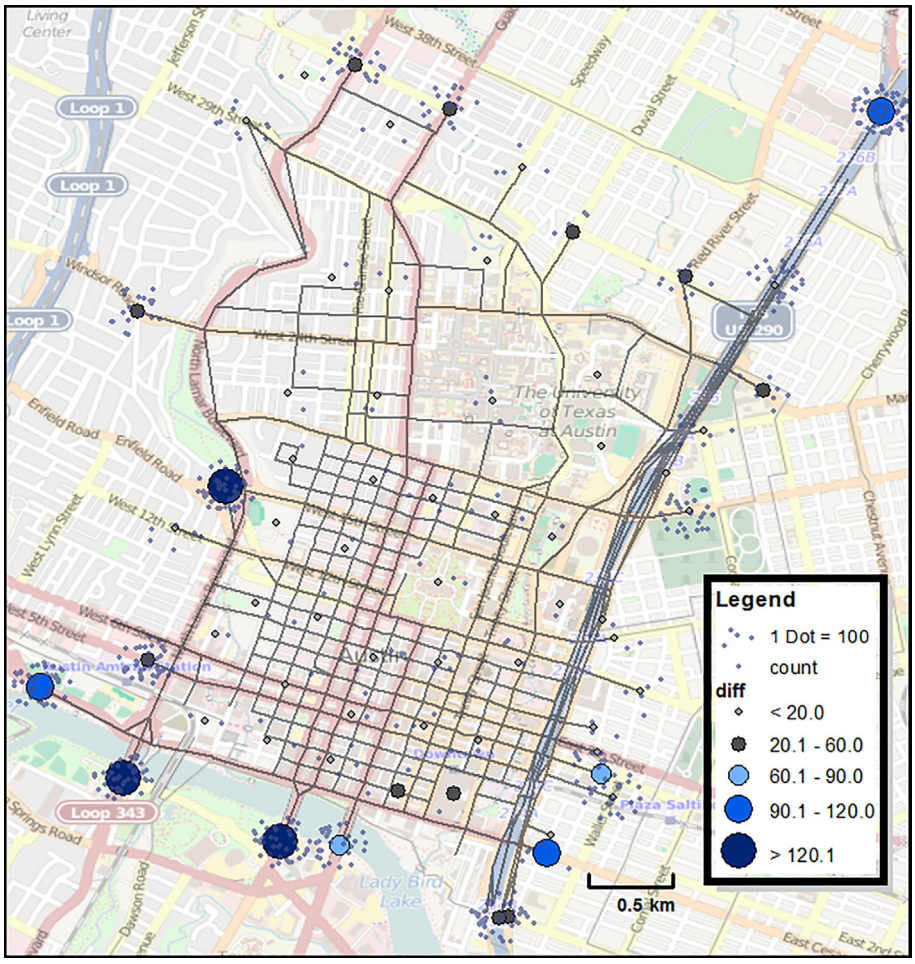

Fig. 5. Average added delay in minutes for each origin zone. (Map imagery© OpenStreetMap contributors).

