Driver Strategy and Traffic System Performance

Terence Kelly*

Department of Computer Science
Princeton University
35 Olden Street
Princeton, NJ 08544 USA

Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501 USA

September 5, 1996

Abstract

As increasingly sophisticated routing, navigation and trip-planning devices are installed in automobiles, it becomes necessary to consider the likely effects of such devices on overall traffic system performance. A simple simulation model of "rush-hour" commuting is presented and the system-level consequences of a variety of agent-level behavior patterns are explored. In the context of this model, increasingly sophisticated agent-level commuting strategies result in decreased system-level performance as measured by several criteria.

1 Introduction

The day is fast approaching when a global positioning system and a navigational computer will together cost less than a spare tire and will be standard equipment on new automobiles. Motorists will soon have the means to compute with ease detailed and accurate statistics on their own commuting times. It is therefore natural to ask what impact these technological developments will have on the overall performance of traffic systems.

This paper presents a very simple agent-based model of "rush-hour" commuting intended to explore this question. The strategies or algorithms used by

*Correspondence should be addressed to Artificial Intelligence Lab, Department of Electrical Engineering and Computer Science, University of Michigan, 1101 Beal Avenue, Ann Arbor, MI 48109–2110 USA, electronic mail: `tpkelly@eecs.umich.edu`.

individual driver-agents to plan future actions based on past experience are varied from simple to sophisticated. The model thus provides a context in which to test empirically the notion that increased cleverness on the part of individuals inevitably results in better outcomes for populations.

The surprising result is that in the context of this model more sophisticated agent strategies result in worse global system performance by a number of very different measures. More precisely, when agents retain more of their past history and perform more elaborate computations on this data in an attempt to attain their goals, the outcome is far *worse* for nearly every agent and markedly better for *no* agent. The straightforward policy implication is that schemes to improve traffic system performance by increasing the information and computational power available to individual motorists must be evaluated with caution.

2 Methods

2.1 The Model

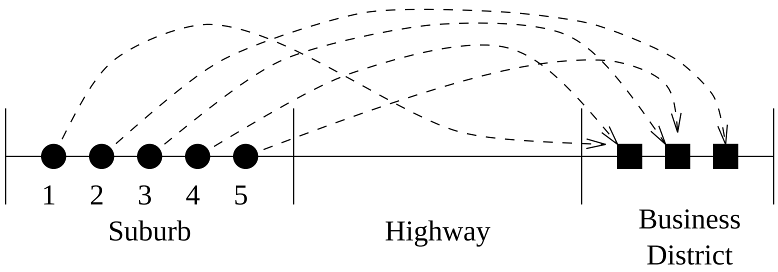

Figure 1: The Rush-Hour World

In Figure 1 we see a depiction of my rush-hour world, a model inspired by but *not* similar to that described in [8]. We have a two-lane road divided into three sections: a "suburb" containing evenly-spaced sources of cars or "houses," a "highway" which contains neither sources nor sinks of cars, and a "business district" containing "offices" or destinations. At the beginning of a simulation run, each car is randomly assigned an office (shown as dashed arrows in the figure). This is the *only* use of the random number generator. For discussion purposes, cars are numbered $1 \ldots N$, with car 1 being the furthest from the business district. On each "day" of the simulation each car drives to its office. Each day every commuter decides upon a time to leave home. A driver's *strategy* is the mapping from its past experience to this decision. The driver population is homogeneous with respect to strategy, i.e. during any given run of the simulation every car uses the same strategy every day. See [8] and the references therein for more general discussions of similar models.

Cars enter and leave the highway at zero speed. At each tick of the clock Δt, each car calculates a new speed and advances at the new speed along the road. The new speed is subject to the following constraints: it must differ from the current speed by no more than $a\Delta t$ where a is a fixed maximum rate of acceleration; it must be between zero and a speed limit u_f; it must be such that a minimum time headway h_{min} is maintained with respect to the car ahead; it must be such that the car can decelerate to a stop at its destination at a rate no greater than a; and it must be as fast as possible so long as the foregoing constraints are satisfied. All cars' positions and velocities are represented as real numbers, *not* discrete values as in cellular automata or "particle hopping" models (e.g., [7]). See [3] for more information on the driving logic and [1] for simulation source code.

