MAT3395

Introduction to Mathematical Models and Mathematical Software

Final Project: Traffic Problems On The Quebec Bridges

Marc-André Théberge [redacted]

Alexandre René [redacted]

Tyler Clemens [redacted]

Alpar Sadr Al Ashrafi [redacted]

Table of contents:

1. Introduction

2. The Model

3. Statistics and Data

4. Neglected Elements

5. Results

6. Conclusion

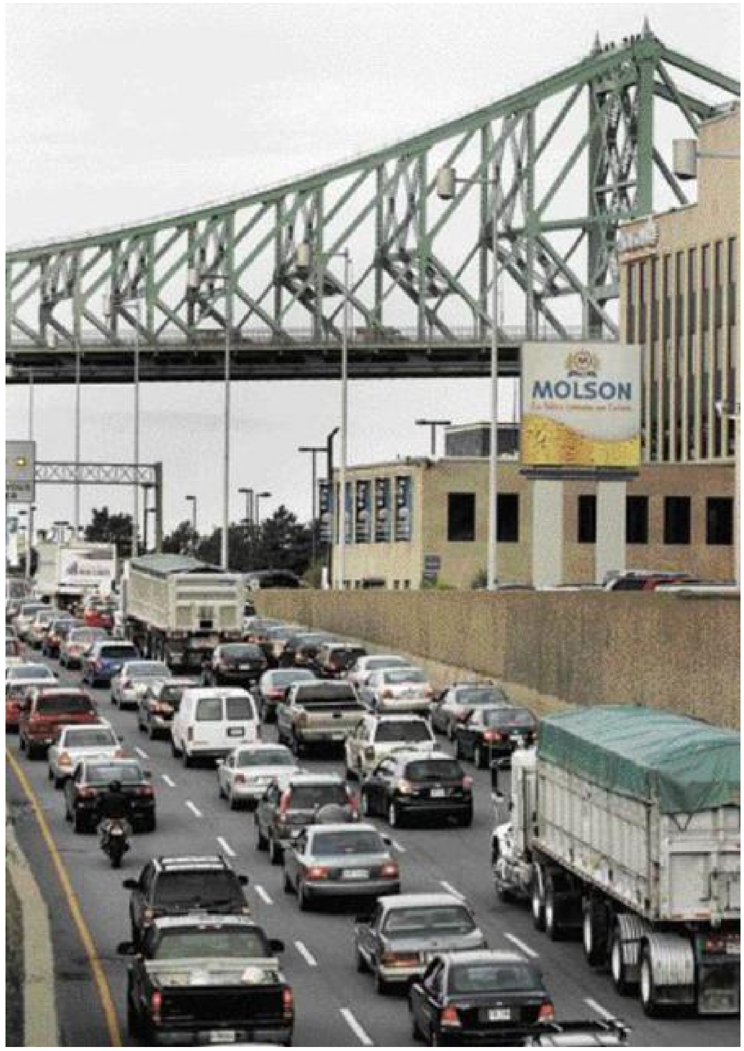

Introduction

The traffic on the bridges has been a big problem for residents for a long time. Some of us experience it everyday, others hear it on the radio. The traffic has been considered heavy every morning for years. Our project consists of reproducing the traffic flow coming into Ottawa from Gatineau-Hull during the morning rush hour. We are mainly concerned about the traffic flow on the 5 bridges leading into the Ottawa region. By the end of the project, we hope to be able to demonstrate different situations to see how they will influence this problem. In particular, we would like to give a prediction of the benefits arising from adding the planned bridge on Kettle Island.

It is worth mentioning that a Canadian company known as Castleglenn Consultants Inc. does this kind of work all the time for cities around Canada. They specialize in transportation planning and traffic engineering and design services in the fields of functional planning, traffic engineering, highway-roadway design and project management throughout communities in Alberta and Ontario. You must have seen the people sitting at busy corners looking bored and counting cars. This is how they collect their data, and then they create models for different cities. Unfortunately, their data or models are not open to the public and we were not able to compare them with our model.

