Traffic Flow Models and the Evacuation Problem

Samuel W. Malone, Carl A. Miller, Daniel B. Neill

February 12, 2001

1 Introduction

In our paper, we consider several models for traffic flow. The first, a steady state model, employs a model for car following distance to derive the a traffic flow rate in terms of empirically estimated driving parameters. From this result, we go on to derive a formula for total evacuation time as a function of the number of cars to be evacuated. The steady state model is analytically convenient, but has the drawback that it does not take the variance in the travelling velocities of vehicles into account. To address this problem, we develop a cellular automata model for traffic flow in one and two lanes, and augment our results through simulation. After presenting the steady state model and the cellular automata models, we derive a space-velocity curve that synthesizes these results. The section following this development of the basic models addresses the issue of restricting vehicle types using several tools for analyzing vehicle velocity variance.

To assess the problem of two lanes converging into one and traffic merging, the next sections address optimal flow issues and explain how congestion occurs. Finally, we bring the collective theory of our assorted models to bear on the five evacuation strategies in question in the section titled "Parallel Paths, and Applications to Evacuation Strategy." Lastly, we present a newspaper article/conclusion summarizing our results clearly and without going into a high level of mathematical detail.

2 Assumptions and Hypotheses

- Driver reaction time is approximately 1 second.
- Drivers tend to maintain a safe following distance; tailgating is unusual.

- All cars are approximately 10 feet long and 5 feet wide.
- Almost all cars on the road are headed to the same destination.

3 Terms

- *Density* - the number of cars per unit distance.
- *Occupancy* - the proportion of the road that is covered by cars.
- *Flow* - the number of cars per time unit that pass a given point on the highway.
- *Travel Time* - the amount of time that a given car spends on the road during evacuation.
- *Total Travel Time* - the sum of the travel times for all evacuated cars.

4 The Steady State Model of Traffic Flow

4.1 Motivation

The subtask of car following has been described successfully by mathematical models, and as authors such as Rothery have noted, understanding this part of the driving process contributes significantly to an understanding of traffic flow. Here, we use a model for the average spacing between vehicles, s, as a function of common velocity, v, which was obtained from data compiled in a series of 23 observational studies of highway capacity (Highway Capacity Manual, p. 4-1 in Rothery, Car Following Models).

4.2 Development

The speed-spacing relations that were obtained from these studies can be represented by the equation

$$s = \alpha + \beta v + \gamma v^2, \tag{1}$$

where α, β, and γ can take on various values and have the following physical interpretations:

$$\begin{aligned}
\alpha &= \text{the effective vehicle length, } L \\
\beta &= \text{the reaction time} \\
\gamma &= \text{the reciprocal of twice the maximum average deceleration} \\
&\quad\ \text{of a following vehicle.}
\end{aligned}$$

The measure of seperation distance given by the above equation denotes the distance between the midpoints of successive cars. Using these relationships, we can obtain the optimal value of traffic density (and velocity) that maximizes the flow rate. We express this fact in the following theorem:

Theorem 4.1. *When coupled with the fundamental equation for traffic flow, $q = kV$, Equation 1 implies the following values for maximal traffic flow (q^*), optimal traffic density (k^*), and optimal velocity (v^*):*

$$q^* = (\beta + 2\gamma^{1/2}L^{1/2})^{-1} \tag{2}$$

$$k^* = \frac{\beta(\gamma/L)^{1/2} - 2\gamma}{\beta^2 - 4\gamma L} \tag{3}$$

$$v^* = (L/\gamma)^{1/2} \tag{4}$$

Proof. First, consider a group of N identical vehicles, each of length L, traveling at a steady state velocity with separation distance given by Equation 1. If we take a freeze frame picture of this group of vehicles spaced over a distance D, then the relation $D = NL + Ns'$, where s' is the bumper to bumper seperation distance, must hold. By definition, $s' = s - L$. Combining these facts, we obtain $k = N/D = N/(NL + Ns') = 1/(L + s') = 1/s$.

Now, with this relationship between density (k) and separation distance (s) in hand, we invoke Equation 1 to get an equation for density in terms of steady state velocity:

$$k = \frac{1}{\alpha + \beta v + \gamma v^2}.$$

This can be rearranged to obtain a quadratic equation in v that has two roots. Taking the positive root to ensure positive values for velocity yields the equation for velocity in terms of density,

$$v(k) = \frac{1}{2\gamma}\sqrt{4\gamma/k + (\beta^2 - 4\gamma L)} - \frac{\beta}{2\gamma}.$$

Using the fundamental equation for traffic flow, $q = kv$, we have the following expression for flow as a function of density:

$$q(k) = \frac{k}{2\gamma}\sqrt{4\gamma/k + (\beta^2 - 4\gamma L)} - \frac{k\beta}{2\gamma}.$$

Differentiating with respect to k, setting the result equal to zero, and wading through a lot of algebra yields the optimal value for density, k^*, which is quoted in the statement of the theorem. The expressions $q(k^*)$ and $v(k^*)$ give, respectively, values for the maximum steady state flow and the velocity associated with optimal steady state density. It is merely a matter of computation to check that $q(k^*)$ is a maximum by evaluating $q''(k)$ at $k = k^*$ for the values of k^* that result from the assignments of α, β, and γ in this paper. $\square$

4.3 Interpretation and Uses

These equations for the maximum flow, the optimal density, and the optimal velocity in the steady state allow us to estimate plausible values for q^*, k^*, and v^* given reasonable assumptions regarding car length (L), reaction time (β), and the deceleration parameter (γ). According to Rothery (Car Following Models), a typical value empirically derived for γ would be $\gamma \approx .023$ sec^2/ft. If we let reaction time $\beta = 1$ second, and assume that car length is approximately $L = 10$ ft., then we obtain the following optimal values for the state parameters of interest:

$$\begin{aligned} q^* &= .510 \text{ cars/sec} \\ k^* &= .024 \text{ cars/ft} \\ v^* &= 20.85 \text{ ft/sec} \end{aligned}$$

Rothery notes that a less conservative estimate may be obtained for γ via the equation $\gamma = \frac{1}{2}(a_f^{-1} - a_l^{-1})$, where a_f and a_l are the average maximum decelerations of the following and lead vehicles, respectively. In principle, one might estimate a less conservative value for γ than the one quoted above (.023 sec^2/ft) by computing the average value of $|a_i^{-1} - a_j^{-1}|$ for all vehicle pairs (i, j) in a representative sample of US vehicles whose deceleration capabilities are known. In practice, we estimate the value for γ by assuming that, instead of being able to stop instantaneously (having an infinite deceleration capacity), the leading car has an deceleration capacity that is twice that of the following car. Thus, instead of using $\gamma = 1/2a = .023$ sec^2/ft., we use this implied value for a to compute $\gamma' = \frac{1}{2}(a^{-1} - 2a^{-1}) = \frac{1}{2}\gamma = .0115$

$\sec^2/\text{ft}$. Unless otherwise stated, we will henceforth assume this value of γ in our computations of maximum flow based on the steady state model. In the case at hand, the less conservative value for γ yields

$$\begin{aligned} q^* &= .596 \text{ cars/sec} \\ k^* &= .020 \text{ cars/ft} \\ v^* &= 29.5 \text{ ft/sec} \end{aligned}$$

The value for v^* given above is about 20 m.p.h., which seems fairly reasonable given that our objective is flow maximization (not velocity maximization). From practical experience, going 20 m.p.h in a regime of high traffic density with an implied bumper to bumper distance of 40 feet is not bad. With regards to the evacuation scenario in question, we can apply the flow maximizing formula for velocity to show that the evacuation of the South Carolina coast was far from optimal. From the fact that the trip from Charleston to Columbia usually takes about 2 hours at 60 m.p.h., we get a distance of about 120 miles. Traversing this distance in 18 hours implies an average velocity of about 7 m.p.h. and a bumper to bumper seperation distance of about 7.3 feet. Clearly, these driving conditions are abyssmal, and the steady state model gives us some sense of how far from optimal they are from the perspective of the individual driver.

4.4 Limitations of the Steady State Model

Although the steady state model allows us to compute a flow maximizing density and the corresponding velocity of a vehicle that is part of a group of vehicles moving in a convoy, it does not take into account the variance of the automobiles' individual velocities. Since situations of high traffic density are especially susceptible to inefficiencies of movement caused by the tendency of individual drivers to overcompensate or undercompensate for the movements of other drivers, understanding the behavior of velocity variance is important. Several of the following sections address this issue.

A second weakness of the steady state model is that the value for maximum flow encoded in it can only give us a first order approximation for the minimum evacuation time. Put another way, the problem of determining maximum flow is distinct from the problem of determining minimum evacuation time, and it is not necessarily the case that we can simply compute

(Number of Cars to be Evacuated) / (Maximum Flow Rate)

to obtain an estimate for the latter quantity. In the next section, we will tackle the problem of minimizing total evacuation time directly using the steady state model, and discuss the relationship of this (central) problem to a class of reasonable performance measures for evaluating evacuation strategies.

5 Minimizing Evacuation Time with the Steady State Model

5.1 Initial Considerations

Before we apply the machinery of the steady state model to the problem of minimizing total evacuation time for a group of vehicles, let us take a step back and consider the general problem of defining what it means for an evacuation to "succeed." The effectiveness of an evacuation is based on several factors. The first goal, as indicated, is keeping evacuation time to a minimum.

Second, it is imperative that the evacuation route be as safe as possible to drive on under the circumstances, which in the case of hurricanes could include inclement weather. While the authorities cannot control the state of the weather, they should be concerned with the elements of the driving conditions that affect the number of accidents that are likely to occur en route. Since the salient factors that affect the occurence of accidents (traffic density, the variance of vehicle velocities, and traffic speed) are modeled more explicitly in later sections, we will delay a discussion of their role in evaluating evacuation success.

The third factor affecting the success of an evacuation is how long, on average, it takes the individual driver to get to a safe destination (Columbia). This issue is related to the problem of minimizing total evacuation time, but not equivalent to it. We now present a general performance measure that relates the problem of minimizing evacuation time to the problems of maximizing traffic flow and minimizing individual driver transit time.

