Creating Complex Congestion Patterns via Multi-objective Optimal Freeway Traffic Control with Application to Cyber-security

Jack Reilly

University of California, Berkeley, 652 Sutardja Dai, Berkeley, CA 94720

Sébastien Martin

Massachusetts Institute of Technology

Mathias Payer

Purdue University, 305 N University Street, West Lafayette, IN 47907

Alexandre M. Bayen

University of California, Berkeley, 642 Sutardja Dai, Berkeley, CA 94720

Abstract

This article presents a study on freeway networks instrumented with coordinated ramp metering and the ability of such control systems to produce arbitrarily complex congestion patterns within the dynamical limits of the traffic system. The developed method is used to evaluate the potential for an adversary with access to control infrastructure to enact high-level attacks on the underlying freeway system. The attacks are executed using a predictive, coordinated ramp metering controller based on finite-horizon optimal control and multi-objective optimization techniques. The efficacy of the control schemes in carrying out the prescribed attacks is determined via simulations of traffic network models based on the cell transmission model with onramps modeled as queue buffers. Freeway attacks with high-level objectives are presented on two illustrative examples: congestion-on-demand, which aims to create precise, user-specified pockets of congestion, and catch-me-if-you-can, which attempts to aid a fleeing vehicle from pursuant vehicles.

Keywords: traffic optimization, ramp metering, cybersecurity, multi-objective optimization

1. Introduction

Public traffic infrastructure is arriving in the cyber age with increasing connectivity between the different segments of roadways. For example, freeways are commonly instrumented with loop detectors that allow for real-time monitoring of roadway speeds [1]. Estimates of road traffic conditions are then fed directly into onramp traffic light metering or variable speed limit algorithms which regulate traffic flow to improve congestion [2, 3]. Finally, these metering algorithms can be coordinated and controlled by a remote command and monitoring center, leading to a regional network of interconnected sensors and controllers [4, 5, 6, 7, 8]. Increased efforts to build systems which understand and utilize the interconnectivity are evidenced by *integrated-corridor-management* (ICM) projects such as *Connected Corridors* [9] and mobile applications which use GPS probe data to improve navigation [10].

Integration of control infrastructure is exemplified by freeway networks instrumented with coordinated, predictive ramp metering control. Coordinated ramp-metering strategies have been investigated extensively [11, 12, 13, 14,

*Corresponding Author

Email addresses: jackdreilly@berkeley.edu (Jack Reilly), semartin@mit.edu (Sébastien Martin), mpayer@purdue.edu (Mathias Payer), bayen@berkeley.edu (Alexandre M. Bayen)

15] as a control scheme for improving and regulating freeway traffic conditions. Traffic management districts have investigated and piloted such schemes [16, 17] due to their ability to anticipate congestion formation and distribute onramp queues across a corridor. In comparison, many ramp metering strategies used in practice [2, 18] are more *reactive* and *decentralized* in nature, but require less communication infrastructure. Notably, the AMOC freeway traffic control tool developed in [4, 19, 20] has been shown to reduce congestion 20-30% in realistic simulation environments.

Controllability analysis allows one to evaluate the effectiveness of integrated schemes such as coordinated metering. Given that onramp metering lights only exist at sparsely-distributed locations along a freeway stretch, one cannot expect metering lights alone to *exactly* achieve any desired objective. Rather, traffic dynamics will dictate congestion patterns at uncontrollable locations. The goal then is to select metering rates which produce congestion patterns *as close as possible* to the prescribed congestion. The deviation between the desired and resulting traffic states serve as an indication of the controllability of the system under study.

There exists a number of applications which make use of the above approach. As an example, one can use congestion pattern replication for traffic model calibration [21]. If one has accurately measured the congestion states at a given instance, then one could search for the optimal set of model parameters which would reproduce the observed congestion.

A consequence of increased controllability is the increased vulnerability and impact of a control system compromise. A compromise at any level of the traffic control infrastructure can lead to both direct access of an attacker to alter traffic lights and changeable message signs, and indirect access via spoofing of sensor readings, which may *trick* the control algorithms to respond to false conditions.

While much research has been conducted on the security of inter-connected vehicles [22, 23], the security of transportation management infrastructure [24, 25] has received less attention. A number of recent compromises underscore the importance of investigating infrastructure security. A man-in-the-middle attack on GPS coordinate transmissions from mobile navigation applications showed it is possible to trick navigation services into inferring non-existent jams [26], while a similar attack used a fleet of mobile phone emulators to mimic the presence of many virtual vehicles on a roadway [27]. A popular vehicle-detection sensor was revealed to use a type of wireless protocol vulnerable to data injection attacks, and a demonstration showed that the access point could be tricked into receiving arbitrary readings [28]. Even insider attacks on command centers have precedent as two Los Angeles traffic engineers in 2009 were found guilty of intentionally creating massive delays by adjusting signal times at key intersections [29].

Given the existence of such vulnerabilities and the scale at which they can be exploited, understanding the nature and costs of such attacks becomes paramount to public safety. In this work, we use freeway controllability to analyze the set of adversarial objectives an attacker could achieve with coordinated metering control.

To do so, we first construct a taxonomy of different vulnerability locations in traffic control systems, defining three distinct layers: physical, close-proximity, and virtual. With each potential attack, we also associate values pertaining to difficulty, impact, and cost of resources expended by the attacker. We motivate our classifications by presenting two scenarios that combine a number of attacks to accomplish a high-level goal.

We then focus our analysis on an in-depth exploration of freeway attacks using coordinated, predictive, ramp metering. The control scheme employed by the attacker uses finite-horizon optimal control based on the adjoint method for finding optimal metering rates to create a desired disruption via complex congestion patterns. We additionally give an overview of multi-objective optimization and discuss how such an approach is useful for solving high-level attack objectives which contain many conflicting sub-goals, such as permitting a fleeing vehicle to escape pursuants on a particular freeway stretch without overly congesting freeway regions irrelevant to the pursuit. While we focus our discussion to ramp metering, the approach is general enough to consider other freeway control attacks, such as modifying variable-speed-limit signs [30, 31, 3] or route-guidance message signs [32, 33, 34].

While counter-measures to the proposed attack methods are outside the scope of this work, potentially relevant approaches have been proposed for physically related problems. In [24], the problem of detecting spoofed GPS probe measurements from vehicles is posed as a mixed-integer linear programming problem. [35, 36] presents methods for detecting stealthy attacks on water SCADA systems is presented. In both scenarios, leveraging knowledge of the physical system under attack is key to detection. Using a similar approach for freeway system security may be promising.

The contributions of this article are enumerated along with its outline. We present a classification of a broad set of attacks on traffic control systems with their relation to the underlying physical and cyber infrastructure (Section 2).

To study the feasibility of such attacks, we develop an adjoint-based multi-objective optimal controller to achieve arbitrary congestion patterns within the dynamical system's limits (Section 3). The controller is demonstrated for two scenarios using a macroscopic flow simulator based on the model in [37] (Section 4); in the first scenario, an attacker wishes to congest precise locations at precise times; in the second, we consider the aforementioned problem of constructing congestion patterns to enable fleeing vehicles to escape pursuants. Simulations are conducted on realistic freeway models, with one experiment conducted on a full-sized model of a 19.4 mile stretch of the I15 South Freeway in San Diego, California. We conclude with future areas of study for traffic system security (Section 5).

2. Traffic System Vulnerabilities

Sections 3-4 detail how one may use optimal control techniques to evaluate the potential impact of a traffic management system compromise. This section reviews standard architectures of such management systems upon which such compromises could occur.