Constructing our model

Our first model consisted of a straight highway, with a constant number of vehicles heading in the same direction at constant speed. Using this model allowed us to verify our equations to determine if they had to be altered (which indeed turned out to be the case).

The first thing we need to determine is the speed of cars on a street as a function of the population of cars P on it. From simple geometry, we know the the average space per car, E, is

$$\text{Average space per car} = E = \frac{DM}{P}$$

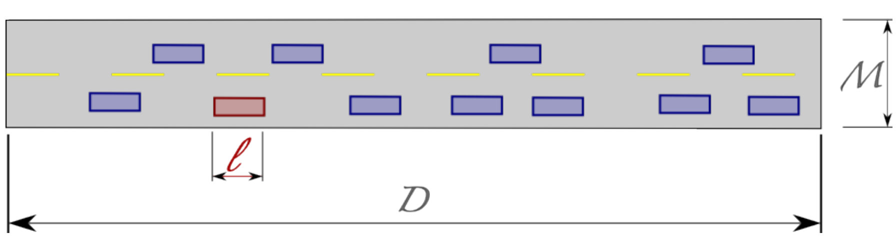

Diagram 1: Geometric argument for determining the average space per car.

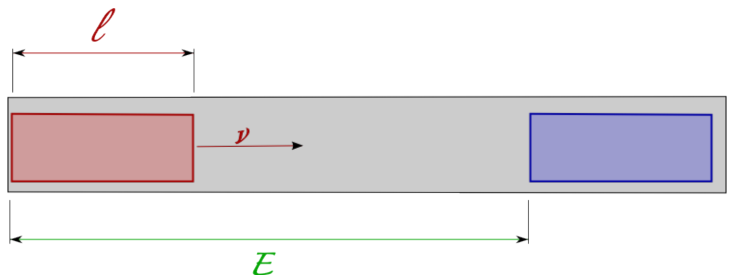

Diagram 2: Applying the two second rule to determine car speed.

Then, if we assume that on average people maintain a speed such that the space in front of them is equivalent to the distance they cover in 2 seconds, their speed V must satisfy

$$\text{V} * 2 \text{ seconds} = E - l = \frac{DM}{P} - l$$

Which gives us the average speed per car,

$$\text{Average car speed on street } i = V_i = Min\left\{\left(\frac{D_i M_i}{l} - P_i\right) \times \frac{3.6}{2},\quad V_i^{max}\right\} \tag{1}$$

where:

D_i is the length of the street (here a "street" is a segment joining two consecutive intersections).

M_i is a multiplying factor, mostly related to the number of lanes.

P_i is the number on vehicles on the street.

l is the average length of a vehicle (in metres).

$V_i^{\max}$ is the speed limit on street i.

3.6 is the conversion factor m/s -> km/h.

We divide by 2 because we use the two-second-rule to determine how fast the vehicles can go without being too close to each other.

Of course, a model composed of one street is of very limited interest. So the next step was to add intersections. Intersections are where cars can enter a street, increasing it's population. For illustration purposes, let's suppose we have an intersection between streets i and j. and that in the direction of traffic flow, street i is after street j. We want to describe the change in population on street i. For this define a parameter Φ_j, the *maximum output* of street j. It is defined as the maximum number of cars that can exit the street in ideal conditions in one minute. Φ is mostly determined by the length of green lights and the intersection's geometry (if there a lane is added for people turning right or left). In cases where there there are more than one streets to which to exit, we assume that Φ is the same in all cases. This of course can be quite far from the truth (such as when comparing the maximum output for cars turning left or continuing) but not doing so would add undue complexity to our model. We further assume that the critical factor in allowing cars to leave a street is the speed of the cars on the street they are moving to and that all other factors are negligible. This seems reasonable, as should correctly model the case where cars can either turn onto a traffic-jammed road or continue to a relatively empty one. Finally, since Φ_j is determined considering the maximum car speed, that also appears in the equation. So the increase in population on street *i* can be written as

$$\frac{dP_i}{dt} = Min\left\{\Phi_j \times \frac{V_i}{V_i^{max}}, P_j\right\} \tag{2}$$

The minimum is only there to avoid adding more cars than the total car population on street *j*.