5.2 A General Performance Measure

In defining a general performance measure, we take into account both the goal of maximizing traffic flow and the desire of individual drivers to minimize their transit time, T, subject to the constraint that their velocity, v, is bounded above by a preferred cruising velocity, v_{cruise}. A natural metric M which captures both of these factors is

$$M = W\frac{N}{lq} + (1 - W)\frac{D}{v}, \tag{5}$$

where $0 \leq W \leq 1$ is a weight factor, D is the distance that must be traversed, l is the number of lanes, and N is the number of cars to be evacuated. This metric makes the assumption that the interaction between lanes of traffic (passing) is negligible, so that we may simply consider total flow to be the flow of an individual lane times the number of lanes. Given a value for W, the problem of minimizing M amounts to solving a one variable optimization problem in either v or k. Setting $W = 1$ corresponds to the problem of maximizing flow, which was solved analytically in the preceeding section. Setting $W = 0$ corresponds to the problem of maximizing velocity subject to the constraint that $v \leq v_{cruise}$. Trivially, this problem has the solution $M = D/v_{cruise}$. Recall that the steady state model is based on a model for car following, and does not apply to situations where cars can travel at the velocity v_{cruise}. In fact, the metric M balances the desire of drivers to go faster than is allowed in heavy traffic conditions with the natural goal of maximizing the flow rate of the traffic stream.

For clarification, observe that setting $K = 1/2$ corresponds to minimizing the total evacuation time (divided by two), since total evacuation time equals $\frac{N}{lq} + \frac{D}{v}$. This expression follows from the assumption that, at the beginning of the path to be traversed, the cars are already traveling at steady state velocity. Thus the approximate evacuation time equals the time it takes the "first" car to traverse the distance D plus the time it takes for the N cars to flow past the endpoint of the path, which is given by N/lq.

To illustrate the fact that the goals of maximizing traffic flow and maximizing velocity are out of sync, we calculate the highest value of K for which minimizing M would result in an equilibrium speed of v_{cruise}. This requires a formula for the equilibrium value v^* that solves the problem

$$\begin{gathered}\text{minimize } M(v) = W\tfrac{N(L+\beta v+\gamma v^2)}{lv} + (1-W)\tfrac{D}{v} \\ \text{subject to } v \leq v_{cruise}.\end{gathered}$$

Using the methods of the previous section, we obtained the formula for $M(v)$ by using Equation 1, the fundamental law of traffic flow $q = kv$, and the relationship $k = 1/s$. Differentiating with respect to v, setting the result equal to zero, and solving for velocity yields

$$v^* = min\left\{ v_{cruise}, \sqrt{\frac{1}{\gamma}\left[L + \frac{(1-W)}{W}Dl/N\right]} \right\}.$$

In order to have $v_{cruise} = \sqrt{\frac{1}{\gamma}\left[L + \frac{(1-W)}{W}Dl/N\right]}$, we would need

$$W = \left(1 + \frac{N}{Dl}(v_{cruise}^2\gamma - L)\right)^{-1}.$$

Using the empirical values $N = 160,000$ cars, $D = 633,600$ ft (120 miles), $l = 2$ lanes, $v_{cruise} = 88$ ft/sec, $\gamma = .0115$ sec^2/ft and $L = 10$ ft, we obtain a value of $W \approx 1/11$. Clearly, then, minimizing evacuation time in situations involving heavy traffic flow is incompatible with allowing drivers to travel at cruise velocity while leaving a safe stopping distance between themselves and the driver in front of them.

The relative weight of the two summands in the metric M depends on the number of cars to be evacuated and the travel distance, as well as the value of K. Because evacuation time is the most reasonable measure to use in evaluating the success of an evacuation, and because minimizing it is slightly different than the problem of maximizing flow, we state analytical solutions for minimum evacuation time and optimal velocity as a function of the number of cars to be evacuated. It turns out that for N large, the problem of evacuation time minimization is essentially equivalent to the problem of flow maximization, a result that will be used later on in the section on parallel paths.

5.3 Computing Minimum Evacuation Time

From the fact that $T = 2M$ when $K = 1/2$, we obtain the following values for the optimal flow, optimal density, and steady state velocity that minimize evacuation time:

$$q^* = k^*v^*$$

$$k^* = \frac{\beta\gamma^{1/2}[L + Dl/N]^{-1/2} - 2\gamma\frac{[L+\frac{1}{2}Dl/N]}{[L+Dl/N]}}{[\beta^2 - 4\gamma L] - \gamma\frac{[Dl/N]^2}{[L+Dl/N]}}$$

$$v^* = \sqrt{\frac{1}{\gamma}\left[L + Dl/N\right]}$$

The minimum evacuation time is

$$T^* = \frac{N}{lq^*} + \frac{D}{v^*}.$$

5.4 Predictions of the Steady State Evacuation Time Model

The chief virtue of the steady state evacuation time model is that it takes a reasonable, empirically tested model of car following and builds from it an estimate for best case evacuation time as a function of the number of automobiles to be evacuated and the number of lanes of traffic. Naturally, the predictions of the model vary based on our assumptions about the values of the constants L, β, γ, and D. For simplicity, we use the values $L = 10$ ft, $\beta = 1$ sec, $\gamma = .0115$ sec^2/ft, and $D = 633,600$ ft (120 mi) that we have used in previous calculations, and assume for the time being that the number of lanes $l = 2$.

To demonstrate visually that the evacuation time minimization problem is effectively equivalent to the flow maximization problem for N large, we compare the functions $T_{min}(N)$ (the minimum evacuation time) and $T_{flow}(N)$ (the evacuation time associated with the maximal flow traffic velocity and density). In the graph shown on the following page, the upper line is the the graph of $T_{flow}(N)$ and the lower line is the graph of $T_{min}(N)$. In fact, it can be show analytically that $\lim_{N\to\infty} \frac{T_{flow}(N)}{T_{min}(N)} = 1$.

The predicted evacuation time for $N = 160,000$ of slightly over 40 hours seems fairly reasonalble. We can evaluate the impact of the strategy of converting Highway I-26 to four lanes by setting $l = 4$ in the equation for minimum evacuation time. This supposition yields the result $T \approx 23$ hours. For the steady state model, this prediction makes sense, since the model does not deal with the effect of a bottleneck that will occur when Columbia is bombarded by the evacuees. Naturally, the bottleneck could be compounded by the use of four lanes instead of two; we will address the bottleneck problem in more detail later by analyzing the relationship between the flow carrying capacity of the highway and the flow carrying capacity of the Columbia entranceway. On balance, however, it would seem from the strong prediction of the steady state model that the strategy of doubling the number of lanes on I-26 would lead to a net decrease in the evacuation time.

As mentioned previously, the steady state model has the drawback of not taking into account the effect of the variance in traveling velocities in areas of high density on traffic flow. As a framework to address this problem, we introduce the following model.

graph 1

6 The One-Dimensional Cellular Automata Model

6.1 Motivation

In heavy traffic cars make repeated stops and starts, and the timing of these is somewhat arbitrary. This randomness has a significant effect on traffic flow, and so a good model of heavy traffic should take it into account. We also desire a model which is simple enough that an explicit formula can be computed for its velocity.

6.2 Development

A single-lane road is divided into a series of sections (cells) of the same length, each slightly longer than the length of one car. Each cell either contains one car or does not contain any cars. A car is said to be *blocked* if the cell directly in front of it is occupied. A fixed probability p is given. At each time state, cars move according to the following rules:

- If a car is not blocked then it advances to the next cell with probability p.
- If a car is blocked then it does not move.

Note that the decisions of each driver to move forward are made independently.

Question 6.1. What is the relationship between traffic density and traffic velocity in this model?

We will obtain an exact answer to this question. First we define the one-dimensional cellular automata model more formally.

A traffic configuration may be represented by a function

$$f\colon \mathbb{Z} \to \{0,1\}$$

where $f(k) = 1$ means that there is car in the kth cell and $f(k) = 0$ means that there is not. Since our model involves randomness, what we are really interested in are probability distributions on the set of all such functions. Such distributions are called *binary processes*.

Given a process X, define a process $I_p(X)$ according to the following rule: If

$$(X(i), X(i+1)) = (1, 0)$$

then

$$(I_p(X)(i), I_p(X)(i+1)) = \begin{cases} (0,1) & \text{with probability } p \\ (1,0) & \text{with probability } 1-p \end{cases}$$

This rule is identical to the traffic flow rule given above. Thus, if X represents the traffic configuration at time t, $I_p(X)$ gives the traffic configuration at time $t+1$.

We are interested in what the traffic configuration looks like after several iterations of I. For convenience, write $I_p^n(X)$ to mean I_p applied n times to X. The formula for traffic velocity in terms of density comes from the following theorem[1]:

Theorem 1. *Suppose X is a binary process of density d. Let $\mathcal{M}_{p,d}$ denote the Markov chain with the following transition probabilities:*

$$0 \longrightarrow \begin{cases} 0 & \text{w/ prob. } 1 - \frac{1-\sqrt{1-4d(1-d)p}}{2p(1-d)} \\ 1 & \text{w/ prob. } \frac{1-\sqrt{1-4d(1-d)p}}{2p(1-d)} \end{cases}$$

$$1 \longrightarrow \begin{cases} 0 & \text{w/ prob. } \frac{1-\sqrt{1-4d(1-d)p}}{2pd} \\ 1 & \text{w/ prob. } 1 - \frac{1-\sqrt{1-4d(1-d)p}}{2pd} \end{cases}$$

The sequence of processes $X, I_p(X), I_p^2(X), I_p^3(X), \ldots$ converges to $\mathcal{M}_{p,d}$.

Here "density" means the frequency with which 1's appear–this is analogous to the average number of cars per cell on the road. (See the appendix for a complete definition, as well as an explanation of the notion of convergence for binary processes.) Essentially, this theorem tells what the traffic configuration looks like after a long period of time has elapsed.

Knowing the transition probabilities allows us to easily compute the average velocity of the cars in $\mathcal{M}_{p,d}$: the average velocity is the likelihood that a randomly chosen car is not blocked and advances to the next cell at

[1] This theorem is the result of previous research done by one of the authors of this paper. The appendix on binary processes (which was written during the contest period) gives a relatively complete proof of the result.

the next time state.

$$
\begin{aligned}
v &= \mathbb{P}\left(I(\mathcal{M}_{p,d}(i)) = 0 \,|\mathcal{M}_{p,d}(i) = 1\right) \\
&= \mathbb{P}\left(\mathcal{M}_{p,d}(i+1) = 0 \,|\mathcal{M}_{p,d}(i) = 1\right) \\
&\quad \cdot \mathbb{P}\left(I(\mathcal{M}_{p,d}(i)) = 0 \,|\mathcal{M}_{p,d}(i) = 1 \text{ and } \mathcal{M}_{p,d}(i+1) = 0\right) \\
&= \left(\frac{1-\sqrt{1-4d(1-d)p}}{2pd}\right) \cdot p \\
&= \frac{1-\sqrt{1-4d(1-d)p}}{2d}
\end{aligned}
$$

Thus,

Corollary 1. In the one-dimensional cellular automata model with a starting configuration of density d, the average velocity tends to

$$v = \frac{1-\sqrt{1-4d(1-d)p}}{2d}$$

as t tends to infinity.

