Decentralization and Security in Dynamic Traffic Light Control

(Extended Abstract)

Jian Lou and Yevgeniy Vorobeychik
Electrical Engineering and Computer Science
Vanderbilt University
Nashville, TN
{jian.lou, yevgeniy.vorobeychik}@vanderbilt.edu

ABSTRACT

Complex traffic networks include a number of controlled intersections, and, commonly, multiple districts or municipalities. The result is that the overall traffic control problem is extremely complex computationally. Moreover, given that different municipalities may have distinct, non-aligned, interests, traffic light controller design is inherently decentralized, a consideration that is almost entirely absent from related literature. Both complexity and decentralization have great bearing both on the quality of the traffic network overall, as well as on its security. We consider both of these issues in a dynamic traffic network. First, we propose an effective local search algorithm to efficiently design system-wide control logic for a collection of intersections. Second, we propose a game theoretic (Stackelberg game) model of traffic network security in which an attacker can deploy denial-of-service attacks on sensors, and develop a resilient control algorithm to mitigate such threats. Finally, we propose a game theoretic model of decentralization, and investigate this model both in the context of baseline traffic network design, as well as resilient design accounting for attacks. Our methods are implemented and evaluated using a simple traffic network scenario in SUMO.

CCS Concepts

•Computing methodologies → Distributed artificial intelligence; •Security and privacy → Embedded systems security; •Computer systems organization → Sensors and actuators;

Keywords

Traffic Control System; Game Theoretical Model; Decentralization and Security; Simulation-Based Method

1. INTRODUCTION

Effective design of large-scale complex traffic control systems, involving many controlled intersections, is fundamental in modern urban centers. As a result, this problem has been considered extensively in prior literature spanning fields such as transportation, operations research, economics, and computer science. Although adaptive, state-aware strategies can offer tremendous gains in traffic control efficiency, they expose an attack surface that can be exploited to substantially increase congestion. For example, a common kind of adaptive control logic involves state captured by vehicle queue lengths in each direction, with light switching between red and green as a function of relative queue lengths. While such state-aware switching can significantly increase efficiency, they also expose a vulnerability of controllers to attacks on sensors from which queue length information is derived.

An additional consideration which is crucial in modern complex traffic networks is that traffic lights on the network are often designed by multiple actors (e.g., municipalities). Consequently, while in principle we may be able to design extremely efficient and resilient controllers for a particular traffic network, this is impractical due to misalignment of interests among the different parties that actually control such networks.

We propose to systematically address the problems described above by considering a multi-intersection scenario in which a) traffic light controllers take into account relative queue lengths to determine red-green state of the traffic lights at an intersection, b) controllers for all lights must be designed to work jointly so as to optimize overall traffic network performance, c) sensors feeding data into the controllers are vulnerable to denial-of-service attacks, and d) intersections are partitioned among a set of players, with own goals pertaining to congestion within their local municipal region, which are in general misaligned with global interests of the entire traffic network. In our paper, we use the "Simulation of Urban MObility" (SUMO) [2] platform to implement, illustrate, and evaluate our approach.

2. TRAFFIC NETWORK MODEL

In the section, we introduce the control logic we use in our paper, and define the metrics to measure the efficiency of a traffic system. The control logic is adapted and revised from [1]. However, they only considered a single-intersection scenario. In our paper, we will generalize the control logic into cases with multiple intersections and correspondingly multiple traffic lights.

2.1 Traffic Control System

Consider a traffic network consisting of n intersections I_1, I_2, $\ldots$, I_n. We assume that each intersection is a cross of two "one-way" roads, and has no left turns.[1] In addition, we assume (as is

[1]Allowing for "two-way" streets and left turns is a relatively straightforward generalization.

HotSoS '16 April 19-21, 2016, Pittsburgh, PA, USA

ACM ISBN 978-1-4503-4277-3/16/04.

