



On Cybersecurity of Freeway Control Systems: Analysis of Coordinated Ramp Metering Attacks

Jack Reilly

652 Sutardja Dai, Berkeley, CA 94720

Sébastien Martin

652 Sutardja Dai, Berkeley, CA 94720

Mathias Payer

305 N University Street, West Lafayette, IN 47907

Alexandre M. Bayen

642 Sutardja Dai, Berkeley, CA 94720

Abstract

This article focuses on cybersecurity of transportation systems and investigates their vulnerability to attacks on the sensing and control infrastructure. An array of different attack points, classified into *physical*, *close-proximity*, and *virtual* layers, are reviewed and investigated. We construct two benchmark *scenarios* which exploit these vulnerabilities to identify the potential harm of a traffic control system compromise. An in-depth analysis is then presented on the takeover of a series of networked onramp metering traffic lights. The analysis is conducted using a method for precise and intelligent onramp metering attacks based on finite-horizon optimal control techniques and multi-objective optimization. It relies on the use of a cell transmission model extended to include onramp buffers. The optimization process uses the adjoint method, directly applied to the cell transmission model. The method is demonstrated for two examples of high-level attack objectives: *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: cybersecurity, traffic optimization, ramp metering, 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

*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)

algorithms which regulate traffic flow to improve congestion [2]. 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 [3, 4].

Increased efforts to build systems which understand and utilize the interconnectivity are evidenced by *integrated-corridor-management* (ICM) projects such as *Connected Corridors* [5] and mobile applications which use GPS probe data to improve navigation [6].

This connectivity offers great potential to better analyze, control and manage traffic but also poses a significant security risk. 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.

A number of traffic-related attacks of infrastructure systems have already been demonstrated in the past few years. 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 [7], while a similar attack used a fleet of mobile phone emulators to mimic the presence of many virtual vehicles on a roadway [8]. 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 [9]. Cyber attacks on a centralized command center remain a serious threat given the frequent discovery of networking vulnerabilities, such as the Heartbleed bug [10]. 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 [11].

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 article, we present a systematic approach to analyzing the topic of traffic control system vulnerabilities and their potential impact.

To do so, we begin by constructing a taxonomy of different vulnerability locations in traffic control systems, defining three distinct layers: physical, close-proximity, and virtual. Difficulty, impact, and cost values are also associated with each potential attack. 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, ramp metering. To achieve this, we develop a method based on adjoint computations and finite-horizon optimal control for finding optimal metering rates to create a desired disruption on the freeway. 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.

The contributions of this article are as follows. We present a classification of a broad set of attacks on traffic control systems with their relation to the underlying physical and cyber infrastructure. Mathematical formulations based on optimal control and adjoint-based methods are used to show exactly how an attacker can exploit these weaknesses. Explicit algorithms using these tools for coordinated ramp metering attacks are derived and presented. Finally, we provide numerical evidence and novel results of the feasibility of these attacks via simulations modeled after actual freeway networks.

The rest of the article is organized as follows. Section 2 summarizes and classifies the vulnerabilities of traffic control systems. Section 3 gives a mathematical approach for carrying out a class of the presented attacks. Section 4 gives two detailed applications of the mathematical approach to ramp metering attacks. The first application shows how ramp metering can allow an attacker to cause congestion in precise locations and at precise moments in time along a freeway. Simulations are applied to a full-sized model of a 19.4 mile stretch of the I15 South Freeway in San Diego, California. Results are shown for both a custom macroscopic flow simulator as well as an Aimsun [12] microscopic model. The second application finds a strategy to solve the aforementioned problem of allowing a fleeing vehicles to escape pursuants. Numerical results are presented, as well as a discussion of the benefits of the multi-objective optimization method. We conclude with some future areas of study for traffic system security.

2. Traffic System Vulnerabilities

In the later part of the article we propose attacks to create congestion based on user-defined needs. This section reviews the current architecture of freeway control systems to show that these attacks can be implemented in practice

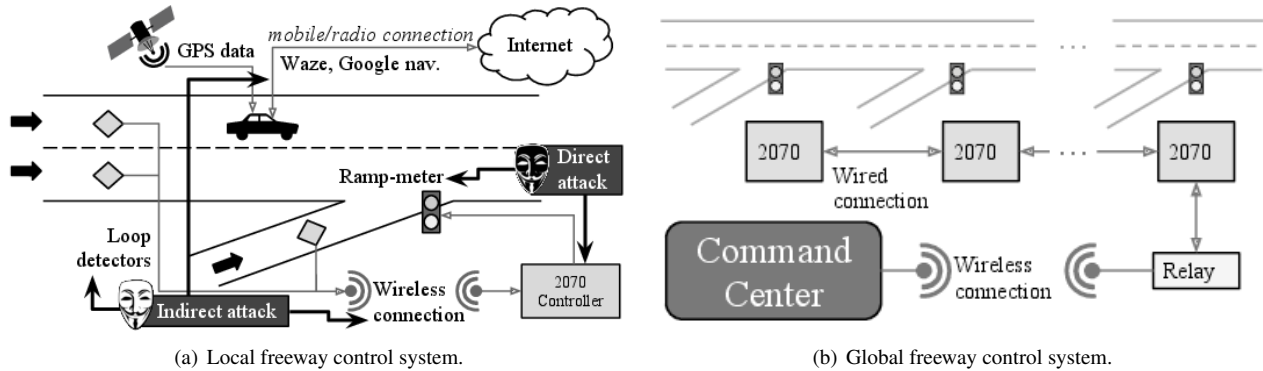


Figure 1: The physical roadway, sensors, connected vehicles and controllers near a freeway/onramp junction in Figure 1(a) form a cyber-physical network we refer to as a local freeway control system. The mask icons (white/black masks for indirect/direct vulnerabilities) denote vulnerability points in 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.

on such systems.