2.2 Notation

Subscript $i \in 1 \ldots N$ is used to index over cars and $j \in 1 \ldots M$ indexes over simulated days, where N and M respectively denote the total number of simulated cars and days in a run of the simulation. Car i begins each day's commute at a "home" location, x_{0i}, and drives to "work" at a destination x_{fi}. Let x_S mark the end of the "suburb" of our roadway, x_H mark the end of the "highway," and x_B denote the end of the "business" district (and of the roadway). Then $0 < x_S < x_H < x_B$, $0 < x_{0i} < x_S \;\; \forall \;\; i$, and $x_H < x_{fi} < x_B \;\; \forall \;\; i$. Note that $x_{0i} = x_B - ix_S/(N+1)$ and thus i is a convenient scale-independent proxy for a car i's starting position.

We define t_{Ii}, the *ideal commuting time* of car i, as the time car i would require to drive to work on an otherwise empty road; t_{Ii} is determined entirely by x_{0i} and x_{fi}. Every motorist is due at work at time the same time t_D (we might say $t_D \equiv$ 9:00 a.m.). The *actual commuting time* of car i on day j is written t_{Aij}. This quantity will vary from day to day depending on traffic congestion. Each day each motorist predicts how long its commute will take on that day. This *expected commuting time*, $E(t_{Aij})$, determines when the motorist leaves home: each motorist departs at time $t_{departure} = t_D - E(t_{Aij})$ (each driver's intent is to arrive at work exactly on time). A driver's *strategy* is simply the procedure used to compute $E(t_{Aij})$ based on $t_{Ai1} \ldots t_{Ai(j-1)}$.

2.3 Performance Metrics

We can now define the performance metrics we will examine. One obvious measure of traffic system performance is total system throughput or flow. I have chosen not to consider this performance measure and instead to focus on other measures associated with the socially important idea of fairness, the economically important notion of lost productivity, and the consistency of traffic system performance as experienced by individual motorists.

2.3.1 Fairness

Most people would say that it is unfair for some motorists to experience more congestion-induced delay than others. We can formalize this notion of fairness as follows: Let $d_{Aij} \equiv t_{Aij} - t_{Ii}$ denote the *absolute delay* of car i on day j, and let $d_{Rij} \equiv d_{Aij}/t_{Ii}$ denote car i's *relative delay* on day j. Define fairness F on day j as

$$F(j) = \sqrt{\frac{\sum_i(\overline{d_{Rj}} - d_{Rij})^2}{N-1}} \tag{1}$$

where $\overline{d_{Rj}}$ denotes the mean of all commuters' relative delays on day j. $F(j)$ is simply the standard deviation of all motorists' relative delays on day j. This measure will be zero if congestion lengthens every driver's commute by the same factor.

A second approach to the fairness issue is to examine each commuter's mean relative delay as a function of i:

$$\overline{D(i)} \equiv \frac{\sum_j(t_{Aij} - t_{Ii})}{Mt_{Ii}} \tag{2}$$

Intuitively, we expect that commuters who live closer to the business district will experience lower relative delays on average, i.e. $a > b$ should usually coincide with $D(a) < D(b)$.

2.3.2 Aggregate Lateness

We define the *absolute lateness* of car i on day j as $l_{Aij} \equiv t_{arrival} - t_D$. Because every car leaves home at time $t_{departure} = t_D - E(t_{Aij})$ we have $l_{Aij} = t_{Aij} - E(t_{Aij})$. Aggregate lateness on day j is defined as

$$L(j) \equiv \sum_i \max(l_{Aij}, 0) \tag{3}$$

We might think of $L(j)$ as a measure of productivity lost due to traffic congestion. This definition of L implies a cost of lateness that is zero for $t_{arrival} \leq t_D$ and increases linearly for $t_{arrival} > t_D$. Compare, e.g., [10].