If many streets *j* are entering street *i*, then we must sum over these streets. Of course, a certain fraction of the cars on each of these entering streets will continue to the street we are studying. This fraction may or may not be 1 for one or many of the streets, depending on the intersection. We write *R*(*j*, *i*) for the fraction of cars going from street *j* to street *i*. We assume this fraction to be independent of time over our period of study (i.e. rush hour). Then equation (2) becomes

$$\frac{dP_i}{dt} = \sum_j R(j,i) \times Min\left\{\Phi_j \times \frac{V_i}{V_i^{max}}, P_j\right\} \tag{3}$$

Of course, we must also consider cars leaving street *i*. The equation for exiting cars is constructed following the same reasoning. The result is very similar to (3), although care must be taken not to misplace the indices. Putting the two together we get the general equation for population variation on any given street:

$$\frac{dP_i}{dt} = \left[\sum_j R(j,i) \times Min\left\{\Phi_j \times \frac{V_i}{V_i^{max}}, P_j\right\}\right] - \left[\sum_k R(i,k) \times Min\left\{\Phi_i \times \frac{V_k}{V_k^{max}}, P_i\right\}\right] \tag{4}$$

where the *j*-streets are entering street *i* and it is exiting to the *k*-streets.

For implementation, we replace the sum over exiting streets by a sparse $n \times n$ matrix containing the *R*(*j*, *i*) values, where *n* is the number of streets in our system. Then the corresponding matrix for the entering streets is simply the transpose of the first. In this way our two sums become two matrix multiplications, which is easier to handle programmatically.

In this way we can describe the rate of change of population on each street by a differential equation. Taken together, the set of equations describes how cars move around in the system. However, we still have no way of getting them in or out of the system. We will now address that.

Defining entry points

Two of us know from personal experience that the latest one can leave while still hoping to be in Ottawa by 8:00 a.m. is between 7:15 and 7:30. As more people leave before that than after, we will take 7:15 a.m. as the time where the amount of people leaving their home peaks.

In this case, since we do not know how these people are distributed, it seems natural to assume that they follow a normal distribution. We further assume that people going to work will leave between 5:00 a.m. and 10:00 a.m., so we want to chose a standard deviation σ_N that will place 90% of them within this time frame. If we scale time in minutes, we find that the appropriate values are $\mu_N = 135$ and $\sigma_N = 62$. That yields the following curve, which represents the rate $N(t)$ at which people leave their homes for work at any given time t.