DOI: http://dx.doi.org/10.1145/2898375.2898384

common) that yellow light cycles are counted as a part of red lights cycles. Each direction j ($j = 1$, or 2) of intersection I_i ($1 \leq i \leq n$) has an exogenously specified minimum *green* light cycle length $\Omega_{i,j,min}$ and maximum *green* light cycle length $\Omega_{i,j,max}$. We assume that each intersection I_i has two sensors in each direction j allowing us to count the number of vehicles, $m_{i,j}(t)$, queued at that intersection in direction j (specifically, an ingress sensor counts incoming vehicles, and an egress outgoing vehicles, with the difference giving us the queue length). For each direction j, we also define a clock variable $c_{i,j}(t)$, which measures the time since the last switch from *red* to *green* of the traffic light for direction j.

2.2 Controllers

For a given intersection I_i, we adopt a two-parameter control logic model from [1], which determines behavior based on a comparison of queue lengths $m_{i,j}(t)$ and corresponding thresholds s_{ij}. Intuitively, when queue length in a particular direction j exceeds the corresponding threshold, this direction is viewed as high-priority and congested.

Consider a set of n intersections, $\{I_1, \ldots, I_n\}$, and associated threshold parameters

$$\mathbf{s} = \{\langle s_{1,1}, s_{1,2}\rangle, \langle s_{2,1}, s_{2,2}\rangle, \ldots, \langle s_{n,1}, s_{n,2}\rangle\}$$

in which $s_{i,j} \in \mathbb{R}^+$ ($1 \leq i \leq n, j = 1, 2$).

2.3 Objective: Weighted Average Latency

Assume there is a vehicle set V in the system, s.t. $|V| = d$. We assume $V = \{v_1, v_2, \ldots, v_d\}$, and every vehicle v_i ($1 \leq i \leq d$) has a corresponding weight w_i which denotes the relative importance of the vehicle. For instance, an ambulance may have higher weight than a common personal car. For each vehicle v_i travelling in the traffic system, *latency* l_i measures the time consumed for the car from entering the system to leaving the system. We can define *Weighted Average Latency (denote as $\mathcal{L}$)* as follows, and minimizing it is also a main goal for the manager of the system.

$$\mathcal{L} = \frac{\sum_{i=1}^{d} w_i l_i}{\sum_{i=1}^{d} w_i}$$

3. DECENTRALIZATION AND SECURITY

3.1 Game Theoretic Model

We now present a game theoretical model, in which multiple defenders determine configurations non-overlapping subsets of traffic lights. There may or may not be an attacker. If the attacker is not considered, we view it as a baseline *decentralized control game*, whereas consideration of an attacker extends the model to a *resilient decentralized control game*.

Formally, assume there is a set of Defenders D ($|D| \leq n$) who are in charge of different districts and corresponding traffic lights in a traffic network. Each defender d is only concerned about the Weighted Average Latency $\mathcal{L}_d$ for her own district d. Let $\mathbf{s}_d$ be the set of parameters controlled by defender $d \in D$, then $\mathbf{s}_d \cap \mathbf{s}_{\mathbf{d}'} = \emptyset$ for $d \neq d'$, and $\bigcup_d \mathbf{s}_d = \mathbf{s}$. Assume there is an attacker $\mathcal{A}$ who can attack a sensor in the system, and her goal is to increase the overall $\mathcal{L}$ of the system. We define see it as a *Multi-Defender Traffic Control Game* in which each defender want to decrease the latency in her own district, and attacker want to increase the latency for the overall system.

3.2 Optimization Problem

If there is only one defender and no attacker in the system, then our goal is to choose the parameters of all intersections $\mathbf{s}$ so as to minimize overall weighted latency, for a given weight vector w:

$$\min_{\mathbf{s}} \mathcal{L}(\mathbf{s}; w). \tag{1}$$

The optimization problem in Equation 1 is intractable because the objective function is a challenge to evaluate even for a fixed parameter vector $\mathbf{s}$, let alone optimize (typically, as below, it is evaluated by running simulations). Rather than exhaustive search, we propose a *coordinate greedy* (or just *CGA*) local search method for efficiently computing an approximately optimal configuration $\mathbf{s}$. The proposed algorithm, Algorithm 1, works by first discretizing parameters for each traffic light i, and then iteratively choosing a particular traffic light, and finding the optimal configuration of parameters of this light, *keeping configuration of the rest fixed*.