Note the distinction between lateness and delay as defined here: a car's lateness depends on its arrival time, whereas its delay depends on its transit time. It is possible to have lateness without delay and vice versa. Note that it is trivial to reduce $L(j)$ to zero: simply arrange for every car to depart for work *very* early. To reduce the absolute delays of all motorists to zero, however, would require co-ordinating their departure times.

2.3.3 Consistency

One property of a traffic system that real commuters care about is *consistency.* A commute that takes 30 ± 20 minutes isn't necessarily preferable to one that takes 40 ± 5 minutes. We define the consistency of commuting experienced by driver i as

$$C(i) \equiv \frac{1}{t_{Ii}}\sqrt{\frac{\sum_j(\overline{t_{Ai}} - t_{Aij})^2}{M-1}} \tag{4}$$

where $\overline{t_{Ai}} \equiv \sum_j t_{Aij}/M$ is the mean actual commuting time of car i. This is simply the standard deviation of car i's absolute delays normalized to the car's ideal commuting time. If motorist i experiences the same relative delay every day, $C(i)$ will be zero.

2.4 Experiments

The independent variable under consideration is adaptive strategy. The only simulation parameter varied in the experiments described here is the rule that drivers use to decide when to leave home for work each day. For each driver strategy the simulation was run three times with different random number seeds. Parameters are as follows: $N = 100$ cars, 10 offices, 2 lanes, $M = 100$ days, $\Delta t = 0.1$ sec, speed limit $u_f = 60$ MPH, following time $h_{min} = 2$ sec, maximum vehicle acceleration $a = 4$ MPH/sec.

The road is five miles long, with a one-mile suburb, a three-mile highway, and a one-mile business district, i.e. $x_S = 1$ mi., $x_H = 4$ mi., and $x_B = 5$ mi. On day 1 of *every* run, regardless of the strategy used, each driver leaves home under the assumption that the road is empty, and therefore $E(t_{Ai(j=0)}) = t_{Ii}$. Note that the driver population is homogeneous in every simulation run: all drivers always use the same strategy. The three strategies tested are the following:

2.4.1 The "Yesterday Rule"

This strategy means that on every day other than the first, the motorist leaves home under the assumption that today's commute will take as long as yesterday's did. So if it took ten minutes to drive to work yesterday, leave home at 8:50 AM today. More formally, $E(t_{Aij}) = t_{Ai(j-1)}$. We call this strategy "the yesterday rule." This strategy is arguably the simplest reasonable way for a driver-agent to adapt its behavior based on past experience: it requires that each agent store only a single number (yesterday's commuting time) and perform a single simple arithmetic operation on the saved state.

2.4.2 Mean of Previous Commuting Times

For this experiment each driver leaves work under the assumption that the commute will take as long as the arithmetic mean of all the driver's previous commuting times. Formally, $E(t_{Aij}) = (\sum_{k=1}^{j-1} t_{Aik})/(j-1)$. We refer to this strategy as "the mean rule."

2.4.3 Least-Squares Regression Prediction

For this experiment, each driver performs an ordinary least squares (OLS) regression on all past commuting times in order to predict how long today's commute will take. The driver linearly extrapolates using the regression coefficients thus obtained in order to predict today's commuting time $E(t_{Aij})$. We call this "the OLS rule."

2.4.4 Windowing

Note that when drivers follow the mean rule and the OLS rule they use *all* of their commuting history rather than, say, only the previous few days' experience. It could be argued that a more rational approach would be to use a small window on the past when predicting $E(t_{Aij})$. I decided against that approach for two reasons: I wanted to avoid arbitrarily selecting a window size or exploring the full range of window sizes; and I wanted the mean and OLS strategies to take advantage of *all* information available to driver-agents.