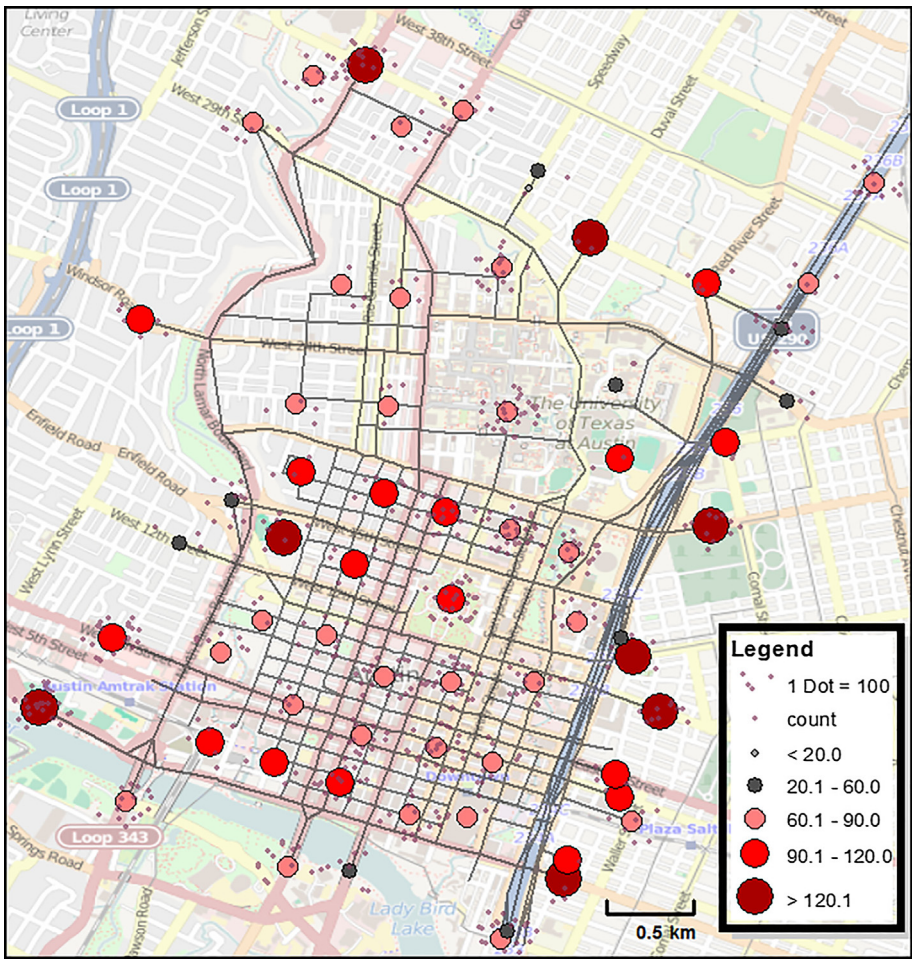

Fig. 6. Average added delay in minutes for each destination zone. (Map imagery© OpenStreetMap contributors).

demand for each origin as a dot density plot centered on each zone. Each dot represents 100 vehicles, and thus, the zones with the greater number of clustered dots have a greater amount of demand and therefore a greater number of vehicles that would be affected by the added delay.

Fig. 5 shows that each origin zone is not equally impacted. The origin zones in the center of the city have little to no added delay, in addition to also having less vehicle demand. The majority of the additional delay is experienced by the vehicles with an origin to the south of downtown Austin or on the west side at Enfield Road.

Fig. 6 shows the same results except for each destination zone. Again, the circles represent the average added delay in minutes for each vehicle in that zone. The small red dots represent the vehicle demand centered on the destination zone, where each dot represents 100 vehicles. Unlike Fig. 5, the added delay in the worst case scenario is more evenly distributed between all the destination zones and the location of the vehicle demand. While the zones in the center of the city appear to have a smaller average added delay, they also have more vehicles, implying that they experience a greater proportion of the total added delay.

Finally, Fig. 7 shows how the added delay in the same scenario as compared to the base case scenario (which total approximately 82,000 h) is distributed among the individual vehicles. The horizontal axis shows the grouping of the minutes of added delay while the vertical axis shows the number of vehicles whose added delay is in that group. As indicated in Fig. 5 and Fig. 6, vehicles are not equally impacted by the failure of the traffic signal system. The majority of vehicles experience delays between 0 and 100 min. However, an unfortunate subset of vehicles experience an increased delay of several hours. Further improvements on behavioral modeling strategies (including dynamic change in path choice), as well as mitigation and repair strategies after signals are disabled would likely reduce the number of travelers who experience the longer delays.

