Optimizing Traffic Flow in Urban Settings

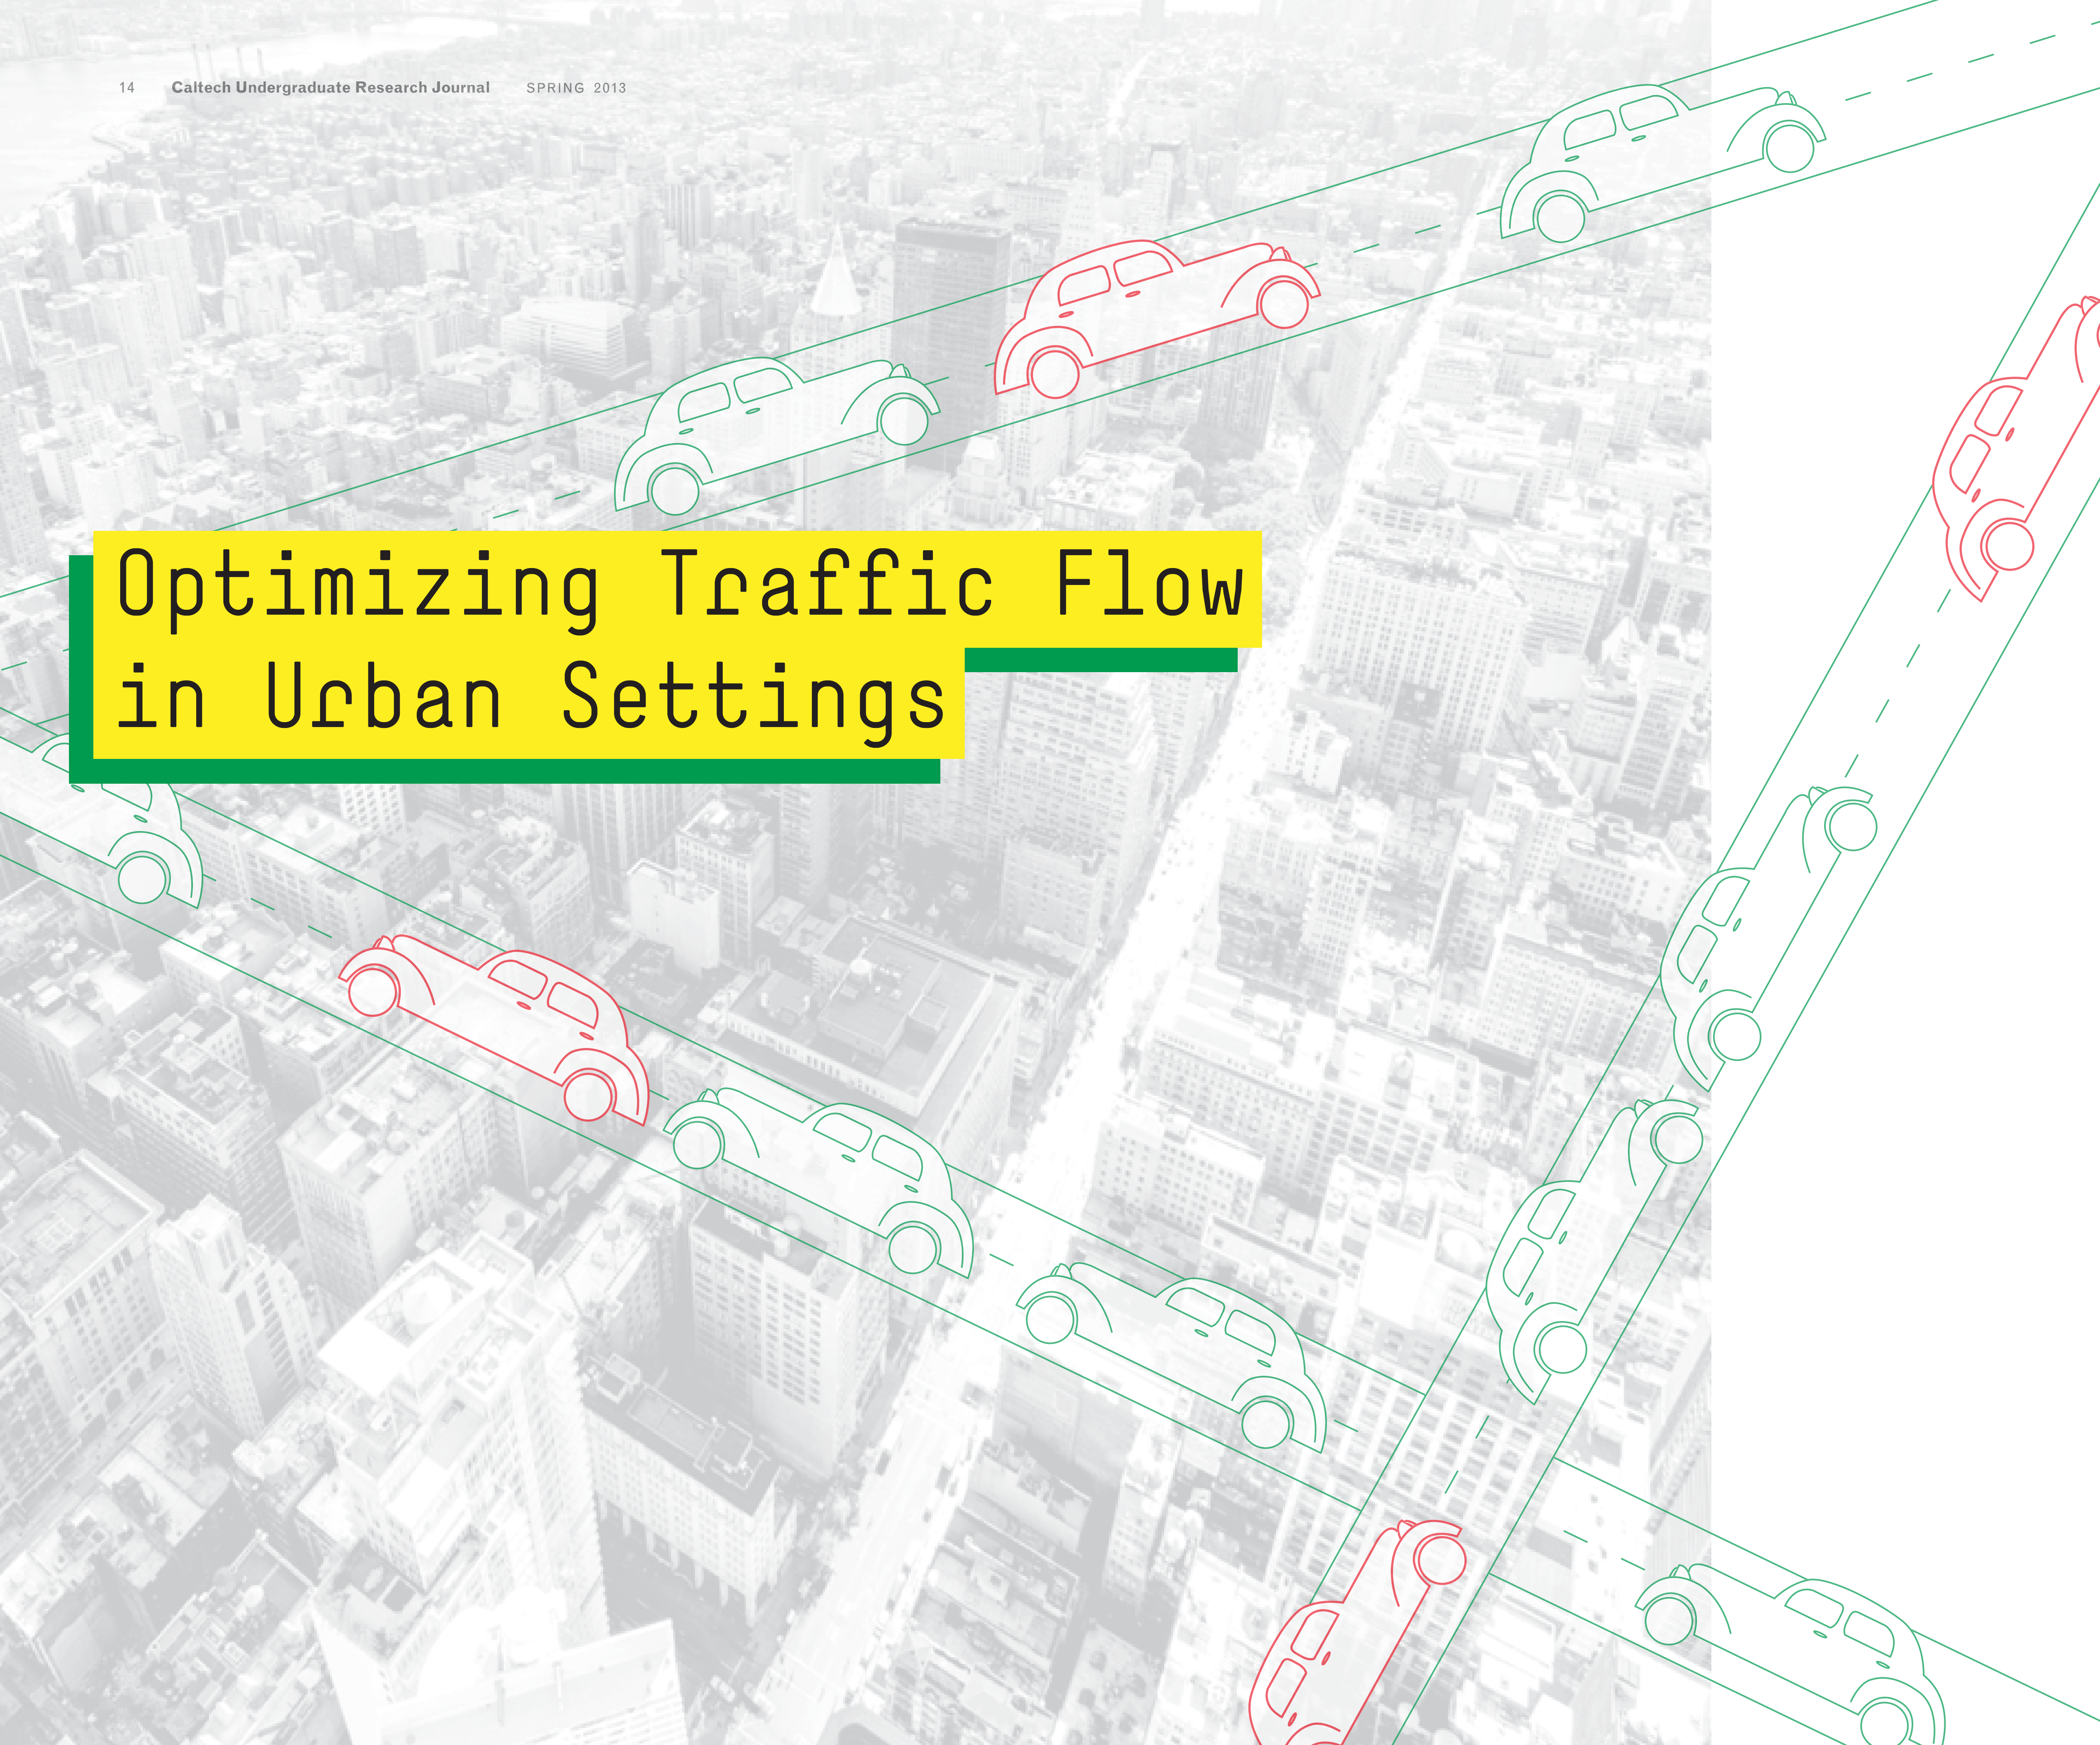

Author: Bertrand Ottino-Löffler
Hometown: Winnetka, Illinois
College: Caltech
Major: Mathematics
(but with strong interest in physics)
Year: Class of 2014
Hobbies: Bertrand's interests include yoga, reading comics, and playing video games, because he doesn't believe in being serious all the time. His plans for the future most likely involve teaching, mathematical physics, and plenty of research.
Mentors: Daniel Abrams, Mark Panaggio

Introduction

Nearly everyone is unhappy with the rage-inducing congestion that plagues their city roads. Improving the traffic situation on a city-wide scale is a difficult task, due to public policy factors and many conflicting interests. Here, we address only one aspect of the problem: optimizing the timing of traffic lights. The challenge appears simple, but it turns out to be a rich problem that presents an opportunity for substantial theoretical analysis.