3 Results

3.1 System-Level Performance Measures

The daily aggregate lateness data $L(j)$ for all nine simulation runs are displayed as time series in Figure 2.

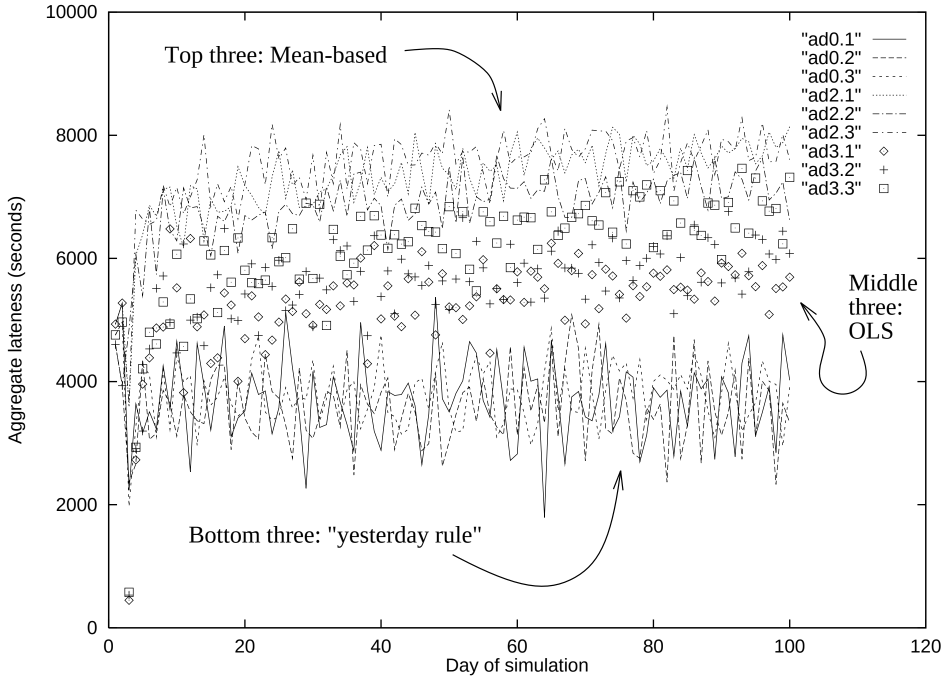

Figure 2: Aggregate Lateness Time Series $L(j)$

Taken at face value, these results point to a surprising conclusion: from the standpoint of total system performance, the simple "yesterday rule" described in section 2.4.1 by far outperforms the more sophisticated agent strategies.

Recall from section 2.3.1 that the fairness $F(j)$ of the traffic system on a given day j is defined as the standard deviation of all drivers' relative delays