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 [6].
- 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 [13] 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 [14] and the older 170 boxes [15], which allows interaction between the sensors and ramp meters.

We assume control boxes are wired to the nearby metering light and have a wireless connection to nearby sensors. Vehicles with navigation devices such as TomTom [16] 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. 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

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	locality	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	locality	low	medium	low
replacing a corridor of control boxes	physical	high	medium	medium
attacking a corridor of control boxes	network	high	high	medium
attacking the control center	network	high	high	high
spoofing GPS data	network	medium	high	medium
attacking navigation software	network	medium	medium	medium

Table 1: List of possible infrastructure attacks with access to different layers that is needed, level of control that the attacker gains, sophistication of the attack, and cost.

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 [17, 18]. Such attacks need low sophistication, are easy to carry out, and are hard to protect against as each device must be physically protected given that software-based protection is not effective against physical attacks. On the other hand, the attack is costly as (i) direct physical access is needed, (ii) the attacker is exposed, and (iii) the attack does not scale (i.e., each piece of equipment is attacked individually). Examples of such an attack in Figure 1(a) include clipping or removing wires between sensors and the 2070 controller, tampering with individual sensors, the ramp meter, or the 2070 controller.

Proximity access (locality): Figure 1(b) depicts multiple control boxes chained together to form a corridor where actuators have a coordinated plan between the different control boxes. 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. Both aspects of the attack are feasible; the first aspect needs additional hardware and an attacker that delivers the hardware, the second aspect needs to find a software vulnerability with a security analysis of the existing devices. These attacks need higher sophistication and knowledge but no longer need direct hardware access to the existing sensors and scales to some extent.

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. Direct access to these centers is usually not given and this attack therefore is highly sophisticated (or needs insider access). This attack is the hardest possible attack as command centers and back links are usually guarded but allows a great scaling effect as many control boxes can be controlled directly.

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 anticipation that the uncompromised metering system will respond to the spoofed sensors in a predictable manner. 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 section¹. As the actual loop-detector sensors low-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 expensiveness 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.

The current traffic control architecture presented above supports the class of attacks described in the next section. Specifically, a mathematical approach to coordinated ramp metering attacks is developed to permit an attacker to effectively exploit vulnerabilities in the metering control system.

3. Theory for Coordinated Freeway Attacks

An attacker can negatively influence the performance of the freeway network or achieve some criminal goal by setting the metering lights to a particular configuration. The impact of such an attack can be maximized by leveraging a discrete dynamical freeway model to compute metering rates which achieve the desired goal using finite-horizon optimal control and multi-objective optimization techniques.

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 links². Flow dynamics along a link i is modeled using a discretized version of the *Lighthill-Whitham-Richards* [20, 21] (LWR) partial differential equation. The continuous LWR equation takes the following form:

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

¹See our link [19] for a Youtube video depiction. This was demonstrated in May 2014 at the White House-hosted SmartAmerica Conference.

²For 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).

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 [22]:

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

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

The discrete model used in this work is adapted from [23, 3] and was chosen for its suitability to ramp metering applications. Following [3], we discretize Equation (1) into cells of spatial size Δx and temporal size Δt using a *Godunov-based* or *cell-transmission-model* (CTM) scheme [24, 22, 25]. 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:

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

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

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

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

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

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

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

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

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. 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*, for solving such constrained optimization problems.

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)-(9) and moving all terms to the left-hand side, one can succinctly express the discrete, controllable dynamical system by:

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

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:

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

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

Gradient Computations via the Adjoint Method. 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.

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 derivative³ 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}, ρ):

³The 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 [26, 27, 28].

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

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* [29, 26], while the system of equations in (20) is referred as the *discrete adjoint system*. It is shown in [3] 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:

$$\begin{aligned}
 \frac{\partial \delta [i, k]}{\partial s} &= \begin{cases} v & s = \rho [i, k], \quad \rho [i, k] v \leq f^{\max} \\ 0 & \text{otherwise} \end{cases} \\
 \frac{\partial \sigma [i, k]}{\partial s} &= \begin{cases} -w & s = \rho [i, k], \quad w (\rho^{\max} - \rho [i, k]) \leq f^{\max} \\ 0 & \text{otherwise} \end{cases} \\
 \frac{\partial d [i, k]}{\partial s} &= \begin{cases} \frac{u[i,k]}{\Delta t} & s = l [i, k], \quad \frac{l[i,k]}{\Delta t} \leq r^{\max} \\ \min \left(\frac{l[i,k]}{\Delta t}, r^{\max} \right) & s = u [i, k] \\ 0 & \text{otherwise} \end{cases} \\
 \frac{\partial f^{\text{in}} [i, k]}{\partial s} &= \begin{cases} \beta [i, k] & s = \delta [i, k], \quad \sigma [i, k] \geq \delta [i, k] \beta [i, k] + d [i - 1, k] \\ 1 & s = d [i, k], \quad \sigma [i, k] \geq \delta [i, k] \beta [i, k] + d [i - 1, k] \\ 1 & s = \sigma [i, k], \quad \sigma [i, k] < \delta [i, k] \beta [i, k] + d [i - 1, k] \\ 0 & \text{otherwise} \end{cases} \\
 \frac{\partial f^{\text{out}} [i, k]}{\partial s} &= \begin{cases} \frac{\frac{\partial \delta [i, k]}{\partial s}}{\frac{\partial (f^{\text{in}} [i+1, k] - d [i+1, k]) / \beta [i, k]}{\partial s}} & \frac{p f^{\text{in}} [i+1, k]}{\beta [i, k] (1+p)} \geq \delta [i, k] \\ \frac{\frac{\partial (p f^{\text{in}} [i+1, k]) / (1+p) \beta [i, k]}{\partial s}}{\frac{f^{\text{in}} [i+1, k]}{1+p}} & \frac{f^{\text{in}} [i+1, k]}{1+p} \geq d [i + 1, k] \\ \text{otherwise} & \text{otherwise} \end{cases} \\
 \frac{\partial \rho [i, k + 1]}{\partial s} &= \begin{cases} 1 & s = \rho [i, k] \\ \frac{\Delta t}{\Delta x} & s = f^{\text{in}} [i, k] \\ -\frac{\Delta t}{\Delta x} & s = f^{\text{out}} [i, k] \\ 0 & \text{otherwise} \end{cases} \\
 \frac{\partial l [i, k + 1]}{\partial s} &= \begin{cases} 1 & s = l [i, k] \\ -\Delta t & s = r [i, k] \\ 0 & \text{otherwise} \end{cases}
 \end{aligned}$$

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 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.

Such a situation with multiple, competing objectives can be described as a *multi-objective optimization problem*.

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).

3.3.2. Decision Maker

A *Decision Maker* (DM) represents the human whose expertise is used to discern a preference between two control values. As we only wish to judge controls which are Pareto optimal, The DM only observes and discerns values on the Pareto front to limit the search space and improve the efficiency of the method. As a consequence, the DM has a hidden objective function: $u(\mathbf{u}, \rho)$, the *utility function*, which can only be indirectly observed through probing the DM. With u , we can reformulate the multi-objective optimization problem as:

$$\min_{x \in P} u(x) \quad (22)$$

The DM is essential to multi-objective optimization problems with large Pareto fronts. There are several ways to interact with him:

- He can evaluate his utility function u on any given Pareto-optimal solution.
- He can give more general preferences on the Pareto front, for example a preference for one of the objective functions, or for a given subset of the Pareto front.

3.3.3. Finite-horizon Optimal Control and Multi-objective Optimization

Scalarization. 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:

$$f(x) = \sum_{i \leq N} a_i f_i(x). \quad (23)$$

The DM can favor a specific objective f_i over other objectives by increasing the a_i coefficient.

It is easy to show that any solution of Equation (23) 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.

A Posteriori Method. Equation (23) allows one to sample the Pareto front by exploring the space of the coefficients which can provide to the DM a representative subset of Pareto-optimal solutions. The DM can then choose *a posteriori* his preferred solutions. And as such this method is called an *a posteriori method*.

This method can be computationally costly as many different optimal control problems need to be solved, but provides a good overview of the Pareto front. In particular, it gives an estimation of the lower and upper bounds of each objective function. Thus one can scale each objective function to take values only between 0 and 1, allowing the different objectives to be easily compared.

Interactive Method. Unlike with the a posteriori method, *Interactive methods* are based upon a repeated interaction with the Decision Maker.

1. The DM gives an indication of how to compute the next Pareto-optimal solution — for example, he may give an idea for the next set of coefficients (a_i) to use and his evaluation of the previous simulation.
2. The interactive scalarization process uses this indication to create a scalar objective — for example using Equation (23), we obtain a scalar objective with the set of coefficients given by the DM.

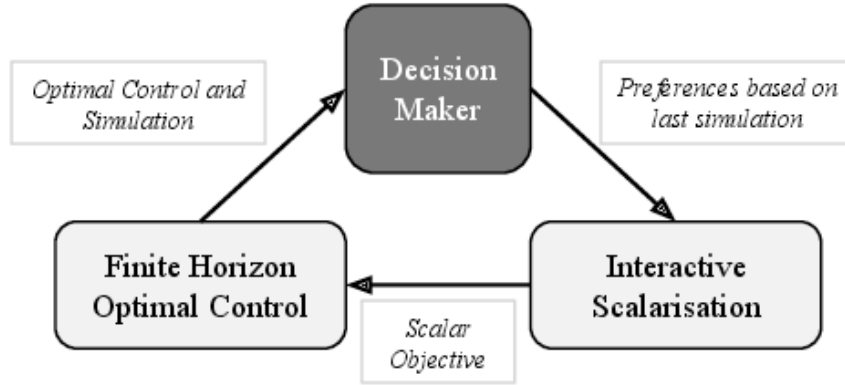


Figure 3: The interactive method for multi-objective optimization embeds the Decision Maker (DM) in the optimization loop, allowing the DM to direct the search of the Pareto front. The optimal controller adapts the advice of the DM to *scalarize* the multiple objectives and solve a new optimization problem. The results of the optimization are then fed back to the DM, and the cycle repeats until satisfaction.