Algorithm 1 Coordinate Greedy Algorithm (CGA)

input: Starting Parameter set $\hat{\mathbf{s}}$
return: Local Minimal Parameter set $\mathbf{s}^*$

1: Copy $\hat{\mathbf{s}}$ to $\mathbf{s}^*$
2: **while** There exists an intersection, such that we could change parameters of the intersection to make $\mathcal{L}$ smaller **do**
3: Make the change to $\mathbf{s}^*$
4: **end while**
5: **return** $\mathbf{s}^*$

3.3 Resilient Control Algorithm

If there is only one defender and one attacker in the system, the goal of resilient network control is to choose $\mathbf{s}^*$ which is best response given attacker plays best response. Since, again, exhaustive search is clearly intractable, we propose an augmented version of the CGA algorithm, shown as Algorithm 2 (RCGA).

Algorithm 2 Resilient CGA (RCGA)

input: Local Optimal Parameter set $\hat{\mathbf{s}}$ given no Attacker
return: Resilient Parameter set $\mathbf{s}^*$

1: Given $\hat{\mathbf{s}}$ is applied in the system, enumerate sensors and find a sensor α, by attacking it we could get maximal $\mathcal{L}$.
2: Given α is attacked, we run Algorithm 1 starting from $\hat{\mathbf{s}}$
3: **return** resulting parameter set $\mathbf{s}^*$

3.4 Approximating Equilibrium in Decentralized Control

When there are multiple defenders and no defender, we consider the game as a Normal-Form game among multiple agents, and our goal is to compute or approximate a *Nash equilibria* among them. In the equilibrium, there is no defender who can switch her configuration to get a higher payoff.

Since computing a Nash equilibrium in our setting is intractable, we propose a simple iterative best response algorithm (Algorithm 3), in which each traffic light is chosen in a given iteration, and the associated defender d optimizes parameters of this traffic light only to minimize $\mathcal{L}_d$, fixing all other parameters. We refer to this algorithm as *BRA* (best response algorithm).

3.5 Approximating Equilibrium in Resilient Decentralized Control

Finally, we consider the decentralized setting, but now allowing for an attacker who will optimally respond to the joint configuration

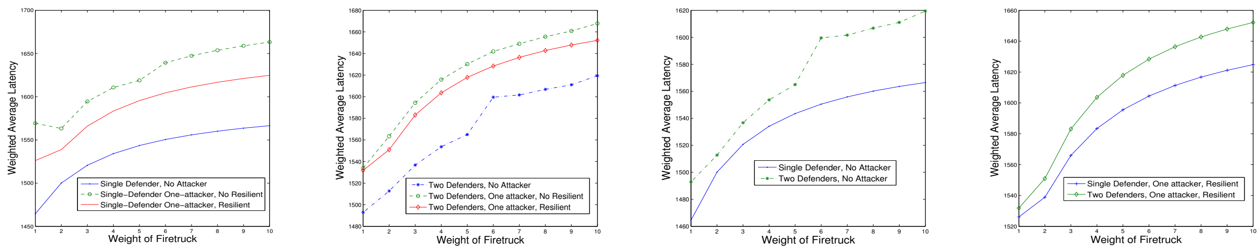

(a) (b) (c) (d)

Figure 1: Comparison between single-defender, no attacker (baseline) configuration, resilient single-defender configuration, as well as decentralized solutions.