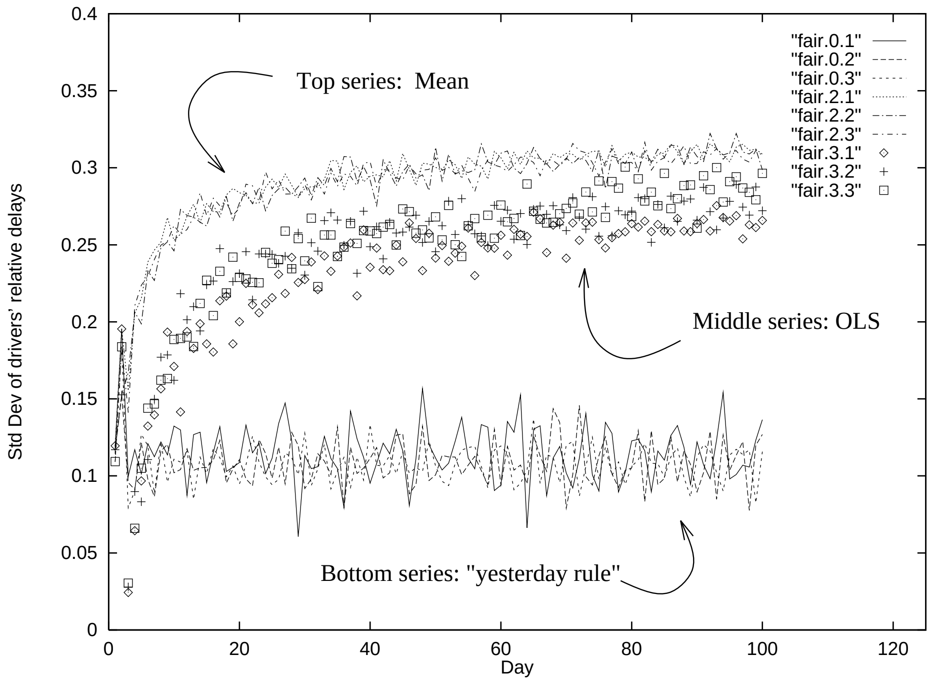

Figure 3: Fairness Time Series $F(j)$

on that day. If all drivers experience roughly the same relative delay, then this measure will be low. A high fairness figure implies that the system is *unfair* on a particular day — some drivers experience very high relative delays, while others experience much less. A time series of system fairness is shown in Figure 3. We see that by this measure the "yesterday rule" again leads to happier overall system behavior than the mean-based and OLS-based agent-level strategies.

3.2 Individual-Level Measures

When we examine the consistency of drivers' commuting times as a function of their home location, $C(i)$, we see a similar picture. In Figure 4 consistency as defined in section 2.3.3 is plotted against i ("car number"). Recall from section 2.2 that i is a scale-independent measure of a commuter's home location: car number 1 is located farthest from the business district, and car 100 lives nearest the business district. As in the case of mean relative delay, we see that when everyone follows the "yesterday rule" the result is global happiness and equitability, whereas the other two rules lead to a situation where long-distance commuters suffer disadvantages, but short-distance commuters do no better than when the yesterday rule is applied.

Another view of the fairness issue involves the advantages associated with living close to one's workplace. We intuitively expect that one is better off