6.3 Relevance

Many factors cause drivers to have different reaction times: distractions, other passengers, level of attention, physical handicaps, etc. As can be seen in a bird's-eye view of heavy traffic, these vagaries alone can cause small traffic jams that appear spontaneously and then dissipate. Any model that does not take this random behavior into account is likely to overpredict the speed of traffic. The one-dimensional cellular automata model accounts for independently random behavior, and it is simple enough that velocity may be computed quickly.

6.4 Limitations

The main limitation of the model is that it does not accurately simulate high-speed traffic. It does not take into account following distance, and the stop-and-start model of car movement is not accurate when traffic is sparse. The model is best applied to slow traffic flows (say, under 15 miles per hour) where the drivers must make frequent stops.

6.5 The One-Dimensional Cellular Automata Model Applied at Low Speeds

To apply the one-dimensional cellular automata model at low speeds requires setting three variables:

Δx	the size of one cell
Δt	the length of one time interval
p	the movement probability

Δx would be equal to the space taken up by a car in a tight traffic jam, so we set it to 15 ft, slightly longer than the length of most cars. Δt should be the time taken by the shortest time taken by drivers to move into the space in front of them after it is vacated. Based on general observation we take this to be $0.5s$. p represents the proportion of drivers who indeed move at close to this speed; we let $p = 0.85$.

We recall the formula from the previous section, inserting the factor $(\Delta x/\Delta t)$ to convert from cells per time-state to feet per second.

$$v = \left(\frac{\Delta x}{\Delta t}\right) \frac{1 - \sqrt{1 - 4d(1-d)p}}{2d}$$

Density d is given in cars per cell. This corresponds to the usual notion of density via

$$d = (15 \text{ ft})K$$

and to the notion of *occupancy* (the proportion of road length taken up by cars) by

$$d = \frac{15\text{ft}}{10\text{ft}} n = \frac{3n}{2}.$$

We calculate v for various values of n and k.

n	$K(\text{ ft}^{-1})$	d	v (mph)
0.60	0.060	0.90	1.90
0.55	0.055	0.83	3.55
0.50	0.050	0.75	5.43
0.45	0.045	0.68	7.51
0.40	0.040	0.60	9.73
0.35	0.035	0.53	11.88
0.30	0.030	0.45	13.67
0.25	0.025	0.38	14.98
0.20	0.020	0.30	15.86

7 Cellular automata simulation: one lane model

In order to explore the one-dimensional cellular automata model, a simple simulation was written in C++. The simulation (*simcell1.cc*, see appendix for source code) consists of a 5000-element long (circular) array of bits, with a "1" representing a car and a "0" representing a (car-sized) empty space. The array is initialized randomly based on a given value of the occupancy n: an element is initialized to "1" with probability n, or to "0" with probability $1-n$. (Recall that n is a unitless value defined as the proportion of highway length taken up by cars: $n = Lk$, where L is the length of a car in feet and k is the density in cars/foot). Once the array is initialized, it is then iterated over 5000 time cycles: on each cycle, a car will move forward with probability p if the square in front of it is empty. The flow q is calculated to be the number of cars N passing the end of the array divided by the number of time cycles, ie. $q = N/5000$, and thus the average velocity of an individual car in cells per time cycle is $v = q/n = N/5000n$.

The simulation was first used to verify the accuracy of the one-dimensional cellular automata equation:

$$v = \frac{1 - \sqrt{1 - 4n(1-n)p}}{2n}$$

where v is the average velocity of a car in cells per time cycle, n is the occupancy, and p is the probability that a car moves forward if it is not blocked by a car in front of it. The simulation was run for various values of the occupancy n and probability p, and simulation results were found to closely match the values predicted by the one-dimensional cellular automata equation (see table).

n	$v_{sim}(p = 1/2)$	$v_{exp}(p = 1/2)$	$v_{sim}(p = 3/4)$	$v_{exp}(p = 3/4)$
0.2	.4334	.4384	.6942	.6972
0.4	.3556	.3486	.5918	.5886
0.6	.2342	.2324	.3922	.3924
0.8	.1078	.1096	.1708	.1743

Table 1: Comparison of simulation and expected values of velocity

Having confirmed the correspondence between theory and simulation, we now consider what a reasonable value would be for the parameter p. Since (on a given time cycle) an unblocked car moves ahead with probability p, it is clear that p should be related to the mean and standard deviation of v_{cruise}, the velocity of a car unimpeded by traffic. To see the relationship, we note that the mean and standard deviation of a binary random variable are p and $\sqrt{p(1-p)}$, respectively. As above, we make reasonable assumptions for the mean and standard deviations of v_{cruise}: $\mu(v_{cruise}) = 60$ mi/hr, $\sigma(v_{cruise}) = 5$ mi/hr. By equating the values of μ/σ for the simulation and observation we obtain:

$$\begin{aligned} \frac{p}{\sqrt{p(1-p)}} &= \frac{60}{5} \longrightarrow \\ p &= 144/145 \end{aligned}$$

For given values of $\mu(v_{cruise})$ and $\sigma(v_{cruise})$, we have:

$$\begin{aligned} \frac{p}{\sqrt{p(1-p)}} &= \frac{\mu}{\sigma} \longrightarrow \\ p &= \frac{1}{1+\left(\frac{\sigma}{\mu}\right)^2} \end{aligned}$$

By equating the mean values of cruise velocity, we also obtain a relationship between the length of a cell and the time step: p cells per time cycle is equivalent to 88 ft/sec, so the length of a cell is 88 ft/sec, multiplied by the length of a time cycle, multiplied by p:

$$L = \mu t p$$

Assuming L equals one car length (10 ft), p equals 144/145, and μ=88 ft/sec, we obtain a time step of 0.113 sec.

We now use the model to predict how fast (on average) a car will move, as a function of the occupancy of traffic. This assumes that there is only one lane of traffic; we expand this model to multiple lanes in the next section. To do this, we consider the "relative velocity", v_{rel}, which is defined as the average velocity of a car in traffic divided by the (mean) cruise velocity. The average velocity is given by the one-dimensional cellular automata equation, and the cruise velocity is p cells per time cycle, so this gives us:

$$\begin{aligned} v_{rel} &= \frac{v_{avg}}{v_{cruise}} \\ &= \frac{1-\sqrt{1-4n(1-n)p}}{2pn} \end{aligned}$$

Using the value of $p = 144/145$ obtained above, we now calculate v_{rel} as a function of occupancy. We also use this value to calculate v_{avg} in feet per second (see table)

n	v_{rel}	v_{avg}
0.1	.9991	87.92
0.2	.9977	87.80
0.3	.9949	87.55
0.4	.9869	86.85
0.5	.9233	81.25
0.6	.6579	57.90
0.7	.4264	37.52
0.8	.2494	21.95
0.9	.1110	9.77

Table 2: Average and relative velocities as a function of occupancy

Thus the model predicts that, for low occupancy, the average velocity will be near the cruise velocity, but for occupancies greater than 0.5, the average velocity will be significantly lower than the cruise velocity. It should be noted that the cellular automata model does not take following distance into account; thus it tends to overestimate v_{avg} for high speeds, and is most accurate when the occupancy is high and velocity low.

Next we calculate the flow rate $q = nv_{avg}/L$, in cars per second, as a function of the occupancy (see table).

It is interesting to note that the flow rate has a maximum at $n = 0.5$, and is symmetric about this line. This follows from the symmetry of the simulation: each car movement can be thought of as switching a car with an empty space, so the movement of cars to the right is equivalent to the movement of holes to the left with equal probabilities.

This model fails, however, to give a reasonable value for the maximum flow rate: a flow rate of 4.063 cars per second equates to 14600 cars/hr, approximately seven times a reasonable maximum rate (CITE SOURCE!). The reason for this discrepancy is the use of a cell size equal to the car length, which is a correct approximation of reality only as the car velocities approach zero and the occupancy approaches 1. In order to correctly approximate the

n	q
0.1	.8792
0.2	1.756
0.3	2.626
0.4	3.474
0.5	4.063
0.6	3.474
0.7	2.626
0.8	1.756
0.9	.8792

Table 3: Flow rate as a function of occupancy

maximum flow rate at occupancy=0.5, we must assume a larger value of the cell size, one which takes following distance into account. Reasonable assumptions are that the cell size is equal to car length plus following distance, and following distance is proportional to the velocity. Assuming a 1 second following distance (more precise values could be determined empirically), we obtain the following expression for cell size:

$$C = L + v_{avg}(1 \text{ s})$$

However, we do not know the value of the velocity v_{avg} until we use the cell size to obtain it. For n large, however, we can assume that v_{avg} is approximately equal to v_{cruise}, and find an upper bound on the necessary cell size:

$$\begin{aligned} C &= L + v_{cruise}(1 \text{ s}) \\ &= 98 \text{ ft} \end{aligned}$$

We then divide the flow rate (originally computed) by the increase in cell size to obtain a more reasonable flow rate:

$$\begin{aligned} q &= \frac{4.063 \text{ cars/second}}{98 \text{ ft}/10 \text{ ft}} \\ &= .415 \text{ cars/second} \end{aligned}$$

This gives us a flow rate of approximately 1500 cars/hr, a much more reasonable figure. However, this is likely to be an underestimate of the actual flow rate: to compute the flow rate more precisely, we must find a method of computing the correct cell size before finding the velocity. We address this problem further in the combined model.