3.4. Effects of limited intervention

The previous section demonstrates the impacts in a scenario where a significant number of traffic signals are disabled during the time of peak morning demand. However, this work also explores the scenario where some traffic signals are protected, called intervention strategies. The intervention strategy is to protect a number of signals, meaning that they cannot be disabled. Fig. 8 shows

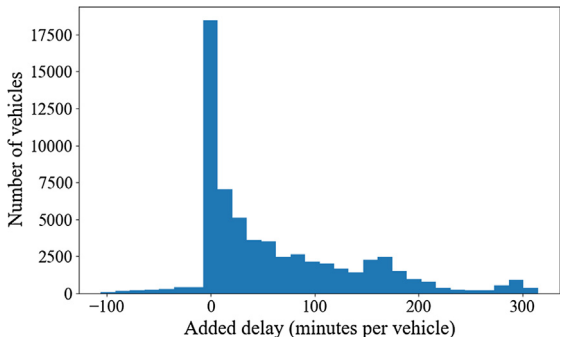

Fig. 7. Frequency distribution of the 100 max-vehicle-flow scenario added delay in minutes for all vehicles.

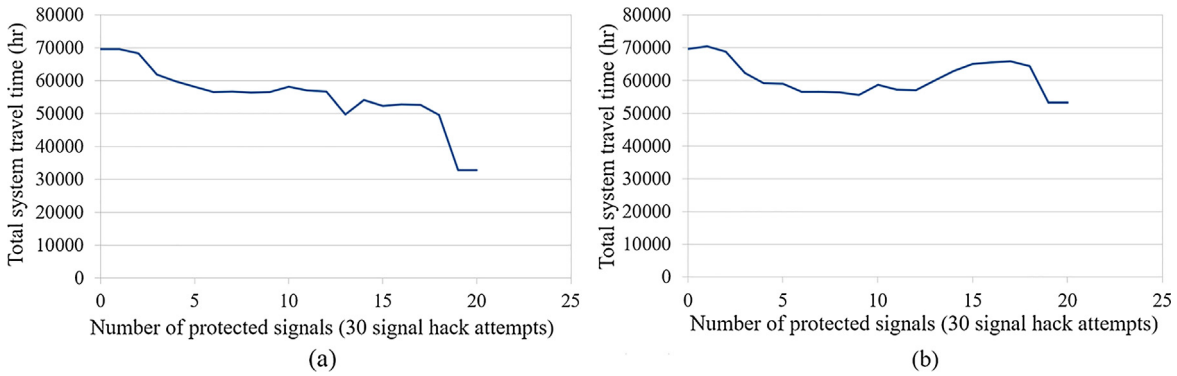

Fig. 8. The impact of (a) intervention where attacker has no information and (b) intervention where attacker has information.

the results for the intervention strategies.

In the first intervention scenario, shown in Fig. 8(a), some subset of traffic signals is protected, shown on the horizontal axis. These are protected in the order of intersection max-vehicle-flow (busiest intersection first), and the attacker does not know that they are protected. If the attacker tries to disable a protected traffic signal, nothing happens. However, the attack scenario assumes that 30 traffic signals are attacked in descending order of intersection max-vehicle-flow—a simple heuristic that an attacker may plan prior to the attack. As shown in Fig. 8(a), if the attacker does not have information on the protected signals, the impacts of the attack can be partially mitigated, although there is still a significant amount of added delay. However, Fig. 8(b) shows the intervention scenario where the attacker has information on the protected signals, and therefore will try to disable the next busiest signals that are not protected. In Fig. 8(b), the intervention strategy is not able to have much positive impact on the added delay in the network; the next busiest signals also heavily influence the traffic network operation. In either case, it is important to consider careful network analysis and other means for prioritizing the protection of signals, as a change in prioritization scheme can affect the outcome of this experiment.

4. Discussion

The research results indicate that significant negative impact can occur on a downtown traffic network when a handful of traffic signals are strategically cyberattacked. It follows that the implementation of cybersecurity policies directly benefits the successful operation of a traffic system, assuming in various cases that the protections are unknown to attackers. This section discusses relevant strategies and research needs that can help in achieving improved cybersecurity.