2.1. The Freeway Control System

Modern freeways encompass control and monitoring mechanisms which enable traffic management to mitigate congestion and improve traffic flow in real-time. While the exact combination of sensors, controllers and transmitters differ from location to location, this article chooses one particular instantiation of a freeway control system, which we find to be representative. Figure 1(a) shows a control system installed near a junction of a freeway and an onramp. We consider three elements of the control system:

- Sensors, used to gather information about the freeway state. For example, loop detectors are used to acquire the flow of vehicles along the freeway and onramps/offramps, while the trajectory of vehicles equipped with GPS (or containing GPS-powered smartphone applications) can be used for estimating real-time traffic conditions [10].
- Actuators, used to influence the evolution and efficiency of the freeway. The most common actuation strategy is *ramp metering*, where traffic lights installed on freeway onramps control the influx of vehicles to the mainline. Other actuators include variable speed limit control [30] and variable message signs. For the purposes of this article, the ramp meters are the only actuators we will consider.
- Local controllers, such as 2070 boxes [38] and the older 170 boxes [39], which allow interaction between the sensors and ramp meters.

Figure 1: The physical roadway, sensors, connected vehicles and controllers near a freeway/onramp junction in Figure 1(a) form a cyber-physical network we refer to as a local freeway control system. The mask icons (white/black masks for indirect/direct vulnerabilities) denote vulnerability points in the local control network. In Figure 1(b), the local controllers are wired together, then connected to a command center via a relay box to form the global control system. This article analyzes vulnerability locations associated with each component.

Attack Description	Access	Control	Complexity	Cost
copper theft/clipping wires	physical	low	low	low
replacing a single sensor/actuator	physical	low	low	low
attacking a single sensor/actuator	proximal	low	medium	low
replacing a single control box	physical	medium	medium	medium
replacing a set of sensors/actuator	physical	medium	medium	medium
attacking a set of sensors/actuator	proximal	low	medium	low
replacing a corridor of control boxes	physical	high	medium	medium
attacking a corridor of control boxes	virtual	high	high	medium
attacking the control center	virtual	high	high	high
spoofing GPS data	virtual	medium	high	medium
attacking navigation software	virtual	medium	medium	medium

Table 1: List of possible infrastructure attacks categorized by the following: access to different layers that is needed, level of control that the attacker gains, sophistication of the attack, and cost of resources an attacker would need to expend.

We assume control boxes are wired to the nearby metering light and have a wireless connection to nearby sensors. We note that sensors are not typically connected in a wireless fashion, but their reduced maintenance costs make them a potentially desirable and viable choice in the near-future. Vehicles with navigation devices such as TomTom [40] automatically analyze radio-broadcasted traffic reports from traffic control centers to improve their navigating functionality.

In order to allow coordinated control and sensing across a freeway stretch with many onramps, the local control systems are connected to allow for a more global configuration. Local controller networking would permit reactive, coordinated schemes as studied in [7]. Figure 1(b) depicts our representative global communication architecture. The local control boxes are wired together along the freeway to form the actuation network, with intermediary *relay boxes* allowing for an uplink and downlink to a remote *command center*. The command center contains instrumentation and personnel for monitoring traffic conditions and setting the metering lights accordingly.

2.2. Infrastructure Weaknesses

The traffic control infrastructure is built up of several layers and each layer poses individual security risks, starting from tampering with the actual devices, cables or wireless signals, to attacking the software of deployed devices or attacking the command center. Attackers can leverage vulnerabilities in the infrastructure to control or disrupt these connected systems. Individual attacks can thereby target the physical layer, the communication layer, the layer of the control center, or any combination thereof.

Direct physical access: The physical layer is the lowest attackable layer and involves direct access to individual wires, opening and accessing the control box, or tampering with individual sensors. Physical attacks involve clipping, tampering, removing, or replacing of wires or hardware. For instance, copper wire theft near freeways is a common occurrence [41, 42].

Proximal access: An attack on the communication layer forges, removes, replaces, or inserts attacker-controlled measurements into the control system, which may then make further decisions based on forged data. An attacker can either replace or add sensors to the current sensor network to inject new measurements or attack the software running on sensors and/or actuators to take over control.

Networked/virtual access: Remote connections from the physical freeway infrastructure to the command center defines another layer with potential vulnerabilities. An attack on this layer can be done by forging or controlling messages from/to the command center and possibly even compromises the command center itself. For this scenario an attacker needs to find software vulnerabilities in the software running in the command center. Table 1 gives a (partial) list of vulnerabilities in our freeway control system along with classifications for each attack.

2.3. Attack Scenarios

We will consider two fictional but realizable attack scenarios and study their consequences on the compromised network. The first scenario involves indirect control of the freeway, through spoofing the sensors, to achieve a local

objective. The second scenario involves direct control of the ramp meters to achieve a global objective along a larger stretch of freeway.

The distinction between direct and indirect control is illustrated in Figure 1(a) via the white mask (indirect) and black mask (direct) icons; direct control can set arbitrary metering rates to a single traffic light or to many lights in a coordinated fashion, while indirect control only modifies sensor readings, with the assumption that the uncompromised metering system will respond to the spoofed sensors in a predictable manner. Many traffic management agencies will rely on numerous data sources to make real-time traffic estimates, so a compromise of multiple data sources may be required to successfully achieve an indirect attack. Examples of direct attacks include a compromise of the 2070 boxes which are directly wired to the meters and a compromise of the command center, which issues upstream metering plans to the 2070 controllers. Examples of indirect attacks include sending fake loop-detector readings to access points and broadcasting false traffic reports to GPS devices which may respond with poor routing advice.

2.3.1. Indirect Attack: VIP-lane

The objective of the attacker is to clear a predetermined section of a regularly congested freeway. The attacker drops low-cost wireless transmitters near the 2070 controllers along the freeway section. This was demonstrated in May 2014 at the White House-hosted SmartAmerica Conference. As the actual loop-detector sensors communicate with the control box wirelessly, the attacker will be able to override the loop-detector signals and send false data that indicates a fully congested freeway. This will indirectly affect the ramp meters, which will respond by limiting onramp flow and thus eliminating significant freeway mainline flow. The attacker will then transmit false GPS location data via a set of hacked cellphones to trick navigation software into believing the freeway is congested. Approaching vehicles using navigation software will then be rerouted around the fake congestion which leads to a further reduction in incoming flow. The net effect of the attack is a congestion-free commute for the attacker: a private VIP lane created purely by indirect, sensor-based attacks.

2.3.2. Direct Attack: catch-me-if-you-can

The objective of the attacker is to escape from pursuants along a large section of freeway. A compromise of all the ramp meters is assumed, as it permits the attacker to selectively congest certain sections of the roadway (see Section 3). One approach is to hack the command center itself, with the downside being the cost and complexity of such an attack (see Table 1). Another solution is to begin by hacking one of the 2070 boxes, and since all the 2070 boxes are networked along the freeway (see Figure 1(b)), a single hacked box can serve as a means of compromising the other nearby boxes, leading to a cascading attack. The attacker can then acquire full control of all the 2070 boxes, and in turn, the ramp metering lights.

Since ramp metering does not enable complete control over the state of traffic, the feasibility of the “catch-me-if-you-can” attack is dependent upon the controllability of the freeway and control infrastructure under consideration. The next section presents a multi-objective, optimal control-based approach to exploring such dependencies.

3. Controllability and Coordinated Freeway Attacks

The attacker is assumed to have full control of a series of ramp metering lights along a freeway section. The particular control scheme is a finite-horizon, optimal controller (Section 3.2). This scheme allows effective coordination of the controls and takes advantage of traffic demand prediction. The controller is then embedded in the context of a multi-objective framework to analyze the feasibilities of individual objectives combined into a single objective (Section 3.3). First, we outline the freeway model used for developing the optimal controller.