8 Cellular automata simulation: two lane model

Our next step was to expand the one dimensional cellular automata model into a simulation of a two lane traffic flow. The simulation (simcell2.cc, see appendix for source code) consists of a two-dimensional (1000x2) array of bits, with a "1" representing a car and a "0" representing an empty space. As in the one lane model, the array is initialized randomly: each element is intialized to "1" with probability n and "0" otherwise. The array is then iterated over 1000 time cycles: on each cycle, a car will move forward with probability p if the square in front of it is empty, or perform a "lane switch" with probability p if the following conditions hold:

1. Cell in front of car is occupied (thus it cannot move forward)
2. Cell beside car is unoccupied
3. Cell diagonally forward from car is unoccupied

A car performing a "lane switch" moves one square forward and changes lanes. For example, consider the following arrangment of cars:

12350
00400

In the next time step, cars 4 and 5 will each move forward with probability p. Car 1 will perform a lane switch with probability p. Cars 2 and 3 are blocked and cannot move. Assuming that all of the unblocked cars choose to move forward, we obtain the following configuration:

02305
01040

As in the one-lane simulation, the flow q is calculated to be the number of cars N passing the end of the array divided by the number of time cycles. Another similarity to the one lane model is the presence of symmetry: the simulation is unchanged if we consider the holes as cars (moving left) and the cars as holes (moving right). As a result of this symmetry, the flow rate $q(n)$ is maximized at $n = 0.5$, and $q(n) = q(1 - n)$ for all n. This condition does not necessarily hold when parameters for the model are changed: for example, if different cars were assigned different values of p as in the following section.

The two-lane simulation suffers from the same flaw as the one-lane simulation: it uses a cell size equal to the car length, which is only correct for cars at very high densities and low velocities. As a result, we do not use the two lane simulation to compute maximum flow rate. However, since the choice of cell size affects the flow rate by a constant factor, we can compare the flow rates obtained when varying parameters of the simulation. In particular, we can use this model to examine how the flow rate changes with the variance of speeds, as discussed in the next section.

9 "But I want to bring my boat": velocity variance and the restriction of vehicle types

We now use our two cellular automata models in order to examine the effects of variance in velocities. There are two main types of variance we must consider: variance of traveling velocity (the random fluctuations in the velocity of a single vehicle over time), and variance of mean velocity (the variation in the mean velocities of all vehicles). We denote these as σ_t^2 (traveling variance) and σ_m^2 (mean variance) respectively. For example, in our original one lane simulation, we assumed that all vehicles travel at the same mean velocity, v_{cruise}=60 mi/hr, so $\sigma_m = 0$. This is clearly an oversimplification: in real life, vehicles' mean velocities may vary significantly. We also assumed that each vehicle's speed fluctuates randomly with standard deviation 5 mi/hr, so $\sigma_t = 5$ mi/hr. This was reflected in the calculation of the probability p, since $p = 1/(1+(\sigma_t/\mu)^2)$. When we take σ_m into account, a different value of the probability p is assigned to each car. We determine this value of p using the following method:

1. Choose a value of the car's mean velocity μ randomly from the normal distribution with mean v_{cruise} and standard deviation σ_m.
2. Thus the car's traveling velocity will be normally distributed with a mean of μ and a standard deviation of σ_t.
3. The car's transition probability p is calculated as:

$$p = \frac{\mu}{v_{cruise} + \lambda\sigma_m}\left(\frac{1}{1+\left(\frac{\sigma_t}{\mu}\right)^2}\right)$$

where λ is a constant best determined empirically. We use $\lambda = 0$ for our simulations, thus presenting a conservative estimate of the change in flow rate as a function of σ_m.

Now we consider what effect σ_t and σ_m are likely to have on the velocity at a given occupancy. As Militzer (1998) states, "small perturbations that increase the density locally lead to a decrease in the flow, which then amplifies the initial perturbation. The *variance* in car velocities is what triggers this process, leading to the formation of traffic jams." (CITE SOURCE!) Considering the cars' movement as a directed random walk, it is clear that increasing σ_t increases the amount of randomness in the system. This will cause the cars to interact with each other (and hence, block each other's movement) more often, decreasing the average velocity in traffic. The effects of σ_m are even more dramatic: cars with low mean velocities will tend to impede the progress of the faster cars behind them. In the one lane model, it is impossible to pass a slow car, so cars will queue up behind it, and their progress will be reduced significantly. In the two lane model, faster cars are able to pass slower cars, reducing the impact of σ_m.

To more precisely examine the effects of these variables, we ran simulations with a number of different values of σ_m, σ_t, and the occupancy n. In particular, we first fixed $n = 0.5$, and varied both σ_m and σ_t from 0 to 15 mi/hr.

For each pair of values, we calculated an average flow rate (in cars/1000 time steps) for the one and two lane simulations. The one lane flow rate was doubled for comparison purposes: this is equivalent to a two lane flow with no lane switching allowed.

	$\sigma_t = 0$	$\sigma_t = 5$	$\sigma_t = 10$	$\sigma_t = 15$
$\sigma_m = 0$	979/976	923/908	856/830	805/776
$\sigma_m = 5$	822/757	817/746	790/729	753/683
$\sigma_m = 10$	699/537	691/518	659/485	642/469
$\sigma_m = 15$	588/393	569/366	540/319	518/292

Table 4: Two lane avg. flow / Twice one lane avg. flow

As can be seen from this table, the maximum flow for the two lane model is more than twice the maximum flow for the one lane model. For low values of σ_t and σ_m, the difference is negligible: allowing cars to switch lanes does not significantly increase the flow rate. For high values of variance, the two lane model has a significantly higher flow rate than the one lane model; this is as expected, since lane switching reduces the problem of queueing behind slow vehicles by allowing faster vehicles to pass them. For both the one lane and two lane models, the average flow rate decreases with increasing σ_t and with increasing σ_m. Each 5 mi/hr increase in σ_m resulted in an 11-16% decrease in flow rate (two lane model, $\sigma_t = 0$), while each 5 mi/hr increase in σ_t resulted in a 5-7% decrease in flow rate (two lane model, $\sigma_m = 0$). Thus the effects of both traveling variance and mean variance dramatically affect flow rate, and the effects of σ_m on flow rate are more significant than the effects of σ_t.

9.1 "So, can I bring my boat?"

We now consider how variations in vehicle type affect the values of σ_m and σ_t, and how this affects flow rate. It is clear that most large vehicles (such as boats, campers, semis, and motor homes) will travel at a slower rate than most normal cars. Thus, if a significant proportion of people bring large vehicles, this results in an increased σ_m and hence a lower flow rate. As a simplified approximation to this, assume there are two types of vehicles: fast cars ($\mu = \mu_1$) and slow trucks ($\mu = \mu_2$). If the proportion of slow trucks is given by $\aleph$, then we can calculate:

$$\begin{aligned}
\sigma_m^2 &= \aleph(\mu_2 - \bar{\mu})^2 + (1-\aleph)(\mu_1 - \bar{\mu})^2 \\
&= \aleph(\mu_2 - (\mu_1 - (\mu_1 - \mu_2)\aleph))^2 + (1-\aleph)(\mu_1 - (\mu_1 - \mu_2)\aleph))^2 \\
&= (\mu_1 - \mu_2)^2(\aleph^2(1-\aleph) + \aleph(1-\aleph)^2) \\
&= (\mu_1 - \mu_2)^2(\aleph)(1-\aleph)
\end{aligned}$$

Thus $\sigma_m = (\mu_1 - \mu_2)\sqrt{(\aleph)(1-\aleph)}$. We now assume that fast cars travel at $\mu_1 = 70$ mi/hr, and slow trucks travel at $\mu_2 = 50$ mi/hr, and find σ_m as a function of $\aleph$. The random fluctuations in vehicle speed are likely to depend more on the psychology of the driver than on the type of vehicle under consideration, so we assume a constant value of $\sigma_t = 5$ mi/hr. This allows us to linearly interpolate from the values given in the table, enabling us to find the flow rate (cars/1000 time cycles) for the two lane and one lane models as a function of the proportion of slow vehicles $\aleph$:

Thus the flow rate is decreased significantly by the presence of slow vehicles: if even 1% of vehicles are slow, the flow rate decreases by 5%, and if 10% of vehicles are slow, the flow rate decreases by 15%. The effects of σ_m are

$\aleph$	σ_t	flow rate	% reduction in flow
0	0	923/908	0/0
.01	1.99	881/844	4.6/7.0
.02	2.80	864/817	6.4/10.0
.05	4.36	831/767	10.0/15.5
.1	6.00	792/700	14.1/22.9
.2	8.00	741/609	19.7/32.9
.5	10.0	691/518	25.1/43.0

Table 5: Flow rate as a function of proportion of slow vehicles

magnified if vehicles are unable to pass slower vehicles, so if the highway went down to one lane at any point (due to construction or accidents, for example) this would reduce the flow rate even further. Thus we recommend restricting the allowed types of evcuation vehicles as follows: no large vehicles (vehicles which may potentially block multiple lanes), and no slow vehicles (vehicles with a significantly lower mean cruising velocity). Exceptions to the "no slow vehicles" rule may be made in two cases:

1. If a family has no other alternative but to drive a slow vehicle (ie. the family owns no fast vehicle), they are still allowed to evacuate.
2. Vehicles which are carrying a large number of people (such as buses) are allowed, since these may significantly reduce the number of vehicles to be evacuated.

Also, slow-moving vehicles should be required to stay in the right lane, allowing faster vehicles to pass by. This will help to reduce the impact of the slow vehicles on the flow rate of traffic. Additionally, families should be encouraged to take as few vehicles as possible in order to minimize the total number of cars to be evacuated; this will help to further reduce evacuation time.

10 The Space-Velocity Curve

In order to determine optimal traffic flow rates we need a good estimate of the relationship between velocity (v) and space per car (s) (or equivalently, velocity and occupancy (n)). We acheive this by combining two of the models previously discussed: the one-dimensional cellular automata model and the steady state model.

The space-velocity function will be denoted by $F(s)$, where s denotes space per car ($s = 1/K$) and $F(s)$ is given in miles per hour. We define $F(s)$ in sections:

$\mathbf{s} \leq \mathbf{15}$ ft:

When spacing is less than 15 feet there is essentially no traffic flow: $F(s) = 0$.

$\mathbf{15}$ ft $\leq \mathbf{s} \leq \mathbf{30}$ ft:

When spacing is between 15 feet and 30 feet, traffic travels at speeds between 0 and 12 miles per hour, and the one-dimensional cellular automata model is appropriate. We have previously obtained the following formula for the one-dimensional cellular automata model:

$$v = \left(\frac{15 \text{ ft}}{0.5 \text{ s}}\right) \frac{1 - \sqrt{1 - 4d(1-d)0.85}}{2d}.$$

Thus,

$$F(s) = (10.2 \text{ mph}) \frac{1 - \sqrt{1 - 4d(s)(1-d(s))0.85}}{d(s)}$$

where $d(s) = (15 \text{ ft})/s$.