Research projects in the field today address many of these security concerns, including documentation of policies, best practices, and the improvement of overall security awareness. More broadly, general computer security incident response and prevention guides have been or are being drafted (U.S. Department of Homeland Security, 2013; Cichonski et al., 2012), including ones that specialize in the transportation field (Frazier et al., 2009; Ramon and Zajac, 2018). While these offer significant value to the domain of cybersecurity, policies and best practices are often expressed in voluminous, generalized ways that may not be immediately accessible or digestible to resource-limited traffic operations personnel or regional network engineers. Along with improved awareness and targeted training, availability of automated signal control network analysis tools can help bridge the gap toward better defenses.

One major function of security analysis is to bring about assurances that security practices already put into effect are working as intended (Bishop, 2003). Automated analysis tools that probe for known vulnerabilities can provide a certain degree of assurance, as well as detection of unexpected security problems. Although automated tools cannot analyze all types of security, the primary

motivation here is to discover and alleviate the worst vulnerabilities.

Emphasis here is placed upon the analysis of existing systems and assurances of intended operation. Indeed, with historic manufacture of proprietary traffic control software, the application level of security cannot be readily addressed or influenced without active intervention of the manufacturer, which can be a lengthy process. Rather, more opportunities for improving today's existing systems lie in the areas of the network level and operating system level (Lampson, 2004). For example, network level encryption can be facilitated by installing VPN hardware, and OS security can be improved by enabling better user authentication and access logging features.

Another tool that can help improve cybersecurity among multiple regions can be the creation of a ranking system that allows sets of automated analysis results to be compared with those of other jurisdictions or representative benchmarks. One example of a similar ranking system for a broader technological domain is the Common Vulnerability Scoring System (CVSS) (U.S. Department of Homeland Security, 2012). Likewise, detailed analysis results can be supplemented by practical documents that describe remedies, in the same spirit of brief notes provided by private industry and government (Moxa, 2014; Econolite, 2014; USDOT/FHWA Office of Operations, 2014). Some documents support the idea that it is generally easier for end users to procure security features than it is to implement security from scratch, where the former leverages proven solutions (Lampson, 2004).

To prepare for unfortunate scenarios where a cyberattack is successfully executed, it is important to also plan out a recovery strategy. Items to address in planning a strategy include prioritizing strategic intersections and corridors, managing and repairing outages, and preventing further attacks while repairs are underway. The prioritization of mitigations is an area for future research. The goal of a well-executed recovery plan is to allow signal operations to be restored expeditiously, limiting travel time delay and safety problems for the most travelers possible.

5. Conclusions

These initial experiments are motivated by a need to improve awareness and policies concerning impacts of possible attacks in new and existing signal control systems. Even though these experiments approximately model reality and focus only on travel delay cause by signal outages, these kinds of modeled results can significantly facilitate useful risk assessments that impact future policymaking.

The overall trends depicted in methodology show increased system-wide delay as the number of disabled signals increases. In the AM peak period model of downtown Austin, TX, if attacks are prioritized according to busy intersections, delay is multiplied by 4.3 times when 26 signals are disabled. If intersection attacks are organized to reach the maximum number of travelers, significantly fewer signals—around 7—are necessary to create a comparable impact. For a scenario in which the attacker chooses signals at random, it was observed that the rate of increase in delay was lower than the corresponding rate for targeted attacks. Importantly, the analysis on the effects of improving cybersecurity for a handful of intersections shows significant reduction in delay at the time of an attack but only in the case of the attacker not knowing which signals are protected. In general, improved cybersecurity leads to a reduction in damage. It follows that vulnerabilities must first be identified before security improvements can be most effectively applied.