Algorithm 3 Best Response Algorithm (BRA)

input: Starting Parameter set $\hat{\mathbf{s}}$
return: Equilibrium Parameter set $\mathbf{s}^*$

Copy $\hat{\mathbf{s}}$ to $\mathbf{s}^*$
while There exists an defender d, such that we could change parameters of the intersection controlled by defender d to make $\mathcal{L}_d$ smaller **do**
    Make the change to $\mathbf{s}^*$
end while
return $\mathbf{s}^*$

of all traffic lights by all defenders, $\mathbf{s}$. Formally, each Defender d wants to minimize $\mathcal{L}_d$ to improve resilience, and attacker want to maximize overall $\mathcal{L}$ by attacking a sensor.

Then we could easily extend *BRA* (we omit it due to limit of space), in which each best response iteration now accounts for the attacker's sensor DoS attack strategy.

4. EVALUATION AND RESULTS

To implement the traffic control algorithm and perform simulation, we employ a simulation suit called SUMO (short for "Simulation of Urban MObility"). SUMO [2] is an open source, highly portable, microscopic road traffic simulation package designed to handle large road networks. SUMO also provides a Traffic Control Interface (TraCI) to let external controllers control the traffic. In our work, we use a Python script to control the simulation through TraCI and implement our control algorithm.

In the paper, we consider an Emergency Vehicle Scenario. In the scenario, there are some common cars traveling from west to east, and some emergency vehicles (firetrucks) traveling from north to south. Assume that common cars have weight 1, and emergency vehicles have higher weights. There are some traffic lights that can be controlled in the intersections of the scenario. Before each direction in an intersection, there are two sensors that count the number of vehicles.

The experiment results can be seen in Figure 1, which shows the overall Weighted Average Latency $\mathcal{L}$ as a function of firetruck weights. When there is a single defender (Figure 1(a)) and no attacker, we obtain a relatively low $\mathcal{L}$ (applying Algorithm 1). However, the figure shows that an attack on the non-resilient configuration can substantially elevate $\mathcal{L}$: ignoring the possibility of a DoS attack can be disastrous for traffic in this scenario. On the other hand, resilient configuration (applying Algorithm 2) performs substantially better under attack.

Next, we split the scenario into two parts, in upper part, defender is in charge of the upper two intersections and another defender is in charge of the lower three intersections. And each defender only cares about Weighted Average Latency of her own district. The result is shown in Figure 1(b). As we can observe, considering resilience is beneficial for the defenders.

From Figures 1(c) and 1(d), we can observe how decentralization impacts the efficiency of a system. By comparing single-defender cases and two-defender cases, we find that the overall $\mathcal{L}$ in two-defender cases is higher than that in single defender cases, with and without attacker. It comes from the *negative externalities* introduced to the system when there are multiple selfish defenders, which make the overall system behavior inefficient.

5. CONCLUSION AND FUTURE WORK

We considered decentralization and security issues in dynamic traffic light control. We proposed a game theoretic model and simulation-based optimization and equilibrium approximation algorithms to address the problem. We then implemented and evaluated our algorithms on the SUMO platform.

There are a number of future research directions that can be considered. One such direction is to investigate scalability of our approach to significantly larger and more complex scenarios. Additionally, we only consider DoS attack on sensors. In future work, it will be important to evaluate resilience in the context of integrity attacks as well.

6. ACKNOWLEDGEMENT

This work was supported in part by System Science of SecUrity and REsilience for cyber-physical systems (SURE), which receives support from the Air Force Research Laboratory under award FA8785-14-2-0180, National Science Foundation under award CNS-1238959 and IIS-1526860, the Office of Naval Research under Award N00014-15-1-2621, Army Research Office under Award W911NF-16-1-0069 and Sandia National Laboratories.

7. REFERENCES

[1] J. L. Fleck and C. G. Cassandras. Infinitesimal perturbation analysis for quasi-dynamic traffic light controllers. In *12th International Workshop on Discrete Event Systems, WODES 2014, Cachan, France, May 14-16, 2014.*, pages 235–240, 2014.

[2] J. E. D. K. Michael Behrisch, Laura Bieker. Sumo – simulation of urban mobility sumo. In *The Third International Conference on Advances in System Simulation*, 2011.