$\mathbf{30}$ ft $\leq \mathbf{s} \leq \mathbf{140}$ ft:

Between 30- and 140-foot spacing, traffic travels between 12 mph and 55 mph, and the steady state model is appropriate. By the steady state equation,

$$s = \frac{1}{2D}F(s)^2 + \beta F(s) + L,$$

thus

$$F(s) = \frac{-\beta + \sqrt{\beta^2 - 4(L-s)\left(\frac{1}{2D}\right)}}{2\frac{1}{2D}}.$$

$L = 10$ ft (the length of a car); $\beta = 1$ s (reaction time); and we take the conservative estimate $D = 21 \text{ ft/s}^2$ for the deceleration rate.

$\mathbf{140}$ ft $\leq \mathbf{s}$:

Above 140-foot spacing traffic velocity will be equal to the speed limit.

$F(s)$ is graphed below (in miles per hour) at two different scales.

graph 2

11 Incoming Traffic Rates

As discussed earlier, the optimal flow of traffic through a route is determined only by the optimal flow through the smallest bottleneck along the route. However the time of travel (which is a more important measure for our purposes) is affected by other factors, including the rate of incoming traffic. If incoming traffic is heavy, congestion occurs at the beginning of the route, thus causing a decrease in velocity and an overall increase in travel time for each car.

How does congestion occur, and how much does it influence travel time? Consider the one-dimensional cellular automata model with $p = 1/2$. Represent the road by the real line, and let $F(x, t)$ denote the density of cars at point x on the road at time t. (For our purposes right now the cells and cars are infinitesimal in length.) Suppose that the initial configuration $F(x, 0)$ is given by the step function,

$$F(x,0) = \begin{cases} 1 \text{ if } x < 0 \\ 0 \text{ if } x \geq 0 \end{cases}$$

This is analogous to a traffic scenario where a dense line of cars is about to move onto an uncongested road.

We will omit units (ft, s, etc.) for the time being. By the formulas from "The One Dimensional Cellular Model" section, the velocity $v(x_0, t_0)$ at position x_0, time t_0 is given by

$$v(x_0, t_0) = \frac{1 - \sqrt{1 - 2F(x_0, t_0)(1 - F(x_0, t_0))}}{2F(x_0, t_0)}$$

while velocity must also be equal to the rate at which the number of cars past point x is increasing; that is,

$$v(x_0, t_0) = \frac{d}{dt}\left(\int_x^\infty F(x, t)dx\right)(t_0) .$$

Thus,

$$\frac{dF}{dt} = -\frac{dv}{dx} = -\frac{d}{dx}\left(\frac{1 - \sqrt{1 - 2F(1 - F)}}{2F}\right)$$

This is a partial differential equation whose unique solution is given by:

$$F(x,t) = \begin{cases} 1 & \text{if } x/t < -\frac{1}{2} \\ \frac{1}{2} - \frac{(x/t)}{\sqrt{2 - 4((x/t))^2}} & \text{if } -\frac{1}{2} \leq x/t \leq \frac{1}{2} \\ 0 & \text{if } \frac{1}{2} < x/t \end{cases}$$

Thus, after a steady influx of cars for a period of Δt, the graph of the resulting congestion is given by

$$\frac{1}{2} + \frac{\frac{x}{\Delta t}}{\sqrt{2 - 4\left(\frac{x}{\Delta t}\right)^2}},$$

and the congestion ends at position $x = \Delta t/2$.

This derivation has many possible uses. For our purposes we are only interested in the fact that the size of the the traffic pattern is *linear* with respect to Δt. In a real-world scenario, this means that if there is an influx of N cars onto a highway, the size of the resulting congestion is directly proportional to the number N. Likewise, the amount of time it takes for the congestion to dissipate is proportional to N.

This allows us to evaluate one of the proposals–staggering evacuation times for different counties. Suppose that there are n counties with populations $P_1, P_2, \ldots, P_n$. If all counties evacuate at the same time, the effect of the resulting traffic jam on total travel time is proportional to the product of the size of the jam and the amount of time before it dissipates:

$$\begin{aligned} \Delta T_{\text{travel time}} &= c_1 \cdot c_2 \left(P_1 + \ldots + P_n\right) \cdot c_3 \left(P_1 + \ldots + P_n\right) \\ &= c_1 c_2 c_3 \left(P_1 + \ldots + P_n\right)^2 \end{aligned}$$

for some constants c_1, c_2, c_3. If the evacuations are staggered, the effect of the jam is

$$\Delta T_{\text{travel time}} = c_1 c_2 c_3 P_1^2 + \ldots + c_1 c_2 c_3 P_n^2 = c_1 c_2 c_3 \left(P_1^2 + \ldots + P_n^2\right)$$

Now, certainly

$$P_1^2 + \ldots + P_n^2 < \left(P_1^2 + \ldots + P_n^2\right);$$

and unless one of the counties has a much larger population than the rest, the difference between these two values is highly significant. We therefore recommend the proposal for staggering counties.

12 The effects of merges and diverges on traffic flow

While the steady state model is a reasonably accurate predictor of traffic behavior on long homogeneous stretches of highway, we must also consider how to deal with the effects of road inhomogeneities: "merges" of two lanes into a single lane, and "diverges" of one lane into two lanes. To do so, we apply the principle of conservation of traffic flow as discussed in Kuhne and Michalopolous ("Continuum Flow Models"). Assuming that there are no sources or sinks in a region, this principle states that:

$$\frac{\partial q}{\partial x} + \frac{\partial k}{\partial t} = 0$$

where q is flow rate (cars/sec), and k is density (cars/ft), x is location (ft), and t is time (sec). Assuming that the merge or diverge occurs at a specific point x, we now consider how this junction behaves in the steady state: that is, for $\frac{\partial D(t)}{\partial t} = 0$. Given $\frac{\partial k}{\partial t} = 0$, this implies that $\frac{\partial q}{\partial x} = 0$. In other words, flow rate is conserved at a junction in the steady state. This implies that the sum of flow rates going into a junction is equal to the sum of flow rates going out of a junction. For a flow q_s diverging into flows q_1 and q_2, we know that $q_s = q_1 + q_2$, and this result also applies for flows q_1 and q_2 merging into a single flow q_s. If we are given the proportion P $(0 < P < 1)$ of the flow q_s going to (or coming from) q_1, and given either density (k_s or k_1), we can use the steady state model to solve for the other density as follows:

$$\begin{aligned} q_1 &= Pq_s \\ k_1 v_1 &= Pk_s v_s \\ k_1 v(k_1) &= Pk_s v(k_s) \end{aligned}$$

From the steady state model, and using the given values of all constants, we know:

$$\begin{aligned} v(k) &= 88 \text{ ft/sec } (k < .0056) \\ v(k) &= 21.7(\sqrt{.08 + \frac{.092}{k}} - 1) \ (.0056 < k < .1) \end{aligned}$$

Assuming both densities are greater than the free-travel density $k = .0056$, we can set:

$$k_1(\sqrt{.08 + \frac{.092}{k_1}} - 1) = Pk_s(\sqrt{.08 + \frac{.092}{k_s}} - 1)$$

Given either k_s or k_1, we can solve numerically for the other with relative ease. Then we can find the velocities associated with each density using the above expression for $v(k)$. There are several subtleties involved: solving the above equation gives two potential values of density, so we make the reasonable assumption that the density is greater on the single lane side of the junction (ie. density increases at a merge and decreases at a diverge). Also, it is possible that solving the above equation produces a velocity v_1 which is larger than v_{cruise}. This problem can be solved by setting $v_1 = v_{cruise}$, and calculating $n_1 = q_1/v_1$.

However, the above discussion does not consider a crucial component of the problem: how is the steady state flow rate determined on a path with merges and diverges? Following the discussion of Daganzo (1997), we consider a “bottleneck” to be defined as an inhomogeneous location (such as a merge or diverge) where queues can form and persist with free flow downstream. The “bottleneck capacity” is the maximum flow rate through the bottleneck; we assume for that bottleneck capacity is constant rather than time-varying, an assumption which Daganzo also makes. If the flow rate on one side of a bottleneck exceeds this capacity, then a queue will form, dissipating predictably when the flow rate decreases. This has interesting results for the steady state model: assume that a steady state flow greater than the bottleneck capacity attempts to enter the bottleneck. This will result in a queue size which continually increases, until it eventually stretches all the way back to its origin. At this point, the steady state flow is blocked by the queue of cars, and decreases to the bottleneck capacity. As a result, we can conclude that the maximum steady state flow rate from point A to point B along a given path is equal to the minimum bottleneck capacity of all bottlenecks along that route.

12.1 Parallel paths, and applications to evacuation strategy

We now consider the case when there are multiple parallel paths $p_1 \ldots p_m$ from point A to point B. Each path p_i has a given bottleneck capacity c_i equal to the minimum of all bottleneck capacities on that route. The maximum steady state flow rate from point A to point B is equal to the minimum of three quantities: the bottleneck capacity of the diverge at point A, the bottleneck capacity of the merge at point B, and the sum of the bottleneck capacities of all paths p_i. If a given path has no bottlenecks, its capacity is equal to the maximum flow rate predicted by the steady state model for that path.

This general framework can be applied to many of the strategies considered by the state officials of South Carolina. We focus here on maximizing flow rate rather than minimizing evacuation time: this is a reasonable assumption since, when the number of cars to be evacuated is very large, maximizing flow gives us an evacuation time near the minimum. We first express a general formula for the maximum flow rate q_{max} from Charleston to Columbia:

$$q_{max} = \min((\sum_i q_i), c_0, c_f)$$

where q_i is the maximum flow rate of the ith path, c_0 is the bottleneck capacity of Charleston, and c_f is the bottleneck capacity of Columbia. The flow rates q_i are defined by the following equation:

$$q_i = \min(b_1 \ldots b_n, q_{i,ss})$$

where $b_1 \ldots b_n$ are the capacities of any bottlenecks along the given route, and $q_{i,ss}$ is the maximum flow along that route predicted by the steady state model.

We first consider the evacuation situation with no strategies implemented, and assume no bottlenecks along I-26. Denoting the steady state value $q_{I-26,ss}$ by q_I, this gives us $q_{max} = \min(q_I, c_0, c_f)$. The important question, now, is which of these factors limits q_{max}. We have determined that $q_I \approx 2000$ cars/hr, but we will not be able to achieve this rate if either $c_0 < 2000$ (traffic jam in Charleston) or $c_f < 2000$ (traffic jam in Columbia). It is a reasonable assumption that, since the traffic in Columbia is split into three different roads, this will result in less congestion than everyone merging onto I-26 in Charleston. Hence we assume $c_0 < c_f$. This implies that c_0 is the limiting factor if $c_0 < 2000$, and q_I is the limiting factor if $c_0 > 2000$. The value of c_0 is best determined empirically, perhaps by extrapolation from Charleston rush hour traffic data, or by careful examination of traffic data from the last hurricane. We now consider the effects of implementing each strategy.