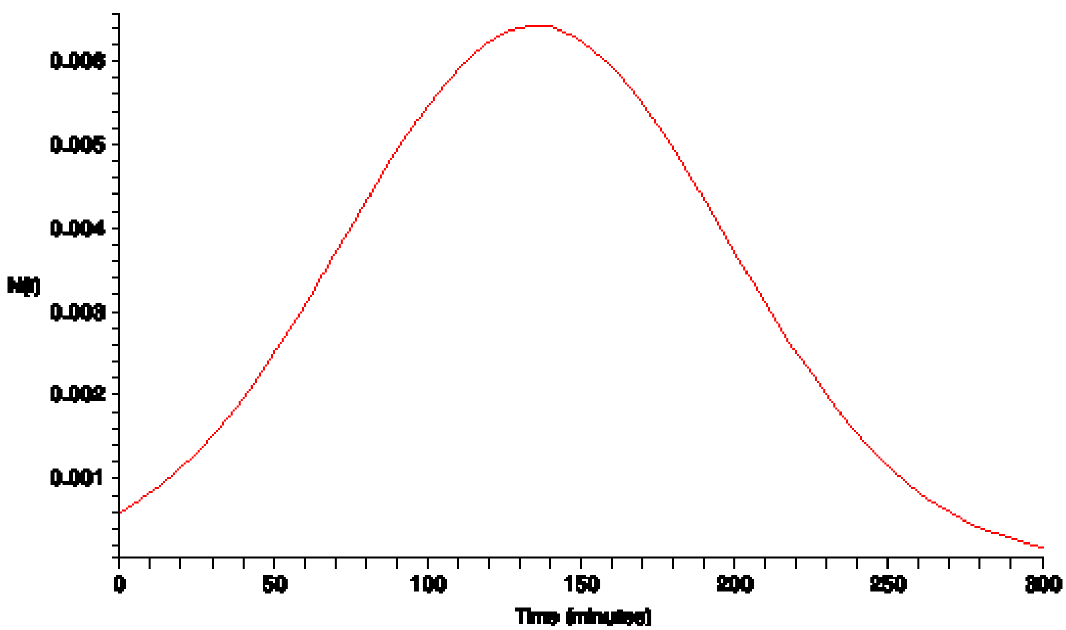

Figure 1: Distribution of the system entry time

Therefore, to add cars into the system, all we do is define a certain number of *entry points*. These are streets with an extra property: the total number of people who work downtown and who will pass through them, which we will note P. Then at any given time t, the rate at which people are being added to this street is simply *P N*(*t*). In this way, incoming people are treated simply by adding a positive term to the differential equations of the entry points.

An important advantage of our definition of entry point is that makes it easy to represent the incoming population of one or more sectors with only one street, helping make the model workable.

Defining exit points

In it's essence, an *exit point* is very similar to a entry point: it will be treated by adding a negative term to it's differential equation. However, it is not quite clear what this term should look like. The first simulations were run using just a constant term. As might be expected, this led to bridges that were much too easy to empty and virtually no traffic jam in the key areas.

We wanted to base the exit rate on data the model was generating to avoid making more or less arbitrary assumptions about the traffic on the Ottawa side. For this reason an attempt was made to use the speed on the bridge as a factor to determine it's exit rate. However this turned out to be no better then the constant term, as the positive feedback between the speed and the exit rate brought them to a complete stop from which they could not recover.

Lacking better options, we finally did make a more or less arbitrary assumption on the density of traffic on the Ottawa side. We figured it would reach it's peak at 8:00 a.m. and should vary according to something resembling a normal distribution *T*(*t*). We settled on a standard deviation of $\sigma_T = 40$, which concentrates most of the traffic between 6:50 and 9:10. The exit rate from the bridge should be smallest when traffic density is greatest. However, even at that point it should not be zero; we fixed it's minimum at 5% of the bridges maximum output Φ, which we defined earlier. This leads us to define the exit rate as follows:

$$\frac{dP}{dt} = -\Phi \frac{1.05\ (\sup\ T(t)) - T(t)}{\sup\ T(t)}$$

or equivalently,

$$\frac{dP}{dt} = -\Phi(1.05 - T(t)\,\sigma_T \sqrt{2\pi})$$

Since this makes the maximum exit rate 1.05 Φ instead of 1 Φ, we simply scale back the maximum output for the bridges by 5%. Otherwise this suits our needs quite well.

Statistics and Data

With our model defined, we need to shape it to the situation at hand, namely the Gatineau road system. The first step is to identify which roads are essential to modeling the morning rush hour traffic. These are shown on Figure 2, overlayed on top of a map of the area. We purposely did not integrate entry points directly into the network. Otherwise, because their extra term does not consider the car speed on the street, too many cars could appear on an already packed street. Some of these entry points actually represent a number of residential streets between two considered intersections.