3.1. Freeway Model

We model the freeway as a sequence of n mainline links (labeled $1, \dots, n$), where both an onramp and offramp are present between consecutive links. For spatial cells which do not have an adjacent onramp (or offramp), one can set the vehicle demand to zero (set the offramp turning ratio to zero). Flow dynamics along a link i is modeled using

a discretized version of the *Lighthill-Whitham-Richards* [43, 44] (LWR) partial differential equation. The continuous LWR equation takes the following form:

$$\frac{\partial \rho_i(t, x)}{\partial t} + \frac{\partial f(\rho_i(t, x))}{\partial x} = 0, \quad (1)$$

with $\rho_i(t, x)$ representing the *density* of vehicles on link i at a particular point in space and time, and f capturing the relationship between the density and *flow* of vehicles, a relationship referred to as a *fundamental diagram* of traffic. We assume f has the following triangular form [45]:

$$f(\rho) = \min(v\rho, w(\rho^{\max} - \rho), f^{\max}),$$

where $v, w, \rho^{\max}$ and $f^{\max}$ are characteristics of the particular freeway section.

The discrete model used in this work is adapted from [37, 5] and was chosen for its suitability to ramp metering applications. Following [5], we discretize Equation (1) into cells of spatial size Δx and temporal size Δt using a *Godunov-based* or *cell-transmission-model* (CTM) scheme [46, 45, 47]. The resulting discrete model has T time-steps, N spatial cells, and N onramps and offramps. The state of cell $i \in [1, N]$ at time $k \in [1, T]$ is given by $\rho[i, k]$, while the number of vehicles on the adjacent onramp is given by $l[i, k]$. We require the following additional variables (specific to time-step k):

- $\delta[i, k]$: Maximum flow of vehicles exiting link i .
- $\sigma[i, k]$: Maximum flow of vehicles entering link i .
- $d[i, k]$: Maximum flow of vehicles exiting onramp i .
- $r^{\max}$: Physical capacity of onramp i .
- $f^{\text{in}}[i, k]$: Actual flow entering link i .
- $f^{\text{out}}[i, k]$: Actual flow exiting link i .
- $r[i, k]$: Actual flow exiting onramp i .
- $\beta[i, k]$: Fraction of total flow from link i entering link $i + 1$ as opposed to offramp i .
- p : Fraction of mainline flow given priority over onramp flow when merging in congestion.
- $D[i, k]$: Flow entering onramp i .

The states of cell and onramp i are advanced from time k to $k + 1$ according to the following equations:

$$\delta[i, k] = \min(v\rho[i, k], f^{\max}) \quad (2)$$

$$\sigma[i, k] = \min(w(\rho^{\max} - \rho[i, k]), f^{\max}) \quad (3)$$

$$d[i, k] = \min(l[i, k] / \Delta t, r^{\max}) \quad (4)$$

$$f^{\text{in}}[i, k] = \min(\sigma[i, k], d[i - 1, k] + \beta[i, k] \delta[i, k]) \quad (5)$$

$$f^{\text{out}}[i, k] = \begin{cases} \delta[i, k] & \text{if } \frac{pf^{\text{in}}[i+1, k]}{\beta[i, k](1+p)} \geq \delta[i, k] \\ \frac{f^{\text{in}}[i+1, k] - d[i+1, k]}{\beta[i, k]} & \text{if } \frac{f^{\text{in}}[i+1, k]}{1+p} \geq d[i+1, k] \\ \frac{pf^{\text{in}}[i+1, k]}{(1+p)\beta[i, k]} & \text{otherwise} \end{cases} \quad (6)$$

$$r[i, k] = f^{\text{in}}[i, k] - \beta[i, k] f^{\text{out}}[i, k] \quad (7)$$

$$\rho[i, k + 1] = \rho[i, k] + \frac{\Delta t}{\Delta x} (f^{\text{in}}[i, k] - f^{\text{out}}[i, k]) \quad (8)$$

$$l[i, k + 1] = l[i, k] + \Delta t (D[i, k] - r[i, k]) \quad (9)$$

Equations (2)-(9) model the merging of onramp and mainline flows, as well as the propagation of congestion waves across the freeway network. The freeway-onramp-offramp junction shown in Figure 2 gives a spatial relation of the state variables.

Figure 2: A freeway-onramp-offramp junction. At time-step k , the upstream mainline density $\rho[i, k]$ at link i and onramp queue $l[i, k]$ merge and either exit the offramp with a *split-ratio* of $(1 - \beta)$ or continue onto the downstream mainline at link $i+1$. The control $u[i, k]$ scales the total demand from onramp $d[i, k]$ by a factor between 0 and 1.

Onramp Metering Model. We introduce a control parameter $u_i[k] \in [0, 1]$, a scaling factor on the demand of onramp i at time-step k . The control $u[i, k]$ represents the influence of onramp traffic lights on the discrete model. In reality, ramp metering controller hardware [38] may restrict metering rates to take only a finite set of values on a subset, limiting the capabilities of the metering effectiveness. Such restrictions are not considered to simplify the presentation of the model. We augment Equation (4) to include the introduced control:

$$d[i, k] = u[i, k] \min(l[i, k] / \Delta t, r^{\max}) \quad (10)$$

3.2. Finite-Horizon Optimal Control and the Adjoint Method.

Using the model in Section 3.1, we seek a method to compute a coordinated ramp metering policy $u[i, k]$ over all space $i \in [1, N]$ and time $k \in [1, T]$, which minimizes (or reduces) some specified objective. We cast the problem as a finite-horizon optimal control problem, and present a method, referred to as the *adjoint method* (mathematically equivalent to the use of *co-state* equations [4]), for solving such constrained optimization problems. To construct the required initial and boundary conditions for the optimal control problem, attackers may use publicly-available traffic data sources such as online loop-detector data services (e.g. PEMS [1]) or crowd-sourced real-time traffic data (e.g. Google real-time traffic API, Waze incident feed); see [10] for approaches to real-time traffic state estimation and [5] for model-predictive-control techniques to handle noisy state estimation in optimal control.

Generally speaking, we consider the minimization of some objective that is a function of both the control variables and the *state* variables. The state variables are assumed to be a deterministic function of the control variables. Let $\mathbf{u}$ be the concatenation of all metering control parameters $u[i, k]$ and let ρ be the concatenation of all state variables (variables not controlled directly, e.g. density and queue length variables). After concatenating all the discrete Equations (2)-(3), (5)-(9), and moving all terms to the left-hand side, one can succinctly express the discrete, controllable dynamical system by:

$$H(\mathbf{u}, \rho) = 0. \quad (11)$$

We note that the min and *if-then-else* expressions in Equations (2)-(10) have deterministic, closed-form solutions, they can be expressed in function notation, just the same as if they had not contained a min or *if-then-else* statement. For instance, there is no problem in the well-posedness of Equation (2) within $H(u, \rho)$:

$$H_{\delta, i, k}(u, \rho) = \delta[i, k] - \min(v\rho[i, k], f^{\max}).$$

Given some objective function $J(\mathbf{u}, \rho)$, our goal is now to find the optimal $\mathbf{u}^*$ which solves the following constrained *finite-horizon optimal control* problem:

$$\min_{\mathbf{u}} J(\mathbf{u}, \rho) \quad (12)$$

$$\text{subject to: Equation (11).} \quad (13)$$

In the above, we omit the $u_i[k] \in [0, 1]$ inequality constraints from the general formulation for clarity of presentation. Box constraints are generally handled in optimal control applications using log-barrier methods or projected gradient-descent; see [5] for additional details.