We start by modeling a two-way street. Assume that all cars exist on a ring with evenly-spaced traffic lights, and an equal number of cars are present on both sides of the street. The cars are granted instantaneous acceleration from zero to velocity V, which is constant across all cars. Next, imagine each traffic light as an arrow rotating continuously around a circle (Fig. 1). The light is red when the arrow is in the red half of the circle, and green when it is in the green half. The angle of the arrow is its phase, and we let 0 to π

Figure 1.

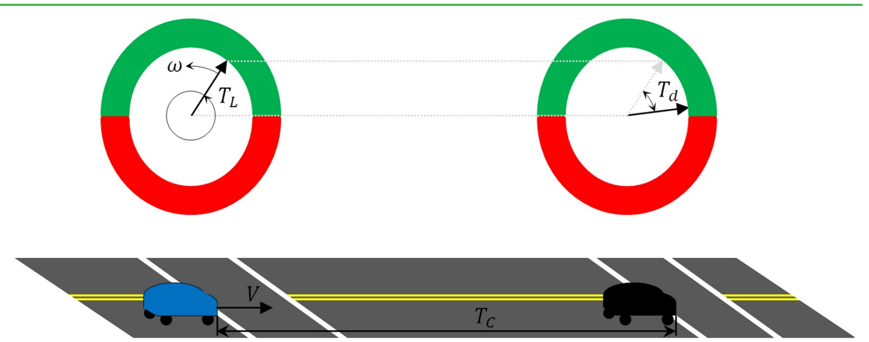

Visual representation of a two-way street. A car moves at velocity V between two traffic lights in time T_C. Traffic lights are depicted as rotating arrows with angular velocity ω, and complete one full cycle in time T_L. The phase of each traffic light is delayed from the previous one by a time T_d.

Table 1.

Key variables used in the theoretical model and simulation.

Variable	Definition
V	Speed of a car in motion
V_{avg}	Average car speed, or $V \times$ (Time Moving)/(Total Time)
N_{LT}	The number of lights a car has passed
ω	Angular frequency of each traffic light
T_L	Time for a traffic light to complete one full cycle
T_C	Time for a car to travel between two lights
T_d	Delay time between lights
r	T_L/T_C
r_d	T_d/T_c
M	$r/(1-r_d)$

> At first, it appears that the overall system is best served by giving one direction a perfect green wave and forgetting about the other side. But […] we actually do better by sacrificing some of the efficiency in one direction for improved conditions in the reverse direction.

radians denote green and π to 2π radians denote red. The time it takes for each arrow to complete full cycle is the period, T_L, which implies an angular frequency ω (Table 1).

The interesting part of the problem comes from modifying the initial phases of the lights. Let T_C be the time it takes for a car to travel between two lights, and T_d be the time delay between the phases of two adjacent lights. If the T_d for every pair of lights is equal to T_C, then we achieve a "green wave": a driver never sees more than one red light, and his takeout food is still warm when he gets home. This is optimal. However, suppose T_d were larger than T_C by the constant value T_L / 2. Now, our driver hits every red light on the road and has to wait the entire half-cycle before moving again. This suboptimal "red wave" leaves him nothing but frustration and cold, mushy takeout.

Therefore, holding T_d constant between all the lights allows us to span the best and worst case scenarios. If we take $r = T_L / T_C$ and rd = T_d / T_C, then the entire system is described with two parameters. Since $T_d < T_L$ by definition, it follows that $0 < r_d < r$. Also, we tend to assume $r > 1$. Note that for cars going the reverse direction, the phase difference between consecutive lights is the opposite of what the forward direction experiences. Thus, their effective rd becomes $r - r_d$. This means that it is generally impossible to set up a green wave in both directions, so maximizing efficiency for the whole street is a messy affair.

Producing a chimera state in a ring of traffic lights is tricky, but could potentially lead to significant improvements in efficiency for cars travelling in both directions.