For future work, extensive game-theoretic analyses (Bell and Cassir, 2002) could determine the minimum number of security interventions necessary to ensure a given level of service after an attack. As mentioned previously, this will depend on the attacker's knowledge of the interventions. Additionally, opportunities exist for improving behavioral modeling, and also modeling mitigation and repair strategies. These can assist agencies in designing recovery plans.

Concern has been expressed about the idea that published vulnerabilities and documented practices for remedies can be used by attackers for malicious purposes. Unless care is taken to protect information, there may be possible exposure of information having to do with previously unknown problems (Landwehr, 1981). In response to these concerns, it is observed that a significant amount of information is already publicly available. Further research should address these concerns in efforts to curtail the possible misuse of information. Nevertheless, fixing flaws in signal controllers and other traffic control systems is a more permanent solution than relying on confidentiality of vulnerabilities.

Acknowledgements

The authors gratefully acknowledge the support of the Data-Supported Transportation Operations & Planning Center.
This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix A. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.tra.2018.12.009.

References

Bayless, S.H., Murphy, S., Shaw, A., 2014. Connected Vehicle Assessment: Cybersecurity and Dependable Transportation. ITS America.
Bell, M.G., Cassir, C., 2002. Risk-averse user equilibrium traffic assignment: an application of game theory. Transp. Res. Part B: Methodol. 36 (8), 671–681.
Bernstein, S., Blackstein, A., 2007. Key Signals Targeted, Officials Say. Los Angeles Times.
Bishop, M., 2003. What is computer security? IEEE Secur. Priv. 1 (1), 67–69.