The set of constraints on $\mathbf{u}$ could be modified to account for other modeling factors such as driver behavior (e.g. maximum time drivers would accept to wait at a red light before ignoring the signal). While the finite-horizon optimal control framework permits it, we do not consider such constraints to simplify the modeling presentation.

As J and H may be non-convex functions of the control and state, it is not always possible to efficiently find the global optimum of J in Problem (12)-(13). Thus, we use a first-order gradient descent approach as a means of reducing the objective value. To handle the non-differentiability of the *if-then-else* functions in Equations (2)-(10), we solve Problem 12 for local-optima using the OpenOpt [48] implementation of Shor's *r-algorithm* [49] for non-smooth, non-linear problems.

We next detail an efficient method for computing gradient values, which is used as a subroutine within the nonlinear solver.

Gradient Computations via the Adjoint Method. We now need to compute the gradient of J with respect to the control variables $\mathbf{u}$ subject to the H constraints. With the partial derivative¹ expressions of H and J , we can compute the gradient of J with respect to $\mathbf{u}$:

$$\nabla_{\mathbf{u}} J(\mathbf{u}', \rho') = \frac{\partial J(\mathbf{u}', \rho')}{\partial \rho} \frac{d\rho}{d\mathbf{u}} + \frac{\partial J(\mathbf{u}', \rho')}{\partial \mathbf{u}} \quad (14)$$

or in abbreviated notation:

$$\nabla_{\mathbf{u}} J = J_{\rho} d_{\mathbf{u}} \rho + J_{\mathbf{u}} \quad (15)$$

It is often prohibitively expensive to compute $d_{\mathbf{u}} \rho$ explicitly and we seek to eliminate the term from the computation. The gradient of H with respect to $\mathbf{u}$ is always zero (since the right hand side is constant for feasible $\mathbf{u}, \rho$):

$$\nabla_{\mathbf{u}} H = H_{\rho} d_{\mathbf{u}} \rho + H_{\mathbf{u}} = 0, \quad (16)$$

we can add it to Equation (15) with a Lagrange-like multiplier λ :

$$\nabla_{\mathbf{u}} J = J_{\rho} d_{\mathbf{u}} \rho + J_{\mathbf{u}} + \lambda^T (H_{\rho} d_{\mathbf{u}} \rho + H_{\mathbf{u}}) \quad (17)$$

$$= (J_{\rho} + \lambda^T H_{\rho}) d_{\mathbf{u}} \rho + (J_{\mathbf{u}} + \lambda^T H_{\mathbf{u}}) \quad (18)$$

The adjoint method chooses the λ value to set the first term to zero (and eliminate $d_{\mathbf{u}} \rho$), and arrive at the following expressing for $\nabla_{\mathbf{u}} J$:

$$\nabla_{\mathbf{u}} J = (J_{\mathbf{u}} + \lambda^T H_{\mathbf{u}}) \quad (19)$$

$$\text{such that: } H_{\rho}^T \lambda = -J_{\rho} \quad (20)$$

The λ variable is commonly referred as the *discrete adjoint variable* [53, 50], while the system of equations in (20) is referred as the *discrete adjoint system*. It is shown in [5] that for freeway traffic network applications, the adjoint method leads to gradient computations which scale linearly with the size of the network and time-horizon, making it especially suitable for real-time applications.

Computation of the Adjoint Equations for Ramp Metering. To apply the adjoint method to the problem of ramp metering, we require the partial derivative terms of the dynamical system H given in Equations (2)-(9), as well as the partial derivative terms of the desired objective J . While the objective partial derivatives will be particular to the application, the dynamical system partial derivatives need only be derived once. The derivations to compute these partial derivatives are quite long, and omitted from this article for brevity. We give the explicit partial derivative expressions below, where a prime (e.g. $d'[i, k]$ vs $d[i, k]$) indicates the current point of computation:

$$\frac{\partial \delta[i, k]}{\partial s} = \begin{cases} v & s = \rho[i, k], \quad \rho'[i, k] v \leq f^{\max} \\ 0 & \text{otherwise} \end{cases}$$

¹The partial derivative terms are not always defined in terms of classical derivatives. We omit this technical detail to simplify the presentation and instead refer the reader to [50, 51, 52].

$$\begin{aligned}
\frac{\partial \sigma[i, k]}{\partial s} &= \begin{cases} -w & s = \rho[i, k], \quad w(\rho^{\max} - \rho'[i, k]) \leq f^{\max} \\ 0 & \text{otherwise} \end{cases} \\
\frac{\partial d[i, k]}{\partial s} &= \begin{cases} \frac{u'[i, k]}{\Delta t} & s = l[i, k], \quad \frac{l'[i, k]}{\Delta t} \leq r^{\max} \\ \min\left(\frac{l'[i, k]}{\Delta t}, r^{\max}\right) & s = u[i, k] \\ 0 & \text{otherwise} \end{cases} \\
\frac{\partial f^{\text{in}}[i, k]}{\partial s} &= \begin{cases} \beta[i, k] & s = \delta[i, k], \quad \sigma'[i, k] \geq \delta'[i, k]\beta[i, k] + d'[i-1, k] \\ 1 & s = d[i, k], \quad \sigma'[i, k] \geq \delta'[i, k]\beta[i, k] + d'[i-1, k] \\ 1 & s = \sigma[i, k], \quad \sigma'[i, k] < \delta'[i, k]\beta[i, k] + d'[i-1, k] \\ 0 & \text{otherwise} \end{cases} \\
\frac{\partial f^{\text{out}}[i, k]}{\partial s} &= \begin{cases} \frac{\frac{\partial \delta[i, k]}{\partial s}}{\frac{\partial(f^{\text{in}}[i+1, k] - d[i+1, k])}{\partial s} / \beta[i, k]} & \frac{p f^{\text{in}}[i+1, k]}{\beta[i, k](1+p)} \geq \delta'[i, k] \\ \frac{\partial(p f^{\text{in}}[i+1, k])}{\partial s} / (1+p)\beta[i, k] & \frac{f^{\text{in}}[i+1, k]}{1+p} \geq d'[i+1, k] \\ \text{otherwise} & \text{otherwise} \end{cases} \\
\frac{\partial \rho[i, k+1]}{\partial s} &= \begin{cases} 1 & s = \rho[i, k] \\ \frac{\Delta t}{\Delta x} & s = f^{\text{in}}[i, k] \\ -\frac{\Delta t}{\Delta x} & s = f^{\text{out}}[i, k] \\ 0 & \text{otherwise} \end{cases} \\
\frac{\partial l[i, k+1]}{\partial s} &= \begin{cases} 1 & s = l[i, k] \\ -\Delta t & s = r[i, k] \\ 0 & \text{otherwise} \end{cases}
\end{aligned}$$

Much of the nonconvexity of the problem arises from the *iffelse* conditions given in the partial derivatives. Furthermore, it is clear the numerical evaluation of the partial derivative terms are a function of the control $\mathbf{u}'$ and corresponding state ρ' values, and thus the adjoint method requires reevaluation of the partial derivative terms for different control values.

3.3. Multiple Objectives: Interactive Multi-objective Optimization

A high-level attack goal often requires satisfaction of many *sub-goals* at once, and often-times the sub-goals can be competing or conflicting. For example, in the *catch-me-if-you-can* scenario, the attacker wants to escape from his or her chasers. Hence the attacker wants to travel the freeway as quickly as possible, but also wants to slow down the chasers behind. As a consequence, we have two simpler but competing objectives. Furthermore, each sub-objective may not be achievable given the level of controllability permitted by the infrastructure. Thus, in the scenario above, one may wish to adjust the relative importance of sub-objectives upon observing some initial outcome.