Figure 2.

Travel efficiency versus r_d at two fixed values of r. 1001 values of r_d were sampled on each plot. The red and green curves show the efficiency of cars moving to the right and left, respectively. The black curve is the average of the green and red curves.

Figure 2. a

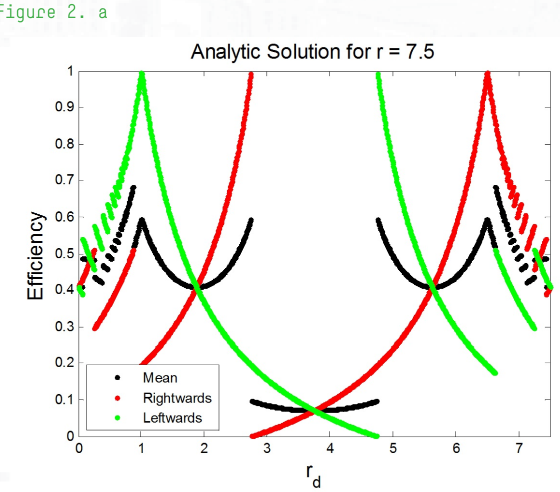

Figure 2. b

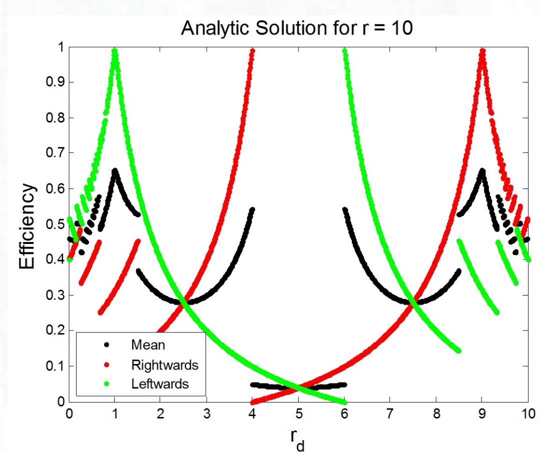

Mathematical Modeling

We start by trying to predict how the average speed of a car varies with r_d, for a fixed r. Constraining the value of r lets us limit the range of rd, as noted before. Let us begin by considering just one car. The average speed of the car is defined as the distance traveled over time, or in other words,

$$V_{avg} = \frac{V \times (\text{Time Moving})}{(\text{Total Time})} \qquad \text{(Eq. 1)}$$

We want to rewrite Eq. 1 in terms of the parameters of our model. It is clear that the Time Moving is simply equal to the travel time between lights (T_C) times the number of lights passed, henceforth known as N_{LT}. Now, let us find the Total Time – the sum of the time moving and the time spent at a red light. When the car reaches the N_{LT} th light, it's journey will end if the light is red; that is, if the phase of this light mod 2π is between π and 2π radians. If we take Z to be some integer, and ΔΦ to be the phase difference between consecutive lights, then the inequality

$$(2Z-1)\pi \;<\; \Delta\phi \times N_{LT} + \omega \times (\text{Time Moving}) \;<\; 2\pi Z \qquad \text{(Eq. 2)}$$

must be satisfied for the car to stop. Because we are looking for the furthest that the car can travel under our constraints, we must find the most restrictive (i.e. smallest) integer value for N_{LT}. By letting $M = r / (1 - r_d)$, we have that

$$Z = \text{ceiling}\left(\frac{N_{LT}}{M}\right) = \text{floor}\left(\frac{N_{LT}}{M} + \frac{1}{2}\right) \qquad \text{(Eq. 3)}$$

must hold as well.

It is possible to solve Eq. 3 for N_{LT} in terms of M. After doing so, we find that the Total Time is equal to $Z \times T_L + N_{LT} \times T_d$. This allows us to rewrite the average velocity function in terms of r and r_d, giving us

$$V_{avg} = \frac{V \times N_{LT}}{r \times \text{ceiling}\left(\frac{N_{LT}}{M}\right) + r_d \times N_{LT}} \qquad \text{(Eq. 4)}$$

We can normalize this expression into an efficiency value, because we know the maximum and minimum possible speeds. V_{avg} cannot be greater than V, the average speed during a green wave. Likewise, it cannot be less than the red-wave speed, which is given by $(V \times T_C) / (T_C + T_{Lt} / 2)$.