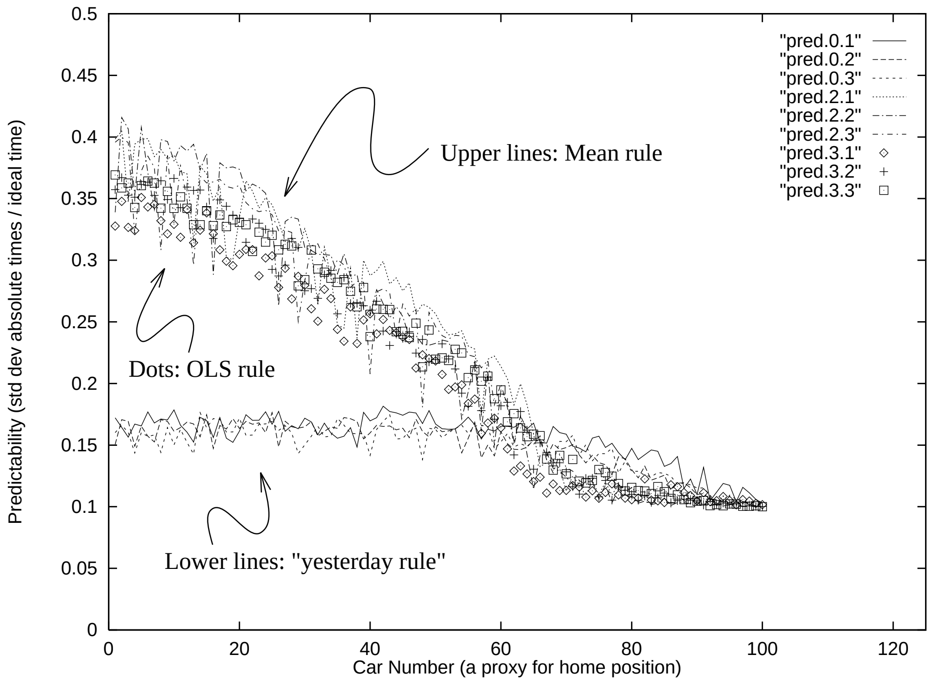

Figure 4: Consistency as function of home location $C(i)$

living near work, but this advantage may depend in part on the driving habits of others. In Figure 5 we plot drivers' mean relative delay $\overline{D(i)}$ over 100 days as a function of i. We see from Figure 5 that when drivers apply the "yesterday rule" the result is that nearly everyone experiences roughly the same low level of mean relative delay. By contrast, when drivers choose departure times based on the mean of their past commuting times, or based on an OLS extrapolation of past experience, the result is a high degree of inequality. The mean rule and OLS rule lead to unhappy outcomes for folks who live far from work, but do not improve the situation for commuters who live near the business district.

4 Discussion

The results presented above are, to the author and a number of others, highly counter-intuitive. Taken at face value, they seem to imply that traffic system performance degrades as driver-agents adopt strategies that take into account more information and perform more elaborate computations. This surprising result demands explanation.

The poor performance of the "mean rule" can perhaps be explained as a system convergence toward a worst-case scenario [6]. For simplicity, assume that all x_{0i} are equal and all x_{fi} are equal (i.e. every driver commutes from the

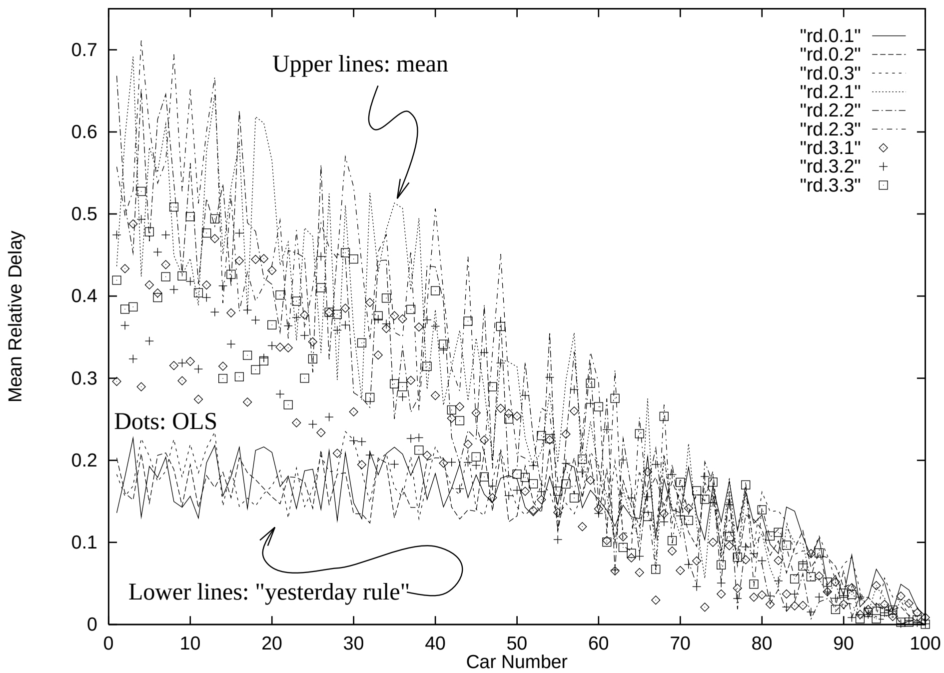

Figure 5: Mean Relative Delay as function of home location $\overline{D(i)}$

same "home" to the same "office") and that we have a one-lane road. On the first simulated day, all drivers plan to depart home and arrive at the office at exactly the same time. But this is impossible; at most one driver can arrive at work on time and all others must be either early or late. The *worst* solution is for everyone to leave at the same time. But this is precisely the solution toward which the system converges when drivers apply the mean rule day after day. It seems plausible that the insight obtained from this simplified scenario generalizes to the more complex experiment described above.