Such a situation with multiple, competing objectives can be described as a *multi-objective optimization problem*. This section introduces concepts for multi-objective problems and details a method called *interactive optimization* [54] which we employ in Section 4 to allow for iteratively updating objectives in order to study controllability in a multi-objective setting.

3.3.1. Multi-objective Optimization and Pareto Front

Definition 3.1 (Multi-objective optimization problem). Given $N \in \mathbb{N}$, let $(f_i(\mathbf{u}, \rho))$ be a set of N objective functions describing the goal of a freeway attack. The *multi-objective optimization problem* we consider is the following simultaneous minimization problem:

$$\min_{x \in X} (f_1(x), f_2(x), \dots, f_N(x)) \quad (21)$$

As we are now minimizing a vector and not a scalar, we need to define how a solution of equation (21) can be “better” than another.

Definition 3.2 (Pareto front). An solution $x \in X$ is said to *Pareto dominate* another solution x' if:

- $\forall i \leq N \quad f_i(x) \leq f_i(x')$
- $\exists j \leq N \quad f_j(x) < f_j(x')$

A solution $x \in X$ is called *Pareto optimal* if there is no other solution x' that dominates it. The set of all Pareto-optimal solutions is called the *Pareto front*, $P \subseteq X$.

Hence, we consider Pareto-optimal solutions to be the solutions of Equation (21). In order to find Pareto-optimal solutions, we will reduce the problem to the common scalar minimization problem, which can be solved with the optimal control tools of Section 3.2. This process is called *scalarization*. As our particular scalarization, we use a linear combination of the individual objective functions:

$$f(x) = \sum_{i \leq N} a_i f_i(x). \quad (22)$$

It is easy to show that any solution of Equation (22) will also belong to the Pareto front. As a consequence, we can explore at least a subset of the Pareto front (with the hope that this subset is representative) by minimizing a linear combination of the objective functions. In Section 4, we embed a human operator in the Pareto-search procedure, a method commonly referred to as *interactive optimization* [54]. We also explore *A-posteriori optimization*, which conducts a Pareto-search independent of a human, and presents a summary of results to the human.

Remark 1. Due to the nonconvexity of Problem (12), globally optimal solutions of the scalarized objectives are not computationally tractable, and we resort to exploring only locally optimal solutions of the scalarized objectives, which may not lie on the Pareto set. Nevertheless, in Section 4, our scalarized solutions are shown to have a strong relationship to the underlying scalarization parameters.

4. Attacks

We will now apply the tools of *adjoint-based finite-horizon optimal control* and *multi-objective optimization* from Section 3 to two families of attacks. The first attack highlights the precision of coordinated ramp metering attacks, while the second showcases the benefits of multi-objective optimization.

Following reproducible research practices [55, 56], the software and data used to produce the numerical results and diagrams in this section is made available [57] to permit the reader to reproduce the presented results.

4.1. First Attack: congestion-on-demand

Congestion-on-demand describes a class of objectives where an attacker wishes to create congestion patterns of a specific nature. The attacks for the first example, *box objective* (to be described), use a macroscopic freeway model of a 19.4 mile stretch of the I15 South Freeway in San Diego California. The model was split into 125 links with 9 onramps and was calibrated [58, 59] using loop-detector measurements available through the PeMS loop-detector system [1]. Figure 3(a) is a *Space-time diagram* of the I15 freeway. There is no ramp metering control applied to the simulation in Figure 3(a), i.e., the ramp meters are always set to green.

4.1.1. Constructing the objective function

In order to achieve the *congestion-on-demand* objective, we will use the finite-horizon optimal control technique introduced in Section 3.2. Therefore, we need to create a class of objective functions able to represent any jam pattern on the freeway. The method we have chosen is to maximize the traffic density where we want to put the congestion, while minimizing it everywhere else.

For every cell density value at position i and time k , we assign a coefficient $a_i^k \in \mathbb{R}$. We can then define the corresponding objective function:

$$J(\mathbf{u}, \rho) = \sum_{i=1}^N \sum_{k=1}^T a_i^k \rho[i, k] \quad (23)$$

When J is minimized, a positive coefficient a_i^k will encourage the minimization of the traffic density at position i and time k , whereas a negative coefficient will encourage congestion. The absolute value of the coefficient represents the importance given to the fulfillment of the objective at the particular time and location of the simulation.

(a) Simulation with no metering.

(b) Trade-off: $\alpha = 0.3$

(c) Trade-off: $\alpha = 0.5$

(d) Trade-off: $\alpha = 0.9$

Figure 3: Figure 3(a) depicts a space-time diagram of vehicle densities on 19.4 mile stretch of I15 Freeway with no ramp metering. The box objective, and example of *congestion-on-demand*, is applied in Figures 3(b)-3(d). The user specifies a “desired” traffic jam between postmile 4.5 and 14, for a duration of 20 minutes between 8:20 and 8:40. For this, the α parameter (introduced in Equation (26)) enables the proper design of tradeoffs in the objective.

4.1.2. Examples

Box Objective. The *box objective* creates a box of congestion in the space-time diagram, i.e., congestion will be created on a specific segment of the freeway during a user-specified time interval.

As we have two competing goals (maximize congestion in the box, minimize congestion elsewhere), we apply the multi-objective optimization procedure in Section 3.3. Indeed, we have the following two objective functions:

$$f_1(\mathbf{u}, \rho) = - \sum_{(i,k) \in \text{Box}} \rho[i, k] \quad (24)$$

$$\text{and } f_2(\mathbf{u}, \rho) = \sum_{(i,k) \notin \text{Box}} \rho[i, k] \quad (25)$$

To solve this multi-objective problem, we balance our two objectives using a linear combination. As we limit ourselves to one degree of freedom, we introduce a single parameter $\alpha \in [0, 1]$ and minimize the following objective function:

$$J_\alpha(\mathbf{u}, \rho) = \alpha f_1(\mathbf{u}, \rho) + (1 - \alpha) f_2(\mathbf{u}, \rho), \quad (26)$$

where α is a trade-off parameter: $\alpha = 1$ is complete priority on the congestion inside the box, while $\alpha = 0$ is complete priority on limiting density outside the box.

Figure 4: Attack to create traffic patterns in the form of Morse code. A coordinated ramp metering attack using the proposed algorithm is able to spell out “C-A-L” in Morse code over successive time-slices of the space-time diagram: C= - · - ·, A= · - ·, L= · - · ·. The entire space-time diagram of the attack is shown in Figure 4(a), while three snapshots of the freeway are shown in Figure 4(b), each slice spelling out one of the letters in “C-A-L” in blocks of congestion. Dots and dashes are depicted, respectively, as short and long line segments overlaid on the density profiles.

The results of the box objective are presented in Figures 3(b)-3(d). We give space-time diagrams for three different values of the parameter α . The box of the objective is shown as a black frame with an actual size of 10 miles and 20 minutes. As the trade-off moves from $\alpha = 0.3$ to 0.9, there is a clear increase in the congestion within the box, at the expense of allowing the congestion to spill outside the desired bounds. In fact, Figure 3(d) ($\alpha = 0.9$) activates the bottleneck near the top-left of the box earlier than Figure 3(b) ($\alpha = 0.3$) to congest the middle portion of the box, which leads to a propagation of a congestion wave outside the bounds of the bottom-right of the box.

Attack to Create Traffic Patterns in the Form of Morse Code.