First, the plan for turning the two coastal-bound lanes of I-26 into two lanes of Columbia-bound traffic clearly doubles q_I to 4000 cars/hr. It is likely to increase c_0 as well, since cars can be directed to two different paths onto I-26 and are thus less likely to interfere with the merging of cars going on the other set of lanes. On the other hand, this implies that twice as many cars will be entering the Columbia area simultaneously, and since the capacity c_f is unchanged, it is possible that this capacity may become the limiting factor. It may be possible to increase c_f by rerouting some the extra traffic to avoid Columbia, or even turning around traffic on some of the highways leading out of Columbia. Thus this strategy is likely to improve evacuation traffic flow, but the extent of success will be affected by the relative values of q_I, c_0, and c_f. If c_0 and c_f are larger than twice q_I, the maximum flow rate will be doubled. If c_0 is the limiting factor, maximum flow will also increase, though possibly to a lesser extent than double. If c_f is the limiting factor, maximum flow may not increase at all, unless the extra traffic is rerouted. Nevertheless, it is likely that (since, as we argued, c_f was not likely to be the limiting factor before) a significant improvement will result, and the strategy should be implemented.

A similar argument applies to turning around the traffic on the smaller highways extending inland from the coast. Each highway will add some capacity to the total $\sum_i q_i$, increasing this term. It is likely, however, that each highway's capacity will be significantly less than q_I. Increasing the number of usable highways has unclear effects on c_0. If residents' evacuation patterns are carefully directed it may increase this capacity by spreading out the residents of Charleston to different roads. If not, the lack of organization may lead to people choosing their evacuation routes arbitrarily, and crossing evacuation routes may lead to traffic jams. As for the reversal of I-26, the reversal of smaller highways does not affect the value of c_f (unless crossing evacuation routes becomes a problem in Columbia). More importantly, the interactions between highways (merges, diverges, etc.) may lead to bottlenecks on each highway, further reducing the highway's capacity. In fact, interactions between these highways and I-26 could cause bottlenecks which slow the flow rate of I-26, offsetting the extra capacity of the smaller highways or even causing a significant problem. Thus it is safer not to turn around traffic on the secondary highways, or to encourage using these as evacuation routes. Selected highways might be used as evacuation routes, and the traffic on these turned around, assuming that the following conditions are met:

1. High capacity (ie. using the highway is worthwhile)
2. Low potential for traffic conflicts with other highways (especially with I-26)
3. Careful direction of Charleston traffic to the secondary highway, and Columbia traffic from the secondary highway, to minimize crossing of evacuation routes and other potential bottlenecks.
4. Multiple (at least two) *adjacent* lanes, since otherwise queues will form behind slow moving vehicles.

We now consider the strategy of establishing temporary shelters in Columbia, to reduce the traffic leaving Columbia. This strategy could be useful if only some of the cars are directed into Columbia: thus the flow of traffic in the Columbia area would be split into four streams rather than three, possibly increasing the value of c_f. Nevertheless, we hesitate to recommend this strategy, since the actual effects are likely to be the opposite. Evacuees entering Columbia (already a bustling city of 500,000 people) are likely to create a large amount of traffic within the city. This traffic congestion will make it difficult for traffic to enter the city, resulting in a major traffic bottleneck. If careful regulation of traffic is not performed, more people will attempt to stay in Columbia than the amount of available housing, and the frantic attempts of individuals to procure housing will exacerbate the traffic bottleneck. Hence it is most likely that c_f will decrease significantly, probably becoming the limiting factor on maximum flow rate. Unless extreme care is taken to regulate the number of cars entering Columbia, and to reduce the extent to which Columbian traffic impairs the influx of refugees, this strategy is likely to result in disaster.

Next, we briefly consider the question of staggering traffic flows. As we showed in a previous section, staggering the evacuation is likely to reduce the time it takes an average car to travel from Charleston to Columbia, while leaving the value of the steady state flow rate q_I unchanged. Thus we concluded that staggering the evacuation will decrease total evacuation time; we show here that staggering may also increase the maximum flow rate. This follows since staggering the evacuation decreases the number of cars that are traveling the city toward I-26 at any one time, reducing the size of the c_o bottleneck. Increasing the capacity c_o, however, will only increase flow rate if c_o is the limiting factor. If the steady state flow rate q_I is the limiting factor, then flow rate will be unchanged. Nevertheless, the reduction in evacuation time makes a staggering strategy worthwhile.

Lastly, we consider the impact of evacuees from Florida and Georgia, and their potential to compund traffic problems. The out-of-state evacuees clearly add to the number of cars to be evacuated, and since evacuation time is proportional to number of cars over flow rate, these evacuees will add to the total evacuation time unless they take a route which does not intersect the paths of the South Carolina evacuees. However, it is very hard to constrain the routes of the out-of-state evacuees, since they may come from a variety of paths, and are unlikely to be informed of the state's evacuation procedures. In particular, two major bottlenecks are likely to occur:

at the intersection of I-26 and I-95, and at the intersection of I-95 with I-20 and U.S. 501. If a large number of cars from I-95 attempt to go northwest on I-26 toward Columbia, this merge is likely to become a bottleneck. As a result, q_{I-26} will no longer be equal to $q_{I-26,ss} = 2000$ cars/hr, but will instead be equal to the capacity of the I-26/I-95 bottleneck. This is likely to significantly reduce q_{I-26}, and makes it extremely likely that q_{I-26} will be the limiting factor of maximum traffic flow. A similar argument suggests that I-95 traffic will impede the flow of traffic west from Myrtle Beach, by causing a bottleneck at the I-95/I-20 junction. Traffic flow from Myrtle Beach is less than from Charleston, and many of the cars from I-95 may have already exited at I-26, so the bottleneck at the I-20 junction is likely to be less severe than at I-26. Nevertheless, it is clear that the flow of evacuees from Florida and Georgia has the potential to dramatically reduce the success of the evacuation. This problem may be reduced by careful regulation of the interactions between highway traffic flows, but is difficult to fully solve.

13 Newspaper Article and Conclusion

Hurricanes pose a serious threat to citizens on the South Carolina coastline, as well as other beach dwellers in Florida, Georgia, and other neighboring states. In 1999, the evacuation effort preceeding the expected landfall of Hurricane Floyd led to a monumental traffic jam that posed other, also serious, problems to the more than 500,000 commuters who fled the coastline and headed for the safe haven of Columbia. Several strategies have been proposed to avoid a future repeat of this traffic disaster.

First, it has been suggested that the two coastal bound lanes of I-26 be turned into two lanes of Columbia bound traffic. A second strategy would involve staggering the evacuation of the coastal counties over some time period consistent with how hurricanes affect the coast, instead of all at once. Third, the state might turn around traffic flow on several of the smaller highways besides I-26 that extend inland from the coast. The fourth strategy under consideration is a plan to establish more temporary shelters in Columbia. Finally, the state is considering placing restrictions on the type and number of vehicles that can be brought to the coast.

In the interest of the public, we have developed and tested several mathematical models of traffic flow to determine the efficacy of each proposal. On balance, they suggest that the first strategy is sound, and should be implemented. Although doubling the number of lanes will not necessarily cut the evacuation time in half, or even double the flow rate on I-26 away from the coast, it will significantly improve the evacuation time under almost any weather conditions.

Our models suggest that staggering the evacuation of different counties is also, on balance, a good idea. Taking such action, on one hand, will reduce the severity of the bottleneck that occurs when the masse of evacuees reaches Columbia, and on the other hand, could potentially increase average traffic speed without significantly increasing traffic density. The net effect of implementing this strategy will likely be an overall decrease in coastal evacuation time.

The next strategy, which suggests turning traffic around on several smaller highways, is not so easy to recommend. The main reason for this is that the unorganized evacuation attempts of many people on frequently intersecting secondary roads is a recipe for inefficiency. In places where these roads intersect I-26, the merging of a heightened volume of secondary road traffic is sure to cause bottlenecks on the interstate that could significantly impede flow. To make a strategy of turning around traffic on secondary roads workable, the state would have to use only roads that have a high capacity, at least two lanes, and a low potential for traffic conflicts with other highways. This would require competent traffic management directed at avoiding bottlenecks and moving Charleston traffic to Columbia with as few evacuation route conflicts as possible.

The fourth proposal, of establishing more temporary shelters in Columbia, is a poor idea. Because it is assumed that travellers are relatively safe once they reach Columbia, the main objective of the evacuation effort should be minimizing the transit time to Columbia and the surrounding area. It is fairly clear that increasing the number of temporary shelters in Columbia would (A) lead to an increased volume of traffic to the city by raising expectations that there will be free beds there, and (B) exacerbate the traffic problem in the city itself due to an increased demand for parking. Together, these two factors are sure to worsen the bottleneck caused by I-26 traffic entering Columbia, and would probably increase the total evacuation time by decreasing the traffic flow on the interstate.

The final proposal of placing limitations on the number and types of vehicles that can be brought to the beach is reasonable. Families with several cars should be discouraged from bringing all of their vehicles, and perhaps required to register with the state if the latter is their intention. Large, cumbersome vehicles such as motor homes should be discouraged, unless they are a family's only option. Although buses slow down traffic, they are beneficial because they appreciably decrease the overall number of drivers. In all cases, slow moving vehicles should be required to travel in the right lane during the evacuation attempt.

In addition to the strategies mentioned above, commuters in the 1999 evacuation were acutely aware of the effect on traffic flow produced by coastal residents of Georgia and Florida travelling up I-95. We have concluded that, when high volume traffic flows such as these compete for the same traffic pipeline, the nearly inevitable result is a bottleneck. A reasonable solution to this problem would be to bar I-95 traffic from merging onto I-26, and instead encourage and assist drivers on I-95 to use the more prominent, inland bound secondary roads connected to that interstate.

To conclude, we think that combining the more successful strategies suggested could lead to a substantial reduction in evacuation time, which, in our view, is the primary measure of evacuation success. Minimizing the number of accidents that occur in en route is also important, but our models directed at the former goal do not make compromises with respect to the latter objective. In fact, the problem of minimizing accidents is chiefly taken care of by ensuring that traffic flow is as orderly, and efficient, as possble.