3. The finite-horizon optimal control method is used to solve the corresponding optimization problem, and gives the result to the DM.

This process is repeated until the DM is satisfied with the results. Figure 3 shows the general process of interactive methods.

The important part of the interactive method is the kind of indications that can be given by the DM, and how the indications and the simulation history will be used in the scalarization process. Section 4.2 gives an example of an interactive method.

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 [30, 31], the software and data used to produce the numerical results and diagrams in this section is made available [19] to permit the reader to reproduce the presented results.

4.1. First Attack: congestion-on-demand

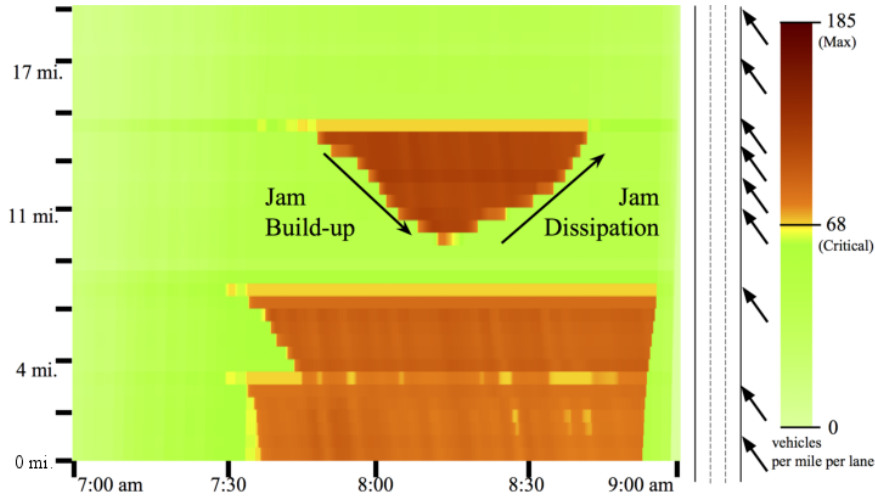
Congestion-on-demand describes a class of objectives where an attacker wishes to create congestion patterns of a specific nature. The attacks for the first example, *box objective* (to be described), use a macroscopic freeway model of a 19.4 mile stretch of the I15 South Freeway in San Diego California. The model was split into 125 links with 9 onramps and was calibrated [32, 33] using loop-detector measurements available through the PeMS loop-detector system [1]. Figure 4(a) is a *Space-time diagram* of the I15 freeway. There is no ramp metering control applied to the simulation in Figure 4(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:

$$J(\mathbf{u}, \rho) = \sum_{i=1}^N \sum_{k=1}^T a_i^k \rho[i, k] \quad (24)$$



(a) Simulation with no metering.




Figure 4: Figure 4(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 4(b)–4(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 (27)) enables the proper design of tradeoffs in the objective.

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.

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:

$$f_1(\mathbf{u}, \rho) = - \sum_{(i,k) \in \text{Box}} \rho[i, k] \quad (25)$$

$$\text{and } f_2(\mathbf{u}, \rho) = \sum_{(i,k) \notin \text{Box}} \rho[i, k] \quad (26)$$

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:

$$J_\alpha(\mathbf{u}, \rho) = \alpha f_1(\mathbf{u}, \rho) + (1 - \alpha) f_2(\mathbf{u}, \rho), \quad (27)$$

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.

The results of the box objective are presented in Figures 4(b)–4(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 4(d) ($\alpha = 0.9$) activates the bottleneck near the top-left of the box earlier than Figure 4(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.