- **Network** Since the I15 network does not have enough controllable onramps for the following attacks to be precise, we now consider a 60 mile freeway network with onramps and offramps spaced every 3.75 miles and a fixed demand on the onramps.
- **Attack** Figure 4 represents the space-time diagram of a *Morse code attack*. The objective is to create the Morse code representation of the three letters “C-A-L”², spelled with congestion blocks on the freeway. The corresponding objective function is the superposition of several box objectives on three thin time stripes of the space-time diagram. Everywhere else, the coefficients are put to zero. The result demonstrates that even with a reasonable number of ramps, one can achieve complex attack patterns. In particular, the optimal control approach was able to identify that creating a single backwards-moving jam was the most effective way to produce the second dash for “C”, the first dash for “A” and the first dot for “L”.

Arbitrary Patterns. Provided the right controllability conditions are satisfied, any congestion pattern may be created if the network has enough control ramps. To work towards this, we can choose the negative and positive coefficients of the *congestion-on-demand* method carefully to match a desired pattern. The following process, as depicted in Figure 5, gives a methodological approach to constructing arbitrary *congestion-on-demand* patterns.

One selects some image file they wish to reproduce in congestion patterns on a space-time diagram. The image is thresholded by color intensity to produce a bitmap of regions of desired congestion (X’s) and free-flow (O’s). Then a *congestion-on-demand* objective (Equation (23)) is constructed from the bitmap and scalarized using the α balance parameter to produce the a_i^k coefficients. A metering policy minimizing the objective is then computed using the optimal control method in Section 3.2. Given sufficient control of the network and optimization time, the resulting space-time diagram from the metering policy will resemble the input image file.

We give an example of the arbitrary *congestion-on-demand* attack in Figure 6, which produces a space-time diagram resembling the *Cal* logo. See [57] for a online video simulation of the *Cal attack*.

²Short for University of California.

Figure 5: Flow-chart for converting an arbitrary image to a *congestion-on-demand* goal. “Converting” an objective of the form in Equation 23 allows an attacker to compute metering rates that produce space-time diagrams resembling the original image.

Figure 6: Space-time diagram obtained following a *congestion-on-demand* attack with a Cal logo as the objective function. The attack was simulated on a 90 miles and 33-onramp freeway, for a 2 hours simulation time and using coordinated ramp metering.

4.2. Attack 2: catch-me-if-you-can

We will now show that the use of the multi-objective optimization methods introduced in Section 3.3 can allow the design of high-level, hard to define attacks. We will consider the example of a vehicle chase, presented in Section 2.3.2. Some vehicles are pursuing the driver along the freeway, while the driver wishes to escape. This objective is distinct from the *congestion-on-demand* attack, as our desired congestion pattern cannot immediately be imagined beforehand and is highly dependent upon the eventual path of the driver.

We translate the attack into a multi-objective problem (see Section 3.3). We can split this attack into four simpler and sometimes conflicting goals, each goal associated with an objective function to minimize:

1. The followers (everyone behind the driver) should travel along the freeway section as slowly as possible — Minimizing f_1 will maximize the traffic density of all freeway sections behind the driver’s trajectory.
2. In particular, those vehicles directly behind the driver should be impeded with increased priority — Minimizing f_2 will maximize the traffic density difference between the cells of the driver’s trajectory and those cells immediately behind.
3. As to not arouse suspicion from monitoring traffic managers, most other travel times should be reduced — Minimizing f_3 will reduce the total travel time of all the vehicles on the freeway to avoid unnecessary congestion.

Figure 7: Space-time diagram with a ternary graph representing the a_1, a_2, a_3 coefficients (here 30%, 55% and 15% respectively) used for the scalarization process in the catch-me-if-you-can example. The trajectory of the driver (blue line) appears to always gain distance in relation to pursuants further upstream (black lines). Best viewed in color.

4. The driver should quickly exit the freeway — Minimizing f_4 will reduce the driver's travel time, to allow him to travel along the freeway as quickly as possible and escape his or her followers.

Constructing a trajectory. f_2, f_3 and f_4 requires the trajectory of the driver, but reconstructing a vehicle's trajectory using a discretized, macroscopic traffic model is not obvious. We have chosen the following algorithm:

1. The driver's trajectory starts at $t = 0$ and in the first "spatial cell" of the freeway section.
2. The driver's current velocity is computed using the current cell's density.
3. The trajectory, assuming the current velocity, is projected to the next spatial cell.
4. If we are not at the end of the trajectory (in space or in time), we go back to step 2.

This algorithm only gives an approximation of the driver's trajectory, as some resolution is sacrificed in order to have a closed-form expression which permits computation of its partial derivatives. The details of the mathematical derivation and representation of the f_i objectives are omitted from the article.

We have four objective functions. In practice, presenting the results is clearer with only three functions, and we have chosen to keep only f_1, f_2 and f_3 in this article, as f_4 was not essential for producing interesting results. We will use the linear scalarization technique presented in Section 3.3, and choose three coefficients $a_1, a_2, a_3 \in \mathbb{R}_+$, so that $\sum_{i=1}^3 a_i = 1$. The objective function we want to optimize is then the following:

$$J(\mathbf{u}, \rho) = \sum_{i=1}^3 a_i f_i(\mathbf{u}, \rho) \quad (27)$$

4.2.1. Implementation

Graphical Representation. The space-time diagram in Figure 7, for a 21 miles freeway with 6 adjacent onramps and a 20 minutes simulation time, is an example output of the optimal control scalarization method. Such plots are useful for the human operator to discern between "good" and "bad" simulations produced from metering rates. The driver's

Figure 8: A grid exploration over the ternary graph. An optimization was conducted for a grid of coefficients regularly spaced on the ternary graph. The resulting scalarized objective is decomposed into the constituent objectives (normalized between 0 and 1) and plotted on separate summary ternary graphs.

trajectory is represented in blue, while the trajectory of three pursuants (a, b, c) are depicted losing ground on the driver once the driver escapes the initially congested region.

Ternary Graph. The triangle in Figure 7 is a visualization of the chosen set of coefficients a_i . The red dot represents the weighted average of the three corners of an equilateral triangle: the closer the red circle is to the a_i corner, the closer a_i is to 1. This is called a *ternary graph*. The top edge will always be a_1 , and the right and left a_2 and a_3 respectively. In this example, we can see that the dominant coefficients are a_1 and a_2 . As a consequence, we have an significant congestion behind the driver, forming immediately behind him.

A posteriori Method - Grid Exploration. Our approach for the a posteriori method is to automatically “explore the triangle of coefficients” to help the human operator find a preferred coefficient solution or region of solutions. Figure 8 presents the result of the a posteriori method. We plot the values of each objective function for the optimal solution associated with all sets of a_i coefficients. The lowest values of each f_i are always reached with the highest values of a_i (where f_i has been normalized to take values between 0 and 1; see Section 3.3). Any non-monotonicity in the graphs are attributed to early terminations of the optimizer’s gradient descent or convergence to sub-optimal local minima. The conflicting nature of the objectives is apparent. Figure 8(b) shows that f_1 is penalized more by high a_3 values than by high a_2 values, i.e., lowering the total travel time at the expense of congesting the region behind the driver.

The a posteriori method provides the human operator with a global overview of the Pareto front, enabling him to immediately locate a desired solution, or at least identify interesting starting points in the Pareto front. For example, Figure 8 gives an indication that the center regions of the triangles have large variations and should be explored further.

Interactive Method. A web application³ (diagram in Figure 9) was developed to allow a full exploration of the interactive method. A human operator first selects his or her desired coefficients (a_i) by clicking on the appropriate spot within triangle b). Then, after a scalarization using the particular coefficients and an optimization of the resultant objective, the interface plots the space-time diagram of the resulting simulation in window a), along with the driver’s trajectory. Any other vehicle’s trajectory can be visualized by clicking at the starting point of the desired trajectory. To enhance the exploration process, the interactive program also chooses two random (but nearby) sets of coefficients and plots their simulation in Figure 9 c1) and c2).