Figs. 2a and 2b plot efficiency versus r_d for forward and reverse traffic, as predicted by this model, for two different values of r. The curve for cars going in the forward direction is the horizontal mirror of the curve for the reverse direction. Let us recall the discussion of red and green waves. We would predict that the velocity of the leftwards car is maximized at $r_d = 1$ and minimized at $r_d = r/2 + 1$, and that the velocity of the rightwards car is maximized at $r_d = r - 1$ and minimized at $r_d = r/2 - 1$. Both plots reflect this prediction. An unexpected feature is the second maximum, which occurs just beyond the red wave minimum. Here, each light turns red just after the car travels through it – a careful distinction that is, as one might imagine, highly sensitive to perturbations. Examining the average curve, we realize something curious. At first, it appears that the overall system is best served by giving one direction a perfect green wave and forgetting about the other side. But in 2a, we actually do better by sacrificing some of the efficiency in one direction for improved conditions in the reverse direction. This compromising behavior is worth noting, and should be investigated upon future work.

Creating a Simulation

To confirm these theoretical results with simulations, adaptive time steps are needed to accurately record the time spent waiting at a red light or travelling between lights. We begin by simulating the "bosonic" case, where the cars do not interact with one another. With some slight modifications, we can extend this simulation to the more realistic "fermionic" situation, in which the cars cannot pass through one another.

Bosonic

Since the cars in the bosonic model have no interaction with one another, it suffices to simulate just one car on a ring of 500 lights for 5000 arbitrary time units. A plot of the simulation efficiency versus the model-predicted efficiency is shown in Fig. 3. The bosonic simulation aligns almost perfectly with the model.

Fermionic

As it turns out, two cars in the real world usually cannot occupy the same place at the same time. How can we tell if our predictions are still valid for cars with non-zero densities? Addressing this question requires a proper simulation of the situation, in which cars cannot pass one other. We expect that as more cars are added to the system, the average velocity will monotonically decrease. Moreover, we expect that cars will start piling up behind red lights, resulting in an additional time delay in reaching the next light. If the first car in the line barely makes it through a green light, all of the following cars will not. This should result in a smoothing of discontinuities in the efficiency graph, especially in the peak right before the red wave.

The plots of fermionic simulation efficiency versus model-predicted efficiency in Figs. 4a and 4b confirm these predictions. 1250 and 6500 cars were simulated on a ring of 500 lights that could support a maximum of 12500 cars. The large number of cars means that multiple parts of the system are being sampled at once; thus, the simulation only needed to be run for 450 arbitrary time units (as opposed to 5000 previously) for convergent behavior to be observed.

Because making a full efficiency versus r_d graph can be taxing at high densities, it is more efficient to hold r_d constant and vary the number of cars in the system. Fig. 5 shows the effects of increasing the number of cars for a fixed r and r_d on the average car speed. Even for a frustratingly high traffic density, Fig. 5a shows that there is little lost in the way of speed until a critical transition point. The critical point arises when so many cars have backed up behind a red light that it takes multiple cycles of the light before a car can pass. This is called a "jamming transition". No complete theory currently exists for the locations of the critical point given a value of r_d. Part of the problem is that, as mentioned before, certain r_d values produce ideal results that are highly sensitive to perturbations. So even though velocity graphs made with these values do contain a transition, the curve is rounded off as seen in Fig. 5b. This makes it difficult to determine the true critical point. Sometimes, as in Fig. 5c, this decay causes a secondary transition that makes critical point detection even harder. All the transitions that are clean enough to estimate a critical point are plotted in Fig. 6.

Figure 3.

Simulated efficiency data (red) superimposed on the analytic efficiency curve (blue) for Bosonic traffic in one direction. Cars were simulated on a loop of 500 lights for 5000 arbitrary time units. 1001 equally spaced values of r_d were taken at r = 7.5, T_C = 1. Frustration refers to a boundary effect which has little effect on the overall performance.