14 Appendix 1

This appendix contains definitions and proofs for the section "THE ONE-DIMENSIONAL CELLULAR AUTOMATA MODEL."

Definition 14.1. A *binary process* X is a probability distribution on binary sequences

$$F: \mathbb{Z} \to \{0, 1\}.$$

Equivalently, X is a sequence of random variables $X(i)$ on the same probability space with range $\{0, 1\}$.

Definition 14.2. A binary process X is of density d if

$$\lim_{n\to\infty} \frac{\sum_{k=1}^{n} X(k)}{n} = d = \lim_{n\to\infty} \frac{\sum_{k=-n}^{-1} X(k)}{n}$$

where lim denotes convergence in probability.

Definition 14.3. Let $\mathcal{J}(X, Y)$ denote the set of joint measures of X and Y.

Definition 14.4. Suppose X and Y are binary processes and $Z \in \mathcal{J}(X, Y)$. Define

$$d_Z(X, Y) = \limsup_{n\to\infty} \mathbb{E}_Z \left(\frac{\sum_{i=-n}^{n} |X_i - Y_i|}{2n+1} \right)$$

Essentially, $d_Z(X, Y)$ measures the frequency with which the terms of X and Y disagree under the joint measure Z. This allows us to define a pseudometric on binary processes:

Definition 14.5.

$$\overline{d}(X, Y) := \min_{Z \in \mathcal{J}(X,Y)} d_Z(X, Y)$$

Since $\overline{d}$ is a pseudometric ($\overline{d}(X, Y) = \overline{d}(Y, X)$ and $\overline{d}(X, Y) + \overline{d}(Y, Z) \geq \overline{d}(X, Z)$), so we may speak of convergence under $\overline{d}$.

Theorem 2. Suppose X is a binary process of density d. Then $I_p^n(X)$ converges to $\mathcal{M}_{p,d}$ under the $\overline{d}$ psuedometric as $n \to \infty$.

Proof. Proving this theorem will require some lemmas. First we must show that $\mathcal{M}_{p,d}$ is indeed of density d and that it is invariant under I_p.

Lemma 14.6. *$\mathcal{M}_{p,d}$ is of density d.*

Proof. The frequencies of 0's and 1's in the Markov chain $\mathcal{M}_{p,d}$ are given by the eigenvector of the transition matrix

$$M = \begin{bmatrix} 1-\frac{1-\sqrt{1-4d(1-d)p}}{2p(1-d)} & \frac{1-\sqrt{1-4d(1-d)p}}{2pd} \\ \frac{1-\sqrt{1-4d(1-d)p}}{2p(1-d)} & 1-\frac{1-\sqrt{1-4d(1-d)p}}{2pd} \end{bmatrix}.$$

corresponding to the eigenvalue 1. Let

$$x = \frac{1-\sqrt{1-4d(1-d)p}}{2pd(1-d)}$$

then

$$\begin{aligned} M\begin{bmatrix} 1-d \\ d \end{bmatrix} &= \begin{bmatrix} 1-xd & x(1-d) \\ xd & 1-x(1-d) \end{bmatrix}\begin{bmatrix} 1-d \\ d \end{bmatrix} \\ &= \begin{bmatrix} (1-xd)(1-d)+xd(1-d) \\ xd(1-d)+(1-x(1-d))d \end{bmatrix} \\ &= \begin{bmatrix} 1-d \\ d \end{bmatrix}. \end{aligned}$$

Thus $\begin{bmatrix} 1-d \\ d \end{bmatrix}$ is the frequency vector for $\mathcal{M}_{p,d}$, and the density of $\mathcal{M}_{p,d}$ (the frequency of 1's) is indeed d. □

Lemma 14.7. *The distributions $\mathcal{M}_{p,d}$ and $I_p(\mathcal{M}_{p,d})$ are identical.*

Proof. The first observation is that if X is an n-state Markov chain, then $I_p(X)$ is an $(n+1)$-state Markov chain. This is clear intuitively, since the functional I_p only introduces dependences between adjacent terms. Therefore $I_p(\mathcal{M}_{p,d})$ is a 2-state Markov chain. To prove the lemma it will suffice to show that the 2-state transition probabilities for $I_p(\mathcal{M}_{p,d})$ are the same as those for $\mathcal{M}_{p,d}$.

Let

$$\begin{bmatrix} d_0 \\ d_1 \end{bmatrix} = \begin{bmatrix} 1-d \\ d \end{bmatrix}$$

$$\begin{bmatrix} m_{00} & m_{01} \\ m_{10} & m_{11} \end{bmatrix} = M = \begin{bmatrix} 1-xd & x(1-d) \\ xd & 1-x(1-d) \end{bmatrix}$$

The following algebraic relations hold:

$$\frac{m_{11}m_{00}}{m_{10}m_{01}} = 1 - p \qquad d_0 + d_1 = 1$$

$$m_{00} + m_{01} = 1 \qquad m_{10} + m_{11} = 1$$

$$d_0 m_{01} = d_1 m_{10}$$

Let $(b_1, b_2, \ldots, b_n)$ denote a binary sequence. From the Markov chain property we have

$$\mathbb{P}\left((\mathcal{M}_{p,d}(1), \mathcal{M}_{p,d}(2), \ldots, \mathcal{M}_{p,d}(n)) = (b_1, b_2, \ldots, b_n)\right) = d_{b_1} m_{b_1 b_2} m_{b_2 b_3} \ldots m_{b_{n-1} b_n}$$

We will use this formula to determine the 2-state transition probabilities for $I(\mathcal{M}_{p,d})$. For convenience we omit commas and parantheses when writing subsequences and also abbreviate $\mathcal{M}_{p,d}(i)$ as M_i, $I(\mathcal{M}_{p,d})(i)$ as $I(M)_i$.

$$\begin{aligned}
\mathbb{P}\left(I(M)_1 I(M)_2 I(M)_3 = 111\right) \quad = \quad & \mathbb{P}\left(M_1 M_2 M_3 M_4 = 1111\right) + \\
& \mathbb{P}\left(M_1 M_2 M_3 M_4 = 1110\right)(1-p) + \\
& \mathbb{P}\left(M_0 M_1 M_2 M_3 M_4 = 10111\right) p + \\
& \mathbb{P}\left(M_0 M_1 M_2 M_3 M_4 = 10110\right) p(1-p).
\end{aligned}$$

To see why this is true, observe that there are four disjoint cases in which $I(M)_1 I(M)_2 I(M)_3 = 111$: one, that $M_1 M_2 M_3 M_4 = 1111$; two, that $M_1 M_2 M_3 M_4 = 1110$ and switching does not occur at the pair (M_3, M_4); three, that $M_0 M_1 M_2 M_3 M_4 = 10111$ and switching occurs at the pair (M_0, M_1); and four, that $M_0 M_1 M_2 M_3 M_4 = 10110$ and switching occurs at the pair (M_0, M_1) and not at the pair (M_2, M_3).

Thus, applying the algebraic relations mentioned above,

$$\begin{aligned}
\mathbb{P}\left(I(M)_1 I(M)_2 I(M)_3 = 111\right) &= d_1 m_{11} m_{11} m_{11} + d_1 m_{11} m_{11} m_{10}(1-p) + \\
&\quad d_1 m_{10} m_{01} m_{11} m_{11} p + d_1 m_{10} m_{01} m_{11} m_{10} p(1-p) \\
&= d_1 m_{11}(m_{11} m_{11} + m_{11} m_{10}(1-p) + \\
&\quad m_{10} m_{01} m_{11} p + m_{10} m_{01} m_{10} p(1-p)) \\
&= d_1 m_{11}(m_{11} m_{11} + m_{11} m_{10}(1-p) + \\
&\quad m_{10} m_{01} m_{11} p + m_{11} m_{00} m_{10} p) \\
&= d_0 m_{11}^2 (m_{11} + m_{10}(1-p) + m_{10}(m_{01} + m_{00}) p) \\
&= d_0 m_{11}^2 (m_{11} + m_{10}(1-p) + m_{10} p) \\
&= d_0 m_{00}^2 (m_{00} + m_{01}) = d_0 m_{00}^2 \\
&= \mathbb{P}\left(M_1 M_2 M_3 = 000\right),
\end{aligned}$$

as desired.

Derivations for the other 7 probabilities are similar. □

We have seen that the distribution $\mathcal{M}_{p,d}$ is stable under the traffic flow functional. To prove convergence we will need the following.

Lemma 14.8. *Suppose X and Y are binary processes. Then*

$$\overline{d}\left(I_p(X), I_p(Y)\right) \leq \overline{d}\left(X, Y\right)$$

Proof. Take a joint measure Z such that

$$\overline{d}\left(X, Y\right) = d_Z\left(X, Y\right).$$

Our method is to demonstrate a joint measure $I'_p(Z)$ of $I_p(X)$ and $I_p(Y)$ such that $d_{I'_p(Z)}\left(I_p(X), I_p(Y)\right) \leq d_Z\left(X, Y\right)$. Essentially this involves "binding" together the switching processes $X \mapsto I_p(X)$ and $Y \mapsto I_p(Y)$.

Z is a distribution on pairs of binary sequences (F, G). Represent adjacent terms of these pairs in matrices like so:

$$\begin{pmatrix} F(k) & F(k+1) & F(k+2) \\ G(k) & G(k+1) & G(k+2) \end{pmatrix} = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 1 \end{pmatrix}$$

Note that $d_Z(X, Y)$ is equal to the frequency with which vertical pairs disagree in this matrix representation. We will thus define $I'(Z)$ with an eye toward minimizing the number of such pairs.

Define $I'_p(Z)$ from Z as follows:

$$\begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \mapsto \begin{cases} \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix} & \text{w/ prob. } p \\ \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} & \text{w/ prob. } 1-p \end{cases}$$

$$\begin{pmatrix} 1 & 0 \\ a & b \end{pmatrix} \mapsto \begin{cases} \begin{pmatrix} 0 & 1 \\ * & * \end{pmatrix} & \text{w/ prob. } p \\ \begin{pmatrix} 1 & 0 \\ * & * \end{pmatrix} & \text{w/ prob. } 1-p \end{cases}$$

when $(a, b) \neq (1, 0)$;

$$\begin{pmatrix} a & b \\ 1 & 0 \end{pmatrix} \mapsto \begin{cases} \begin{pmatrix} * & * \\ 0 & 1 \end{pmatrix} & \text{w/ prob. } p \\ \begin{pmatrix} * & * \\ 1 & 0 \end{pmatrix} & \text{w/ prob. } 1-p \end{cases}$$

when $(a, b) \neq (1, 0)$.