Byres, E., Lowe, J., 2004. The myths and facts behind cyber security risks for industrial control systems. Proc. VDE Kongress 116.
Cerrudo, C., 2014. Retrieved November 2017, from Hacking US (and UK, Australia, France, etc.) Traffic Control Systems: <http://blog.ioactive.com/2014/04/hacking-us-and-uk-australia-france-etc.html>.
Chiu, Y.-C., Bottom, J., Mahut, M., Paz, A., Balakrishna, R., Waller, T., Hicks, J., 2011. Dynamic traffic assignment: a primer. Transportation Research E-Circular (E-C153).
Cichonski, P., Millar, T., Grance, T., Scarfone, K., 2012. Computer Security Incident Handling Guide: Recommendations of the National Institute of Standards and Technology. National Institute of Standards and Technology.
Daganzo, C.F., 1994. The cell transmission model: a dynamic representation of highway traffic consistent with the hydrodynamic theory. Transp. Res. Part B: Methodol. 28 (4), 269–287.
Daganzo, C.F., 1995. The cell transmission model, Part II: network traffic. Transp. Res. Part B: Methodol. 29 (2), 79–93.
Dresner, K., Stone, P., 2004. Multiagent Traffic Management: A Reservation-Based Intersection Control Mechanism, vol. 2, 530–537.
Econolite, 2014. AN2152: Added Security for your Traffic Signal Network to Protect Your Traffic Control Devices. Retrieved from Added Security for your Traffic Signal Network To Protect Your Traffic Control Devices.
Faturechi, R., Miller-Hooks, E., 2014. Measuring the performance of transportation infrastructure system in disasters: a comprehensive review. J. Infrastruct. Syst. 21 (1).
Feng, Y., Huang, S., Chen, Q., Liu, H., Morley, M., 2018. Vulnerability of traffic control system under cyberattacks with falsified data. Transp. Res. Rec.
Fletcher, D., 2014. Retrieved November 2017, from NCHRP 20-59 (48): Effective Practices for the Protection of Transportation Infrastructure from Cyber Incidents: <http://trbcybersecurity.erau.edu/files/NCHRPProject.ppt>.
Frazier Sr, E.R., Nakanishi, Y., Lorimer, M.A., 2009. Surface Transportation Security: Security 101: A Physical Security Primer for Transportation Agencies. NCHRP Report, vol. 14.
Ghena, B., Beyer, W., Hillaker, A., Pevarnek, J., Halderman, J.A., 2014. Green lights forever: analyzing the security of traffic infrastructure. Proceedings of the 8th USENIX Workshop on Offensive Technologies.
Goodspeed, T., 2008. Reversing the Econolite ASC/3 Traffic Light Controller. ToorCon Seattle.
Halsey, A., 2009. Traffic Signals Disrupted, Creating Chaos in Montgomery. The Washington Post.
Hoo, K.J., 2000. How Much is Enough? A Risk Management Approach to Computer Security. Stanford University, Stanford, California.
Infosec Institute, 2014. Retrieved November 2017, from Hacking Traffic Light Systems: <http://resources.infosecinstitute.com/hacking-traffic-light-systems/>.
Institute of Transportation Engineers, 2003. Retrieved November 2017, from Survey Results: Traffic Signal Systems Requirements Survey: <http://library.ite.org/pub/e267a385-2354-d714-517d-dc4cd091d64d>.
Lampson, B.W., 2004. Computer security in the real world. Computer 37 (6), 37–46.
Landwehr, C.E., 1981. Formal models for computer security. ACM Comput. Surv. (CSUR) 13 (3), 247–278.
Levin, M.W., Boyles, S.D., 2015. Intersection auctions and reservation-based control in dynamic traffic assignment. Transportation Research Board 94th Annual Meeting, 15-2149.
Li, H., Tian, Z., Deng, W., 2011. Capacity of multilane all-way stop-controlled intersections based on the conflict technique. Transp. Res. Rec. J. Transp. Res. Board 2257, 111–120.
Moxa, 2014. Retrieved November 2017, from Moxa Connection: Cybersecurity for Centralized Advanced Traffic Management Systems: <http://www.moxa.com/newsletter/connection/2014/01/feat_01.htm>.
Ramon, M.C., Zajac, D.A., 2018. Cybersecurity Literature Review and Efforts Report. Southwest Research Institute.
Reilly, J., Martin, S., Payer, M., Bayen, A.M., 2015. On cybersecurity of freeway control systems: analysis of coordinated ramp metering attacks. Transportation Research Board 94th Annual Meeting, No. 15-5248.
Sawin, D., 2010. Retrieved June 2015, from Control Systems Security Program: Transportation: <http://www.volpe.dot.gov/sites/volpe.dot.gov/files/docs/CSSPT_Conference_Presentation.pdf>.
Stein, W.J., Neuman, T.R., 2007. Mitigation Strategies for Design Exceptions.
Sullivan, J.L., Aultman-Hall, L., Novak, D.C., 2009. A review of current practice in network disruption analysis and an assessment of the ability to account for isolating links in transportation networks. Transp. Lett. 1 (4), 271–280.
Sullivan, J.L., Novak, D.C., 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: Policy Pract. 44 (5), 323–336.
Tampère, C.M., Corthout, R., Cattrysse, D., Immers, L.H., 2011. A generic class of first order node models for dynamic macroscopic simulation of traffic flows. Transp. Res. Part B: Methodol. 45 (1), 289–309.
U.S. Department of Homeland Security, 2012. A Closer Look at CVSS Scoring. ICS-CERT Monitor, 2.
U.S. Department of Homeland Security, 2013. NIPP 2013: Partnering for Critical Infrastructure Security and Resiliance.
U.S. Department of Transportation, 2016. The Value of Travel Time Savings: Departmental Guidance for Conducting Economic Evaluations Revision 2 (2016 Update). U.S. Department of Transportation. Retrieved November 2018, from The Value of Travel Time Savings: Departmental Guidance for Conducting Economic Evaluations, Rev. 2 (2014 Update): <https://www.transportation.gov/sites/dot.dev/files/docs/vot_guidance_092811c.pdf>.
USDOT/FHWA Office of Operations, 2014. Retrieved November 2017, from Cyber Security Advisory: <http://trbcybersecurity.erau.edu/files/CSAdvisory-20148-Final.pdf>.
Wang, J., Dixon, K., Li, H., Ogle, J., 2004. Normal acceleration behavior of passenger vehicles starting from rest at all-way stop-controlled intersections. Transp. Res. Rec. J. Transp. Res. Board 1883, 158–166.
Wang, Z., Chan, A.P., Yuan, J., Xia, B., Skitmore, M., Li, Q., 2015. Recent advances in modeling the vulnerability of transportation networks. J. Infrastruct. Syst. 21 (2).
Wu, N., 2000. Determination of capacity at all-way stop-controlled intersections. Transp. Res. Rec. J. Transp. Res. Board 1710, 205–214.
Wu, N., 2002. Total capacities at all-way stop-controlled intersections: validation and comparison of highway capacity manual procedure and addition-conflict-flow technique. Transp. Res. Rec. J. Transp. Res. Board 1802, 54–61.