Figure 10 shows an overview of the results obtained while using the interactive interface. The first column shows simulations for the corners of the ternary graph, i.e., only one objective is active at a time. The results are intuitive in that optimizing f_1 (Figure 10.1) produces congestion everywhere behind the driver, optimizing f_2 (Figure 10.2)

³Interactive web application demo available at [57].

(a) Diagram of web application functionality.

(b) Actual web application [57] created for the purpose of this article. The triangle is replaced by 4 sliders, to match the 4 objective functions. Web application [57] is available online for the reader's convenience.

Figure 9: Interface of the interactive optimization system used to solve the multi-objective optimization problem to produce the attacks presented in the article.

creates a distinct increase in congestion behind the driver, and optimizing f_3 (Figure 10.3) maintains critical density everywhere, equivalent to maximizing throughput at maximum freeway speeds.

The second column (Figures 10.A-C) shows an interactive shift from favoring f_3 (minimize travel times) to favoring f_2 (trajectory boundary congestion). The shift progressively limits congestion formation, and intelligently removes more congestion *ahead* of the driver, as to not decrease the delay of pursuant vehicles.

The last column of Figure 10 demonstrates how the interactive process allows for fine-tuning of the balance of the objectives. Figure 10.a appears to be overly congested within the driver's trajectory. An interactive progression towards lower total travel times concludes with a desirable congestion boundary in Figure 10.c.

5. Conclusion

This article presents an analysis of the controllability of freeway traffic control systems and its ramifications in the context of physical and cyber-attacks. The impact of an attack is understood via the response of the control system, with direct attacks on the metering lights being potentially more effective than indirect attacks on the sensing infrastructure. Coordinated ramp metering attacks, being the highest level compromise considered, are extensively analyzed using methods from the fields of optimal control and multi-objective optimization. The mathematical approach to coordinated attacks on the freeway is explicitly derived for ramp metering applications. Detailed numerical simulations of coordinated ramp metering attacks were conducted to demonstrate the hazards of such compromises and the utility of optimal control tools in not only the hands of traffic managers, but also of adversaries.

As future work, we will develop methods that leverage knowledge of freeway dynamics to detect when a compromise of the traffic control system has occurred and how to mitigate the potential harm. For instance, as already demonstrated on water SCADA systems [35], one can detect when sensor readings lie outside those expected given the dynamical assumptions and classify such a sensor as faulty or compromised. We also wish to extend our method to freeway *calibration* [21] where one interactively optimizes over model parameters to allow model simulations to match recorded observations.

Figure 10: Summary of *catch-me-if-you-can* simulations generated via the interactive method. Column 1 shows optimizations over individual objectives. Column 2 shows a transition from favoring f_3 to favoring f_2 . Column 3 shows a progression across all three objectives.

Acknowledgements

This work was supported in part by FORCES (Foundations Of Resilient CybEr-Physical Systems), which receives support from the National Science Foundation (NSF award numbers CNS-1238959, CNS-1238962, CNS-1239054, CNS-1239166), and by the California Department of Transportation under the Connected Corridors program.

- [1] Z. Jia, C. Chen, B. Coifman, P. Varaiya, The PeMS algorithms for accurate, real-time estimates of g-factors and speeds from single-loop detectors, in: *Intelligent Transportation Systems, 2001. Proceedings. 2001 IEEE*, IEEE, 2001, pp. 536–541.
- [2] M. Papageorgiou, H. Hadj-Salem, J. Blosseville, ALINEA: A local feedback control law for on-ramp metering, *Transportation Research Record* 1320 (1991) 58–64.
- [3] D. Chen, S. Ahn, A. Hegyi, Variable speed limit control for steady and oscillatory queues at fixed freeway bottlenecks, *Transportation Research Part B: Methodological* 70 (2014) 340–358.
- [4] A. Kotsialos, M. Papageorgiou, F. Middelham, Optimal coordinated ramp metering with advanced motorway optimal control, *Transportation Research Record: Journal of the Transportation Research Board* 1748 (1) (2001) 55–65.
- [5] J. Reilly, S. Samaranyake, M. L. Delle Monache, W. Krichene, P. Goatin, A. Bayen, Adjoint-based Optimization on a Network of Discretized Scalar Conservation Laws with Applications to Coordinated Ramp Metering, *Journal of Optimization Theory and Applications* (under review).
- [6] S. Timotheou, C. G. Panayiotou, M. M. Polycarpou, *Transportation Systems: Monitoring, Control, and Security*, in: *Intelligent Monitoring, Control, and Security of Critical Infrastructure Systems*, Springer, 2015, pp. 125–166.
- [7] I. Papamichail, M. Papageorgiou, V. Vong, J. Gaffney, Heuristic ramp-metering coordination strategy implemented at Monash freeway, Australia, *Transportation Research Record: Journal of the Transportation Research Board* 2178 (1) (2010) 10–20.
- [8] D. Pisarski, C. Canudas-de Wit, Optimal balancing of freeway traffic density: Application to the grenoble south ring, in: *European Control Conference (ECC), IEEE*, 2013, pp. 4021–4026.
- [9] M. A. Miller, A. Skabardonis, San Diego I-15 Integrated Corridor Management (ICM) System: Stage II (Analysis, Modeling, and Simulation, California PATH Program, Institute of Transportation Studies, University of California at Berkeley, 2010.
- [10] D. B. Work, S. Blandin, O. P. Tossavainen, B. Piccoli, A. M. Bayen, A traffic model for velocity data assimilation, *Applied Mathematics Research eXpress* 2010 (1) (2010) 1.