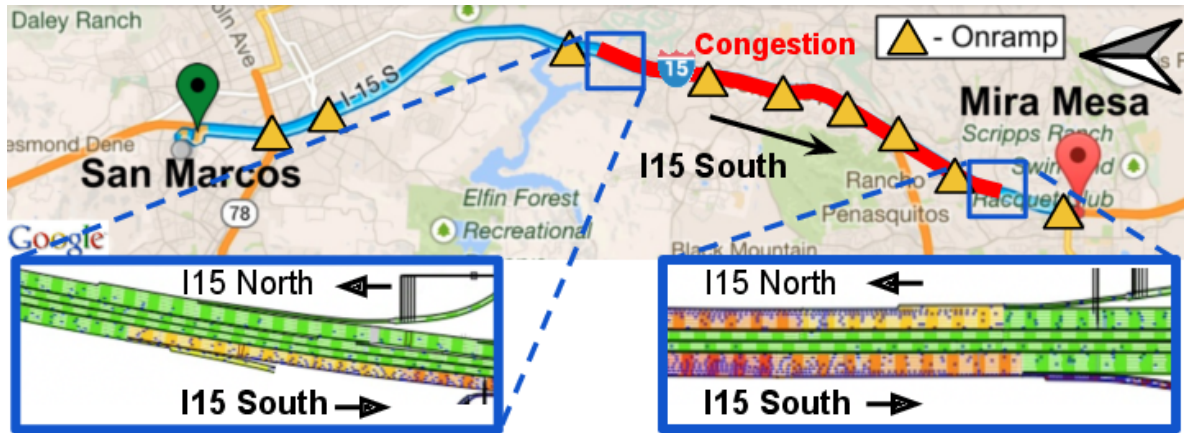


Figure 5: Box objective attack implemented on a microscopic model of I15 Freeway produced with Aimsun software. The metering lights were set using the *congestion-on-demand* strategy. Snapshots of traffic at the north and south extents of the box show that the strategy maintains congestion within the box and free-flow conditions outside the box. A link to a video of the microsimulation is provided [19]. Best viewed in color.

As an illustration of the practical nature of coordinated ramp metering attacks, we also implemented the box objective on an Aimsun microsimulation model [12] of the I15 freeway network. This model originates from the I15 integrated corridor management project ran in San Diego in 2010, ref [5]. The geographical location of the I15 network is given in Figure 5 and shows San Marcos as the southbound start and Mira Mesa as the end, with the desired box of congestion placed approximately 5 miles before Mira Mesa. A snapshot of the northern and southern extents of the box at the time of 8:30 are shown below the map. The south-bound lanes in the snapshot indicate that congestion was more or less confined to the desired box. A summary video [19] of the I15 microsimulation shows the formation and dissipation of congestion within the predetermined freeway section.

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 6 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”⁴, 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”.

⁴Short for University of California.

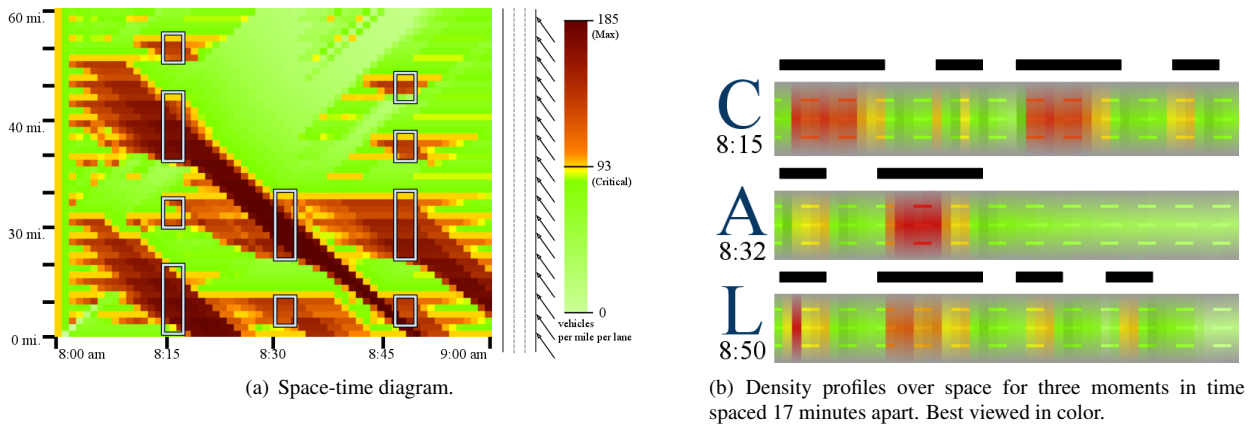


Figure 6: Attack to create traffic patterns in the form of Morse code. A coordinated ramp metering attack using the proposed algorithm is able to spell out “C-A-L” in Morse code over successive time-slices of the space-time diagram: C= - · - ·, A= · - ·, L= · - · - ·. The entire space-time diagram of the attack is shown in Figure 6(a), while three snapshots of the freeway are shown in Figure 6(b), each slice spelling out one of the letters in “C-A-L” in blocks of congestion.




Figure 7: Flow-chart for converting an arbitrary image to a *congestion-on-demand* goal. “Converting” an objective of the form in Equation 24 allows an attacker to compute metering rates that produce space-time diagrams resembling the original image.

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 7, 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 (24)) 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 8, which produces a space-time diagram resembling the *Cal* logo. See [19] for an online video simulation of the *Cal* attack.

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 more realistic and 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:




Figure 8: 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.

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.
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 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.

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 chose 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:

$$J(\mathbf{u}, \rho) = \sum_{i=1}^3 a_i f_i(\mathbf{u}, \rho) \quad (28)$$

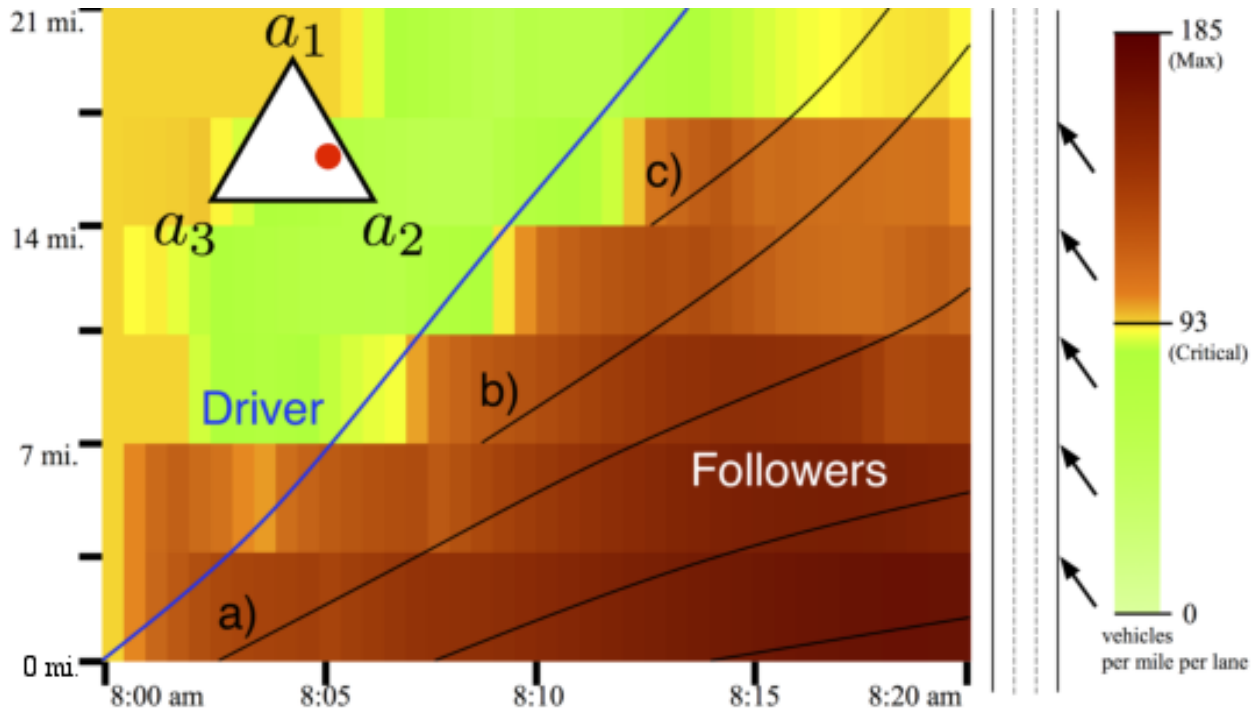


Figure 9: Space-time diagram with a ternary graph representing the a_1, a_2, a_3 coefficients (here 30%, 55% and 15% respectively) used for the scalarization process in the catch-me-if-you-can example. The trajectory of the driver (blue line) appears to always gain distance in relation to pursuants further upstream (black lines). Best viewed in color.

4.2.1. Implementation

Graphical Representation. The space-time diagram in Figure 9, 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 DM to discern between “good” and “bad” simulations produced from metering rates. The driver’s trajectory is represented in blue, while the trajectory of three pursuants (a, b, c) are depicted losing ground on the driver.

Ternary Graph. The triangle in Figure 9 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 *Decision Maker* find a preferred coefficient solution or region of solutions. Figure 10 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 10(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 DM 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 10 gives an indication that the center regions of the triangles have large variations and should be explored further.

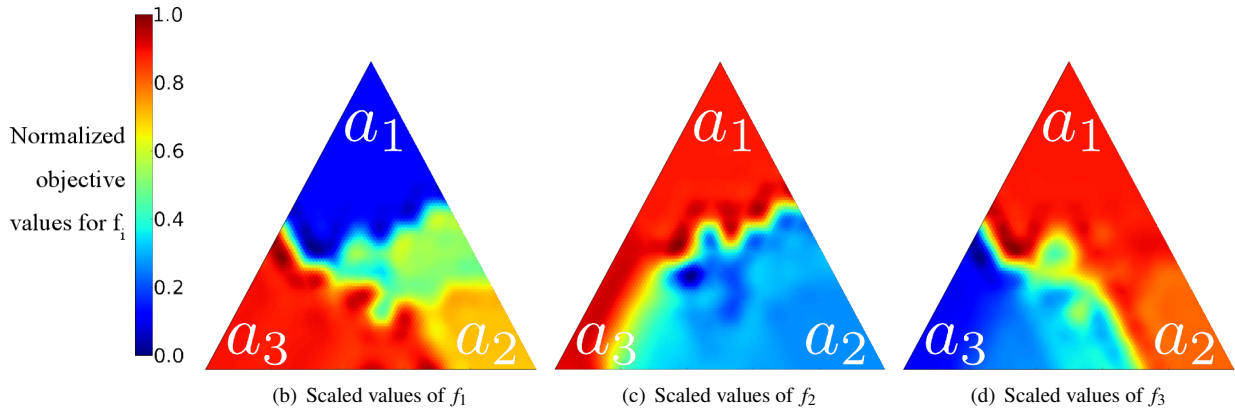


Figure 10: A grid exploration over the ternary graph. An optimization was conducted for a grid of coefficients regularly spaced on the ternary graph. The resulting scalarized objective is decomposed into the constituent objectives (normalized between 0 and 1) and plotted on separate summary ternary graphs.

Interactive Method. A web application⁵ (diagram in Figure 11) was developed to allow a full exploration of the interactive method. The DM first selects his 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 c1) and c2).