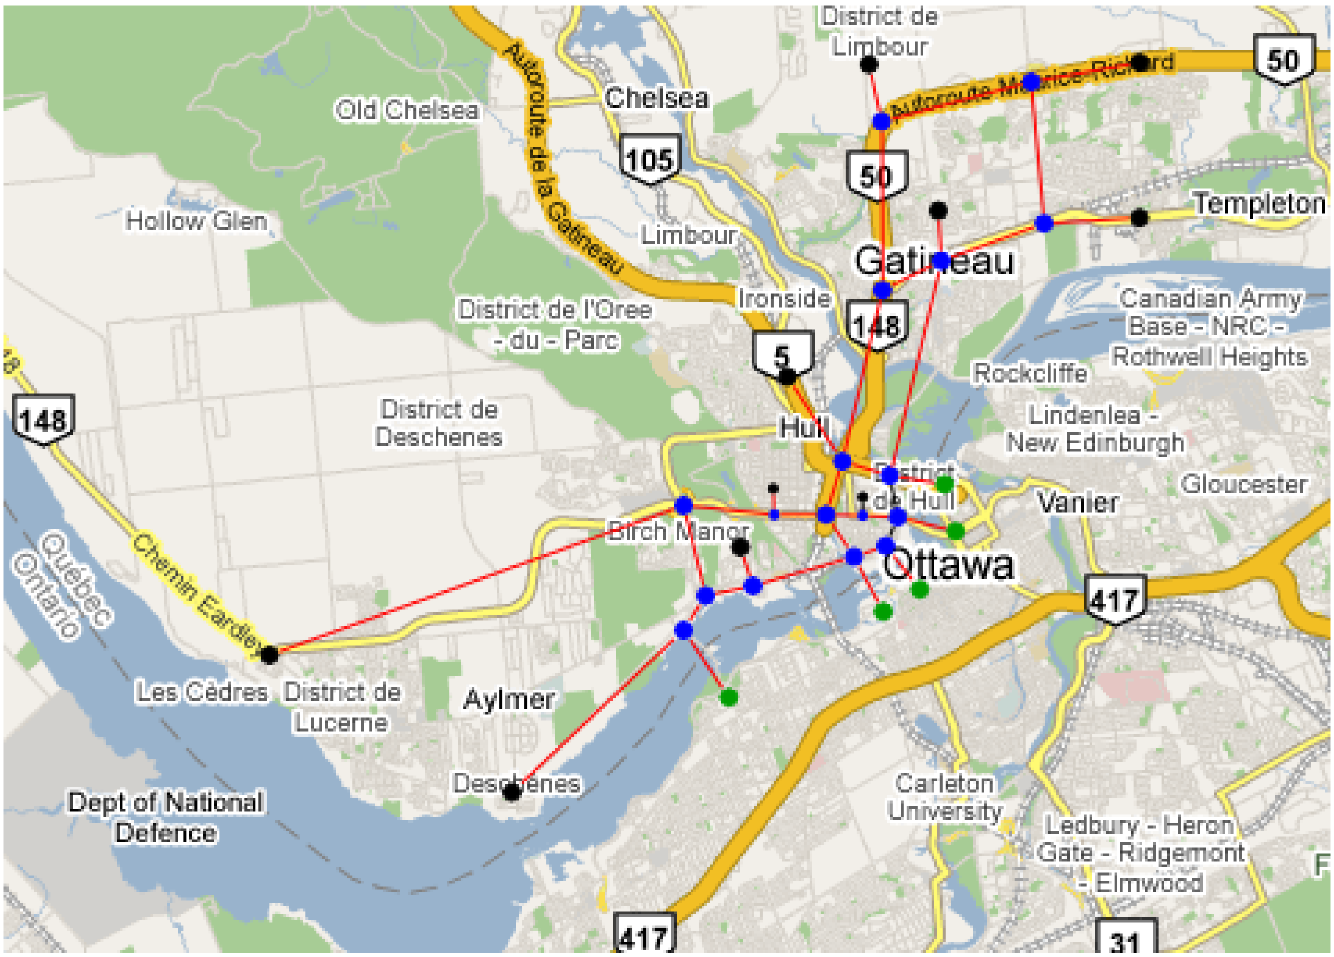

Figure 2: Basic map of Gatineau we used. The road system was divided into 42 streets for internal representation, drawn in red. Blue points are standard intersections. Cars enter from the black points, and exit to the green ones.

Now we need to define how many people will enter from each of our entry points. We know that the active population of Gatineau is 176 100 people. However, a sizable portion of these do not require to cross a bridge to reach their workplace, and thus would not influence the traffic we are studying. We have no data on the number of people who are in this situation, but we know that there are too many to simply ignore them. Therefore we fixed the proportion of active people that work downtown at 45%. Some of these work in downtown Hull, contributing up to there but not onto the bridges. For simplicity, and not having actual data to estimate their number, we did not account for them and supposed every considered worker crosses a bridge to Ottawa. That puts at 80 000 the number of Gatineau workers who will cross a bridge at rush hour. There are also some coming from two smaller municipalities; when we apply the same proportions we did for the Gatineau workers, these add another 4000 to the system.