- [11] L. Zhang, D. Levinson, Optimal freeway ramp control without origin–destination information, *Transportation Research Part B: Methodological* 38 (10) (2004) 869–887.
- [12] A. Hegyi, B. De Schutter, H. Hellendoorn, T. Van Den Boom, Optimal coordination of ramp metering and variable speed control-an mpc approach, in: *American Control Conference*, 2002. Proceedings of the 2002, Vol. 5, IEEE, 2002, pp. 3600–3605.
- [13] M. Ben-Akiva, D. Cuneo, M. Hasan, Evaluation of freeway control using a microscopic simulation laboratory, *Transportation Research Part C: Emerging Technologies* 11 (1) (2003) 29–50.
- [14] J. Haddad, M. Ramezani, N. Geroliminis, Cooperative traffic control of a mixed network with two urban regions and a freeway, *Transportation Research Part B: Methodological* 54 (2013) 17–36.
- [15] M. J. Cassidy, J. Rudjanakanoknad, Increasing the capacity of an isolated merge by metering its on-ramp, *Transportation Research Part B: Methodological* 39 (10) (2005) 896–913.
- [16] S. Ahn, R. L. Bertini, B. Auffray, J. H. Ross, O. Eshel, Evaluating benefits of systemwide adaptive ramp-metering strategy in portland, oregon, *Transportation Research Record: Journal of the Transportation Research Board* 2012 (1) (2007) 47–56.
- [17] C. Arthur MacCarley, S. P. Mattingly, M. G. McNally, D. Mezger, J. E. Moore, Field operational test of integrated freeway ramp metering/arterial adaptive signal control: lessons learned in irvine, california, *Transportation Research Record: Journal of the Transportation Research Board* 1811 (1) (2002) 76–83.
- [18] E. Smaragdis, M. Papageorgiou, E. Kosmatopoulos, A flow-maximizing adaptive local ramp metering strategy, *Transportation Research Part B: Methodological* 38 (3) (2004) 251–270.
- [19] A. Kotsialos, M. Papageorgiou, M. Mangeas, H. Haj-Salem, Coordinated and integrated control of motorway networks via non-linear optimal control, *Transportation Research Part C: Emerging Technologies* 10 (1) (2002) 65–84.
- [20] R. C. Carlson, I. Papamichail, M. Papageorgiou, A. Messmer, Optimal motorway traffic flow control involving variable speed limits and ramp metering, *Transportation Science* 44 (2) (2010) 238–253.
- [21] D. Jacquet, C. C. De Wit, D. Koenig, Traffic control and monitoring with a macroscopic model in the presence of strong congestion waves, in: *Decision and Control, 44th IEEE Conference on*, IEEE, 2005, pp. 2164–2169.
- [22] G. Yan, D. Wen, S. Olariu, M. C. Weigle, Security challenges in vehicular cloud computing, *Intelligent Transportation Systems, IEEE Transactions on* 14 (1) (2013) 284–294.
- [23] D. Ward, I. Ibarra, A. Ruddle, Threat analysis and risk assessment in automotive cyber security, Tech. rep., SAE Technical Paper (2013).
- [24] E. S. Canepa, C. G. Claudel, Spoofing cyber attack detection in probe-based traffic monitoring systems using mixed integer linear programming, in: *Computing, Networking and Communications (ICNC), 2013 International Conference on*, IEEE, 2013, pp. 327–333.
- [25] B. Ghena, W. Beyer, J. Hillaker, J. Pevarnek, J. A. Halderman, Green lights forever: Analyzing the security of traffic infrastructure, in: *8th USENIX Workshop on Offensive Technologies (WOOT’14)*, 2014.
- [26] T. Jeske, Floating car data from smartphones: What google and waze know about you and how hackers can control traffic, *Proc. of the BlackHat Europe*.
- [27] N. Tufnell, Students hack Waze, send in army of traffic bots, *wired.co.uk*.
URL <http://www.wired.co.uk/news/archive/2014-03/25/waze-hacked-fake-traffic-jam>
- [28] K. Zetter, Hackers Can Mess With Traffic Lights to Jam Roads and Reroute Cars, *wired.com*.
- [29] S. Grad, Engineers who hacked into L.A. traffic signal computer, jamming streets, sentenced, *Los Angeles Times*.
- [30] A. Muralidharan, R. Horowitz, Optimal control of freeway networks based on the link node cell transmission model, in: *American Control Conference (ACC)*, IEEE, 2012, pp. 5769–5774.
- [31] J. Reilly, A. M. Bayen, Distributed Optimization for Shared State Systems: Applications to Decentralized Freeway Control via Subnetwork Splitting, *IEEE Transactions on Intelligent Transportation Systems* (In Review).
- [32] A. K. Ziliaskopoulos, A linear programming model for the single destination system optimum dynamic traffic assignment problem, *Transportation science* 34 (1) (2000) 37.
- [33] S. Samaranayake, J. Reilly, W. Krichene, M. L. Delle Monache, P. Goatin, A. M. Bayen, Multi-commodity real-time dynamic traffic assignment with horizontal queuing., *Operations Research* (In Review).
- [34] H. K. Lo, W. Szeto, A cell-based variational inequality formulation of the dynamic user optimal assignment problem, *Transportation Research Part B: Methodological* 36 (5) (2002) 421–443.
- [35] S. Amin, X. Litrico, S. Sastry, A. M. Bayen, Cyber security of water scada systems—part I: analysis and experimentation of stealthy deception attacks, *Control Systems Technology, IEEE Transactions on* 21 (5) (2013) 1963–1970.
- [36] S. Amin, X. Litrico, S. S. Sastry, A. M. Bayen, Cyber security of water scada systems—part ii: Attack detection using enhanced hydrodynamic models, *Control Systems Technology, IEEE Transactions on* 21 (5) (2013) 1679–1693.
- [37] M. L. Delle Monache, J. Reilly, S. Samaranayake, W. Krichene, P. Goatin, A. M. Bayen, A PDE-ODE model for a junction with ramp buffer, *SIAM Journal on Applied Mathematics* 74 (1) (2014) 22–39.
- [38] AASHTO, ITE, NEMA, Model 2070 Controller Standard Version 03, Tech. rep. (2012).
- [39] FHWA, Type 170 Traffic Signal Controller System - Microcomputer Based Intersection Controller, Tech. rep., Federal Highway Administration (1978).
- [40] TomTom (2014).
URL <http://www.tomtom.com/>
- [41] J. Sutton, Copper wire stolen from traffic signal, street lights in Oklahoma City, *The Oklahoman*.
- [42] M. Rosenberg, Underground copper wire heist causes San Jose freeway flood, *San Jose Mercury News*.
- [43] M. J. Lighthill, G. B. Whitham, On kinematic waves. II. A theory of traffic flow on long crowded roads, *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 229 (1178) (1955) 317.
- [44] P. Richards, Shock waves on the highway, *Operations research* 4 (1) (1956) 42–51.
- [45] C. F. Daganzo, The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory, *Transportation Research Part B: Methodological* 28 (4) (1994) 269–287.
- [46] S. K. Godunov, A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics, *Matematicheskii*

- Sbornik 89 (3) (1959) 271–306.
- [47] J. P. Lebacque, The Godunov scheme and what it means for first order traffic flow models, in: Internaional symposium on transportation and traffic theory, 1996, pp. 647–677.
 - [48] D. Kroshko, et al., Openopt, Software package downloadable from <http://openopt.org>.
 - [49] N. Z. Shor, Minimization methods for non-differentiable functions, Vol. 3, Springer Science & Business Media, 2012.
 - [50] M. Gugat, M. Herty, A. Klar, G. Leugering, Optimal Control for Traffic Flow Networks, *Journal of Optimization Theory and Applications* 126 (3) (2005) 589–616. doi:10.1007/s10957-005-5499-z.
 - [51] M. B. Giles, S. Ulbrich, Convergence of linearized and adjoint approximations for discontinuous solutions of conservation laws. Part 2: Adjoint approximations and extensions, *SIAM Journal on Numerical Analysis* 48 (3) (2010) 905–921.
 - [52] S. Ulbrich, A sensitivity and adjoint calculus for discontinuous solutions of hyperbolic conservation laws with source terms, *SIAM journal on control and optimization* 41 (3) (2002) 740–797.
 - [53] M. B. Giles, N. A. Pierce, An introduction to the adjoint approach to design, *Flow, Turbulence and Combustion* 65 (3-4) (2000) 393–415. doi:10.1023/A:1011430410075.
 - [54] K. Deb, Multi-objective optimization using evolutionary algorithms, Vol. 16, John Wiley & Sons, 2001.
 - [55] D. L. Donoho, A. Maleki, I. U. Rahman, M. Shahram, V. Stodden, Reproducible research in computational harmonic analysis, *Computing in Science & Engineering* 11 (1) (2009) 8–18.
 - [56] V. Stodden, Enabling reproducible research: Licensing for scientific innovation, *Int'l J. Comm. L. & Pol'y* 13 (2009) 1–55.
 - [57] J. Reilly, S. Martin, SmartRoads Website (2014).
URL <http://traffic.berkeley.edu/smartroads>
 - [58] G. Dervisoglu, A. Kurzhanskiy, G. Gomes, R. Horowitz, Macroscopic freeway model calibration with partially observed data, a case study, in: American Control Conference (ACC), 2014, IEEE, 2014, pp. 3096–3103.
 - [59] A. Muralidharan, R. Horowitz, Imputation of Ramp Flow Data for Freeway Traffic Simulation, *Transportation Research Record: Journal of the Transportation Research Board* 2099 (-1) (2009) 58–64.