As before, all of these choices are made independently.

$I'_p(Z)$ as defined above is a joint measure for $I_p(X)$ and $I_p(Y)$. This is clear if one simply ignores one of the rows: the rules for the top row are equivalent to those defining $I_p(X)$ in terms of X, and likewise for $Y \mapsto I_p(Y)$ on the bottom row. The effect of these rules is that when $(X_i, X_{i+1}) = (Y_i, Y_{i+1}) = (1, 0)$, the switching decisions at (X_i, X_{i+1}) and (Y_i, Y_{i+1}) are coupled together; all other switching decisions are made independently.

The key assertion is that *the density of unalike vertical pairs does not increase with the application of* I'_p. To verify this claim simply involves checking each possible 2×3 binary matrix to see that the expected number of unalike terms remains the same or decreases; we omit this part of the proof.

So applying I'_p does not increase the value:

$$\limsup_{n\to\infty} \mathbb{E}_Z \left(\frac{\sum_{i=-n}^{n} |X_i - Y_i|}{2n+1} \right).$$

Thus,

$$\begin{aligned} \overline{d}\,(I_p(X), I_p(Y)) &\leq d_{I'_p(Z)}\,(X, Y) \leq d_Z\,(I_p(X), I_p(Y)) \\ &= \overline{d}\,(X, Y) \end{aligned}$$

as desired. $\square$

Therefore $\left(\overline{d}\left(I_p^n(X), I_p^n(Y)\right)\right)_n$ is a monotone decreasing sequence for any binary processes X and Y. In particular,

$$\overline{d}\left(I_p^n(X), I_p^n(\mathcal{M}_{p,d})\right) = \overline{d}\left(I_p^n(X), \mathcal{M}_{p,d}\right)$$

is monotone decreasing.

The final lemma is weak but necessary: we need to rule out the case where $\left(\overline{d}\left(I_p^n(X), \mathcal{M}_{p,d}\right)\right)$ is bounded below by a positive value.

Lemma 14.9. *There exists a continuous function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ with $f(x) < x$ such that for any binary processes X and Y with the same density, if*

$$\overline{d}\left(X, Y\right) < \epsilon$$

then

$$\overline{d}\left(I_p^N(X), I_p^N(Y)\right) < f(\epsilon)$$

for some sufficiently large N.

Proof. Omitted. □

Now we can complete the proof of the theorem. Let X be a binary process of density d. Then $\overline{d}\left(I_p^n(X), \mathcal{M}_{p,d}\right)$ is monotone decreasing with respect to n, so it must approach some limit δ. Suppose for the sake of contradiction that $\delta > 0$. $f(\delta) < \delta$, so by continuity there exists $\lambda > 0$ such that $f(\delta + \lambda) < \delta$. Choose n large enough that

$$\overline{d}\left(I_p^n(X), \mathcal{M}_{p,d}\right) < \delta + \lambda.$$

Apply the lemma above to $\epsilon = \delta + \lambda$ to obtain a value N such that

$$\overline{d}\left(I_p^{n+N}(X), \mathcal{M}_{p,d}\right) < f(\delta + \lambda) < \delta.$$

This is a contradiction.

Thus

$$\limsup_{n\to\infty} \overline{d}\left(I_p^n(X), \mathcal{M}_{p,d}\right) = 0,$$

as desired. □

15 Appendix 2

```
// simcell1.cc
// one lane cellular automata simulation

#include <iostream.h>
#include <math.h>
#include <stdlib.h>

bool randprob(float p)
{
  float r=float(rand()+0.5)/32768;
  return (p>r);
}

int main()
{
  bool carmoved=false;
  int k;
  cin >> k;
  float chunks[10000];
  float newchunks[10000];
  int i,t,last;
  for (i=0;i<10000;i++)
    if (rand()<k) chunks[i]=1;
  int car_to_watch=0;
  while (chunks[car_to_watch]==0) car_to_watch++;
  cout << car_to_watch << endl;
  for (t=0;t<5000;t++)
  {
     for (i=9999;i>=0;i--)
     {
        if (i==0) last=9999; else last=i-1;
        if (!(carmoved) && (chunks[i]==0) && (chunks[last]==1) &&
           (randprob(144.0/145)))
        {
            chunks[i]=1;
            chunks[last]=0;
            carmoved=true;
            if (last==car_to_watch) car_to_watch=i;
```

```
        }
        else carmoved=false;
      }
      if (t%100==99) cout << car_to_watch << endl;
  }
}
```

16 Appendix 3

```
// simcell2.cc
// Two lane cellular automata simulation
// includes characteristic velocity differences and random fluctuations

#include <iostream.h>
#include <math.h>
#include <stdlib.h>

#define PI 3.14159265359
#define MAX_LENGTH 1000
#define PRINT_LENGTH 50
#define MAX_TIME 1000
#define OUTPUT false
#define PREVENT_SWITCH false // set to true to prevent switching lanes

float sigma_v; float sigma2;

float phi(float x)
{
   // Abramowitz & Stegun 26.2.19
   float
      d1 = 0.0498673470,
      d2 = 0.0211410061,
      d3 = 0.0032776263,
      d4 = 0.0000380036,
      d5 = 0.0000488906,
      d6 = 0.0000053830;

   float a = fabs(x);
   float t = 1.0 + a*(d1+a*(d2+a*(d3+a*(d4+a*(d5+a*d6)))));
```

```
   // to 16th power
   t *= t;  t *= t;  t *= t;  t *= t;
   t = 1.0 / (t+t);  // the MINUS 16th

   if (x >= 0)  t = 1 - t;
   return t;
}

float phiinv(float p)
{

   // Odeh & Evans. 1974. AS 70. Applied Statistics. 23: 96-97
   float
      p0 = -0.322232431088,
      p1 = -1.0,
      p2 = -0.342242088547,
      p3 = -0.0204231210245,
      p4 = -0.453642210148E-4,
      q0 =  0.0993484626060,
      q1 =  0.588581570495,
      q2 =  0.531103462366,
      q3 =  0.103537752850,
      q4 =  0.38560700634E-2,
      pp, y, xp;

   if (p < 0.5)  pp = p;  else  pp = 1 - p;

   if (pp < 1E-12)
      xp = 99;
   else {
      y = sqrt(log(1/(pp*pp)));
      xp = y + ((((y * p4 + p3) * y + p2) * y + p1) * y + p0) /
               ((((y * q4 + q3) * y + q2) * y + q1) * y + q0);
      }

   if (p < 0.5)  return -xp;
   else  return  xp;
}
```

```
float getrand(float mu,float sigma)
{
  float r=rand()+0.5;
  return mu+phiinv(r/32768)*sigma;
}

bool randprob(float p)
{
  float r=float(rand()+0.5)/32768;
  return (p>r);
}

float getprob()
{
  float mu=getrand(88,sigma_v);
  return (mu/88)/(1+(sigma2/mu)*(sigma2/mu));
}

int main()
{
  int count=0;
  bool * carmoved=new bool[2];
  bool * canmove=new bool[2];
  carmoved[0]=false; canmove[0]=true;
  carmoved[1]=false; canmove[1]=true;
  int k;
  cin >> k;
  cin >> sigma_v;
  cin >> sigma2;
  float chunks[MAX_LENGTH][2];
  int i,j,t,last;
  for (i=0;i<MAX_LENGTH;i++)
    for (j=0;j<2;j++)
       if (rand()<k) chunks[i][j]=getprob(); else chunks[i][j]=0;
  for (t=0;t<MAX_TIME;t++)
  {
     for (i=MAX_LENGTH-1;i>=0;i--)
     {
        if (i==0) last=MAX_LENGTH-1; else last=i-1;
        for (j=0;j<2;j++)
```

```
{
   canmove[j]=((chunks[i][j]==0) && (!carmoved[j]));
   carmoved[j]=false;
}
if ((canmove[0]) && (canmove[1]))
{
    for (j=0;j<2;j++)
    {
        if ((chunks[last][j]!=0) && (randprob(chunks[last][j])))
        {
           if (OUTPUT) cout << "Moving ahead " << last << " " << j << endl;
           chunks[i][j]=chunks[last][j];
           chunks[last][j]=0;
           carmoved[j]=true;
           if (last==MAX_LENGTH-1) count++;
        }
    }
}
else if (canmove[0])
{
    if ((chunks[last][0]!=0) && (randprob(chunks[last][0])))
    {
        if (OUTPUT) cout << "Moving forward " << last << " " << 0 << endl;
        chunks[i][0]=chunks[last][0];
        chunks[last][0]=0;
        carmoved[0]=true;
        if (last==MAX_LENGTH-1) count++;
    }
    else if ((chunks[last][0]==0) && (chunks[last][1]!=0) &&
             (randprob(chunks[last][1])) && (!(PREVENT_SWITCH)))
    {
        if (OUTPUT) cout << "Switching lanes " << last << " " << 1 << endl;
        chunks[i][0]=chunks[last][1];
        chunks[last][1]=0;
        carmoved[1]=true;
        if (last==MAX_LENGTH-1) count++;
    }
}
else if (canmove[1])
{
```

```
            if ((chunks[last][1]!=0) && (randprob(chunks[last][1])))
            {
                if (OUTPUT) cout << "Moving forward " << last << " " << 1 << endl;
                chunks[i][1]=chunks[last][1];
                chunks[last][1]=0;
                carmoved[1]=true;
                if (last==MAX_LENGTH-1) count++;
            }
            else if ((chunks[last][1]==0) && (chunks[last][0]!=0) &&
                     (randprob(chunks[last][0])) && (!(PREVENT_SWITCH)))
            {
                if (OUTPUT) cout << "Switching lanes " << last << " " << 0 << endl;
                chunks[i][1]=chunks[last][0];
                chunks[last][0]=0;
                carmoved[0]=true;
                if (last==MAX_LENGTH-1) count++;
            }
        }
      }
      if (OUTPUT) cout << "TIME " << t << ": count = " << count << endl;
      if (OUTPUT)
      {
         for (i=0;i<PRINT_LENGTH;i++)
         if (chunks[(0+i)%MAX_LENGTH][0]>0)
            cout << 1;
            else cout << 0;
         cout << endl;
         for (i=0;i<PRINT_LENGTH;i++)
         if (chunks[(0+i)%MAX_LENGTH][1]>0)
            cout << 1;
            else cout << 0;
         cout << endl << endl;
      }
  }
  cout << count << endl;
}
```