Glossary

Capacity: The maximum vehicle flow that a roadway can support
Cell Transmission Model (CTM): A popular technique for modeling traffic flows, where roadway links are divided into discrete sections, or cells. A free-flow vehicle is modeled to traverse one cell within a fixed time interval, such as 6 s
Conflict Monitor or Malfunction Management Unit (MMU): A device in a traffic signal control system that verifies proper operation of a traffic controller. If an anomaly is detected, then the conflict monitor will place the intersection into flashing operation
Demand: Vehicles that are poised to travel through a path or roadway segment within a traffic model. Demand may be based upon analysis of the number of vehicles expected to travel from a specific origin to a specific destination around a particular time of day
Dynamic Traffic Assignment (DTA): A family of route choice models in which travel time over each link may vary depending upon time of day. The optimal route from one origin to one destination may change because of varying traffic volume and demand. DUE is achieved after a number of iterations
Dynamic User Equilibrium (DUE): A characteristic of some DTA models that defines a state of convergence, where all paths taken by vehicles from a given origin to a given destination around a specific time of day have equivalent travel time
Field: The environment of streets used by daily travelers
Flashing Operation: A mode of traffic signal control systems that involves flashing red and possibly yellow displays in such a way that safe traffic control is maintained
Flashing Red: A signal indication within the USA and other countries where drivers are to stop as though the intersection is controlled by a two-way or four-way stop

sign configuration

Flashing Yellow: A signal indication within the USA and other countries where drivers are to proceed through an intersection with caution

Four-way Stop: A stop sign configuration where vehicles from all approaches of an intersection are to stop and proceed according to a defined order. In the USA, the stopped vehicle in the approach to the right is to proceed first.

Free-flow Speed: The speed that vehicles travel through a roadway in light traffic conditions, and without interruptions caused by traffic control devices

Link: An element within a graph representation of a roadway network that represents a section of roadway, often from one intersection to the next

Method of Successive Averages (MSA): A technique within DTA for reaching DUE by gradually decreasing the number of vehicles that are eligible to be reassigned to faster paths at each DUE iteration. This is usually according to a probability defined by $1/n$, where n is the DTA iteration number

Node: An element within a graph representation of a roadway network that joins adjacent links together. An intersection would be represented by a node

Signal Controller: A programmable and sometimes remotely controlled device that continuously determines the displays within vehicle and pedestrian signal heads. While very old devices are electromechanical, newer devices are specialized networkable computers

Signal Control Network: A computer network that connects traffic signal controllers, often involving miles of fiber optic cables or wireless technologies. The network provides capabilities for monitoring and remote-controlling traffic signal controllers in the field

Traffic Control Devices: Signs and signals that affect the flow of traffic

Total Delay: The sum of delay experienced by all vehicles during the entire analysis period

Turning Movement: A traversal of a vehicle through two links adjoining an intersection within a graph representation of a roadway network

Vehicle Flow: A measure of roadway usage, often measured as vehicles per hour as observed from a specific location