Figure 12 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 12.1) produces congestion everywhere behind the driver, optimizing f_2 (Figure 12.2) creates a distinct increase in congestion behind the driver, and optimizing f_3 (Figure 12.3) maintains critical density everywhere, equivalent to maximizing throughput at maximum freeway speeds.

The second column (Figures 12.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.

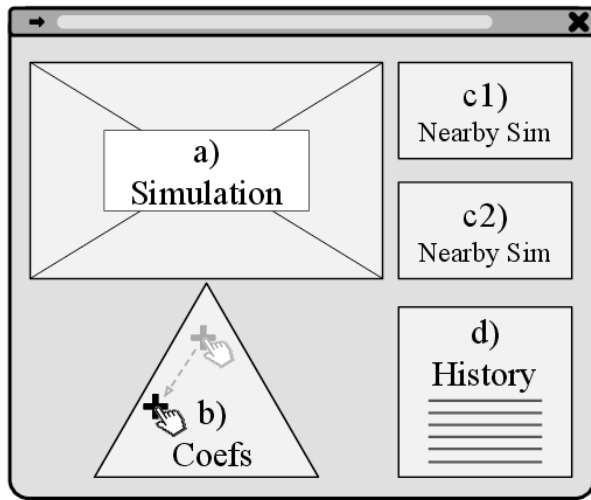
The last column of Figure 12 demonstrates how the interactive process allows for fine-tuning of the balance of the objectives. Figure 12.a appears to be overly congested within the driver's trajectory. An interactive progression towards lower total travel times concludes with a desirable congestion boundary in Figure 12.c.

5. Conclusion

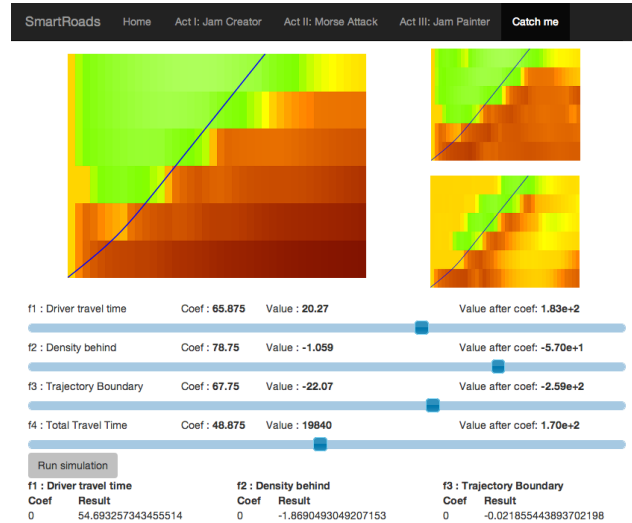
This article presents an overview of freeway traffic control systems and their vulnerability to physical and cyber-attacks. The impact of an attack is understood via the response of the control system, with direct attacks on the metering lights being potentially more effective than indirect attacks on the sensing infrastructure. Coordinated ramp metering attacks, being the highest level compromise, 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

⁵Interactive web application demo available at [19].



(a) Diagram of web application functionality.



(b) Actual web application [19] created for the purpose of this article. The triangle is replaced by 4 sliders, to match the 4 objective functions. Web application [19] is available online for the reader's convenience.

Figure 11: Interface of the interactive optimization system used to solve the multi-objective optimization problem to produce the attacks presented in the article.

demonstrated on water SCADA systems [34], one can detect when sensor readings lie outside those expected given the dynamical assumptions and classify such a sensor as faulty or compromised.

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. The authors would like to thank Dimitris Triantafyllos and others in Transport Simulation Systems (TSS) for their collaboration on the Aimsun microsimulations included within the article.

- [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] J. Reilly, W. Krichene, M. L. Delle Monache, S. Samaranyake, P. Goatin, A. Bayen, Adjoint-based optimization on a network of discretized scalar conservation law PDEs with applications to coordinated ramp metering, *Journal of Optimization Theory and Applications* (under review).
- [4] 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.
- [5] 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.
- [6] 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.
- [7] 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*.
- [8] 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>
- [9] K. Zetter, Hackers Can Mess With Traffic Lights to Jam Roads and Reroute Cars, *wired.com*.
- [10] Codenomicon, The Heartbleed Bug (2014).
URL www.heartbleed.com
- [11] S. Grad, Engineers who hacked into L.A. traffic signal computer, jamming streets, sentenced, *Los Angeles Times*.