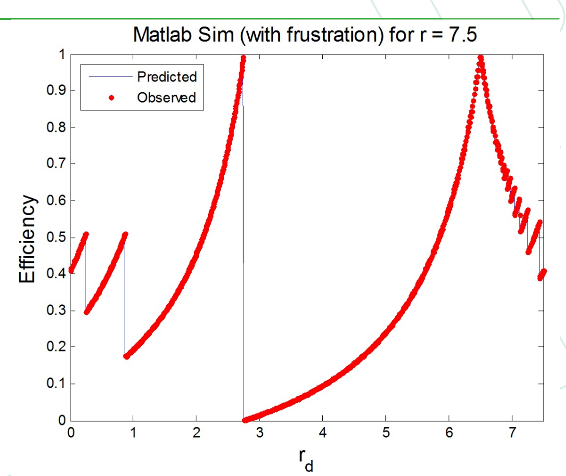

Figure 4. a

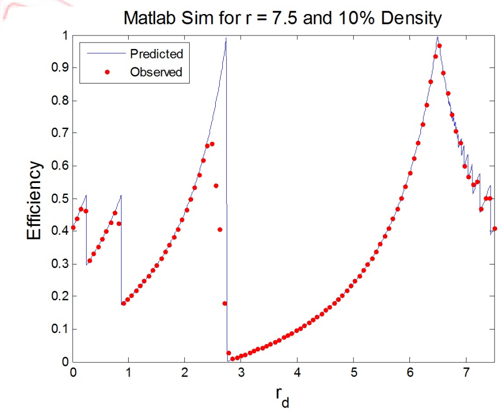

Figure 4. b

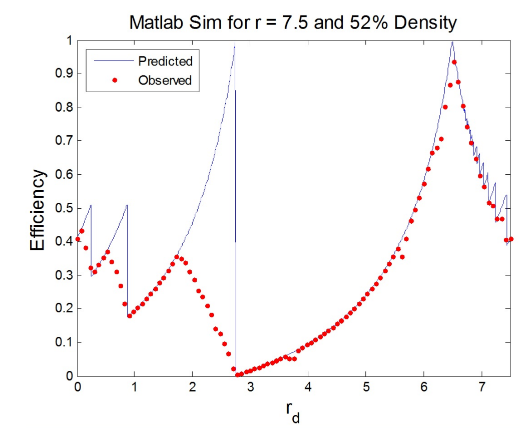

Figure 4.

Simulated efficiency data (red) superimposed on the analytic efficiency curve (blue) for fermionic traffic in one direction. The simulated road was a ring of 500 lights with room for 12500 cars. 1250 cars were simulated in 4a, and 6500 cars were simulated in 4b. In both simulations, 101 values of r_d are simulated at r = 7.5 and T_C = 2 for 450 time units.