It is obvious that in terms of both aggregate lateness and individual relative delay, nearly every day's commuting experience is sub-optimal. It is *not* true that every day's experience is "Pareto inefficient," i.e. some drivers' experience is optimal, because on quite a few days a small number of lucky drivers arrive at work with no delay whatsoever. Nonetheless, it is clear that nearly everyone would be better off if departure times were globally assigned by a "central planner." An interesting question is how a central planner might go about assigning departure times so as to optimize aggregate lateness, fairness, or whatever. Another intriguing issue is how robust a centrally-planned schedule would be in the face of stochastic fluctuations in compliance with the plan. (Of course, an essential feature of automotive traffic is precisely the absence of a central planner, and we want to understand how agent-level behavior gives rise to macro-scale dynamics.)

A major goal of traffic simulation research is to provide models with predictive power to policymakers and public officials. In some cases, a model may instead provide a qualitative evaluation of the wisdom of a traffic management strategy. For instance, in [8] the authors present simulation results suggesting that advanced traffic management systems that push traffic systems toward maximal flow may actually *degrade* traffic system performance. The results presented here imply that systems which enable commuters to formulate plans based on a detailed analysis of their past experience may have negative consequences.

If the Rush Hour model presented here is an appropriate abstraction of real commuting from which robust, general results can be obtained, then a wider range of agent adaptive strategies might be explored. Windowed mean and OLS strategies are obvious first candidates for testing.

Furthermore, it is natural to ask what would happen in a mixed simulation, where a heterogeneous population of drivers apply different strategies. Indeed, this is perhaps the most interesting kind of question to ask. The results reported above suggest that the "mean rule" and "OLS rule" are outright losers, conferring no advantage on individual drivers and leading to dismal global system performance. But if *nearly* all drivers applied the "yesterday rule," would a small number of commuters using the other rules obtain any advantage? This is perhaps the most interesting open question surrounding this simulation.

A Acknowledgements

I am indebted to Ken Steiglitz, Perry Cook, Mariusz Jakubowski, Jim Roberts, and Kai Nagel for advice and support. I was supported by NSF grant PHY99300206 during the last stages of writing.

References

[1] `C` source code available at `http://www.santafe.edu/~tpkelly/rh/`.

[2] Daniel L. Gerlough and Matthew J. Huber. *Traffic Flow Theory.* Transportation Research Board, National Research Council, Washington, D.C., 1975.

[3] Terence Kelly. COS 496 midterm report: Rush hour (unpublished). 1996. Available at `http://www.santafe.edu/~tpkelly/rh/`.

[4] Wilhelm Leutzbach. *Introduction to the Theory of Traffic Flow.* Springer-Verlag, second edition, 1988.

[5] Adolph D. May. *Traffic Flow Fundamentals.* Prentice Hall, 1990.

[6] Kai Nagel. Personal communication.

[7] Kai Nagel. Particle hopping models and traffic flow theory. *Physical Review E*, 53:4655, 1996.

[8] Kai Nagel and Steen Rasmussen. Traffic at the edge of chaos. In Rodney A. Brooks and Pattie Maes, editors, *Artificial Life IV*, pages 222–235. MIT Press, 1994.

[9] Gordon Frank Newell. *Traffic Flow on Transportation Networks.* MIT Press, 1980.

[10] Arnott R., De Palma A., and Lindsay R. A structural model of peak-period congestion: A traffic bottleneck with elastic demand. *American Economic Review*, 83(1):161–179, 1993.

[11] Philip D. Straffin. *Game Theory and Strategy.* The Mathematical Association of America, 1993.