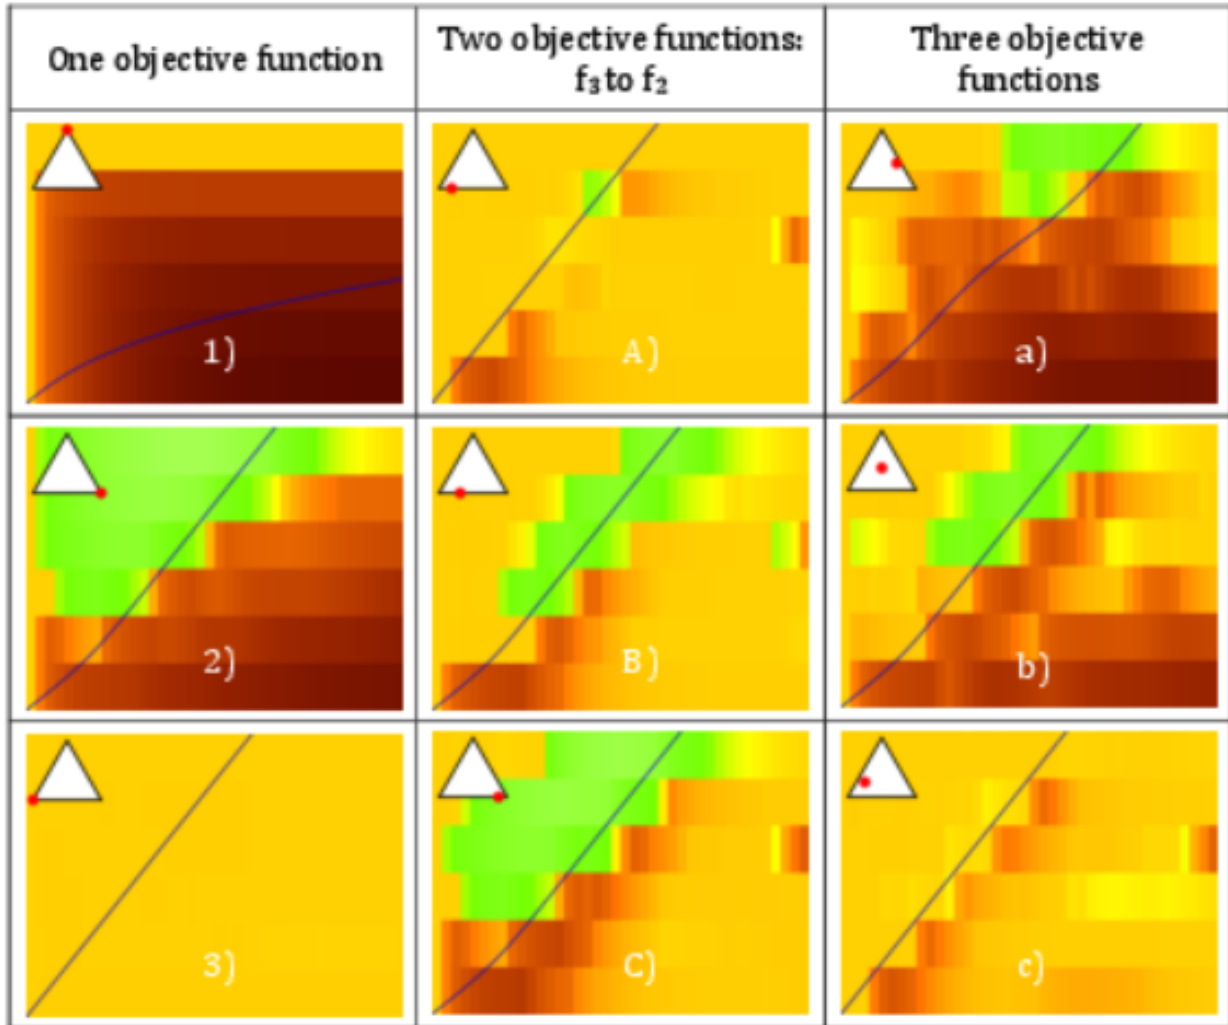


Figure 12: Summary of *catch-me-if-you-can* simulations generated via the interactive method. Column 1 shows optimizations over individual objectives. Column 2 shows a transition from favoring f_3 to favoring f_2 . Column 3 shows a progression across all three objectives.

[12] J. Barceló, Microscopic traffic simulation: A tool for the analysis and assessment of its systems, in: Highway Capacity Committee, Half Year Meeting., TSS - Transport Simulation System, Barcelona, 2002.

[13] 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.

[14] AASHTO, ITE, NEMA, Model 2070 Controller Standard Version 03, Tech. rep. (2012).

[15] FHWA, Type 170 Traffic Signal Controller System - Microcomputer Based Intersection Controller, Tech. rep., Federal Highway Administration (1978).

[16] TomTom (2014).
URL <http://www.tomtom.com/>

[17] J. Sutton, Copper wire stolen from traffic signal, street lights in Oklahoma City, The Oklahoman.

[18] M. Rosenberg, Underground copper wire heist causes San Jose freeway flood, San Jose Mercury News.

[19] J. Reilly, S. Martin, SmartRoads Website (2014).
URL <http://traffic.berkeley.edu/smartroads>

[20] 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.

[21] P. Richards, Shock waves on the highway, Operations research 4 (1) (1956) 42–51.

[22] 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.

[23] 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.
- [24] S. K. Godunov, A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics, *Matematicheskii Sbornik* 89 (3) (1959) 271–306.
- [25] J. P. Lebacque, The Godunov scheme and what it means for first order traffic flow models, in: *International symposium on transportation and traffic theory*, 1996, pp. 647–677.
- [26] 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.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] V. Stodden, Enabling reproducible research: Licensing for scientific innovation, *Int'l J. Comm. L. & Pol'y* 13 (2009) 1–55.
- [32] 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.
- [33] 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.
- [34] 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.