Figure 5. a

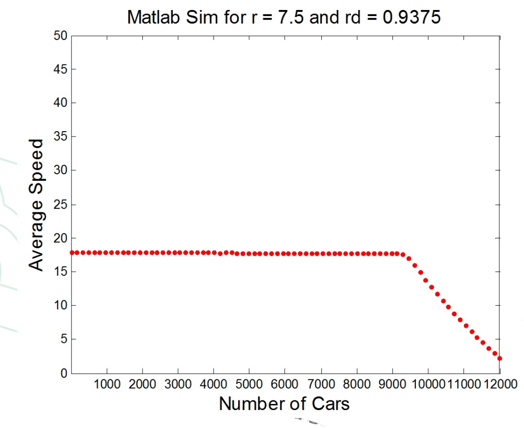

Figure 5. b

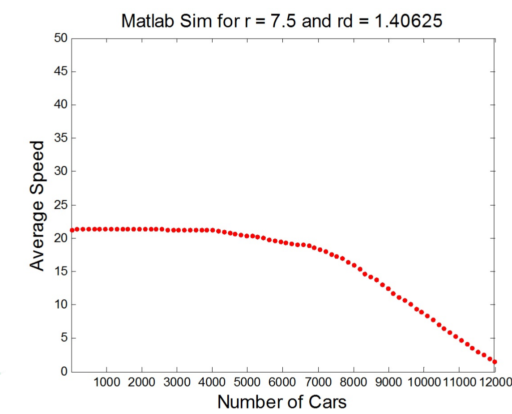

Figure 5. c

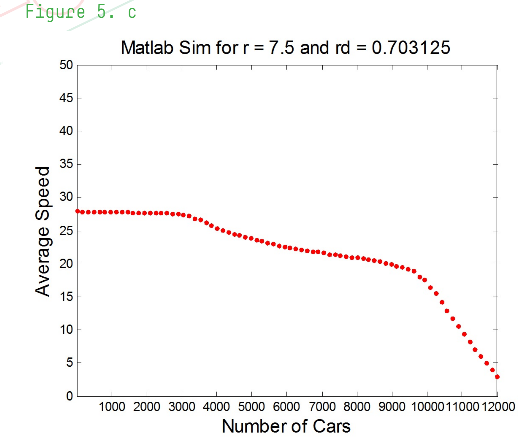

Figure 5.

Figure 5 a, b, c. The value of r_d is held constant at three different values. Plots show the average car velocity versus number of cars. Simulations were run in a system with a maximum capacity of 12500 cars, with r = 7.5 and T_C = 2. 76 simulations were each run for 500 arbitrary time units.

Future work

The results presented cover only a fraction of possible questions. But the model is fully set up to address many different optimization problems.

In the real world, cars travel at different speeds. How does randomness affect the efficiency of the cars? Since the simulated traffic is moving around a ring, all the cars would likely bunch up behind the slowest moving car. These results would be uninteresting and uninformative. More meaningful results can be obtained if the cars re-randomize their velocities periodically throughout the simulation. However, we predict that this randomization will not drastically alter the shape of the efficiency graph.

A more exotic possibility is to change the fundamental behavior of the lights, which is the long-term goal of this research. Under the new schema, each light would be coupled to every other light in the following manner. For a system with N lights, we define a coupling strength K and a phase delay α.

If light i has phase Φ i and natural frequency ω i, then its behavior is described by

$$\frac{d\phi_i}{dx} = \omega_i + \frac{K}{N}\sum_{j=1}^{N}\sin(\phi_j - \phi_i + \alpha). \qquad \text{(Eq. 5)}$$

If the lights all have the same ω and start at random phases, a Kuramoto model predicts that the lights will eventually attain identical phases for any nonzero K – a drab result.[1] If the coupling decays with distance, however, the results are intriguing. With the right initial conditions, then we may arrive at a “chimera state”, in which the lights are divided into two categories: coherent oscillators (which move together) and incoherent oscillators (which move in a randomized fashion).[2] Producing a chimera state in a ring of traffic lights is tricky, but could potentially lead to significant improvements in efficiency for cars travelling in both directions.

Figure 6.

Figure 6. Estimated critical densities at various values of r_d. These values were inferred from plots similar to the one in Fig. 5a.

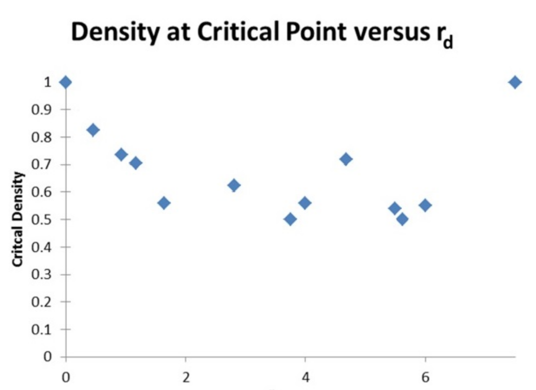

Acknowledgements

I would like to thank Daniel Abrams, Mark Panaggio, the Student Faculty Programs office of Caltech, and Northwestern University. I would also like to thank the Mellon Mays Fellowship for their continued support.

Further Reading

1. Strogatz, Steven. From Kuramoto to Crawfrord: exploring the onset of synchronization in populations of coupled oscillators, Physica D, 143:1-20, 2000.

2. Abrams, D.M., Rennie Mirollo, Steven H. Strogatz, and Daniel A. Wiley. Solvable model for chimera states of coupled oscillators. Physical Review Letters, 101(8):084103, 2008.