Now, to determine how many of these 84 000 workers will come from each entry point, we used the following population distribution given in Figure 3. We then had to further divide those workers among our identified entry points; this was done according to our knowledge of the area.

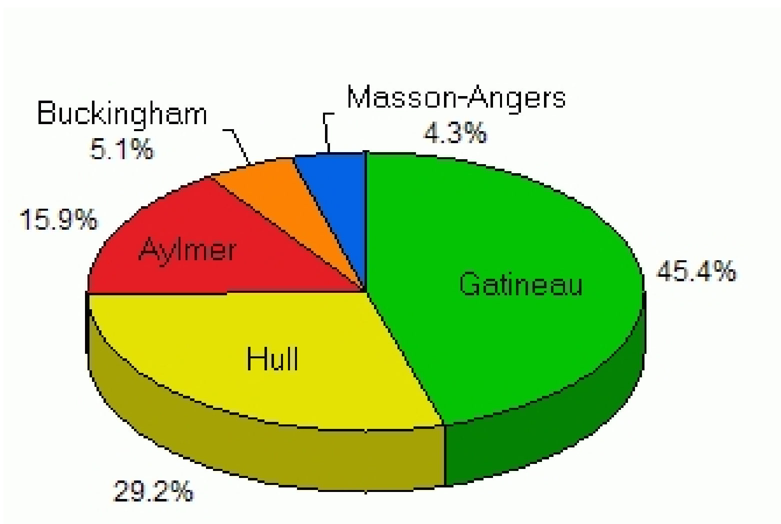

Figure 3: Population distribution among the Gatineau sectors

Finally, we need to determine the properties of each street which used in constructing our model. This includes: street length (*D*) and width (*M*), speed limit (V^{max}), maximum output (Φ) and, for every other street, the proportion of cars that will be moving to it. The first three components can be determined with satisfactory precision. So can the proportion of cars moving into each other street, in the majority of cases (for many of them, the is only one subsequent street). That leaves the maximum output as the only other ingredient (along with the number of downtown workers) in our formula with a high guesswork content.

Results

Assumptions and neglected elements

Before discussing our results, we must point out some of the aspects that might affect them. For example, the assumptions upon which are built the entry point and exit point equations are potentially important sources of error. We have also already mentioned the guesswork involved in determining how many people contribute to traffic and the maximum output of the streets. These last two factors have the greatest potential to modify results. Notably, there are 42 output factors to determine, which puts a lot of variability in our results.

The other principal assumptions/neglected elements are :

- Every vehicle was assumed the same length.
- Although a very approximate factor is added for Ottawa traffic, our model essentially stops after the vehicles cross their respective bridge, and we consider them to have arrived at their destination and their influence on the model negligible. We realize that the cars continue to move in traffic into the city, and that this will affect traffic flow on the bridge, but we do have to stop at some point.
- We have neglected reserved lanes, as in taxi and bus lanes. We assumed that every lane on the studied street are used for common traffic.
- We have studied the traffic going into Ottawa and not coming out. That being said, especially during rush hour, this should not be particularly significant.
- As mentioned before, we have neglected the population that work in the Gatineau-Hull area, as we assumed that the majority of the population work in the Ottawa region.
- We have also neglected the Alonzo bridge as a route to get into Ottawa.
- We did not consider random events and/or acts of God.
- We do not consider the time it takes to drive a street's length, because during rush hour you don't really have a chance to roll your normal speed. For this reason, however, this model would be no good at modeling light traffic. Then again, that hardly requires an elaborate model, as the only things you really need to consider are distance and speed limit.

Results and Analysis

Although we were worried at first that MatLab might balk at having to solve 42 simultaneous differential equations, it turns out this is not a problem at all. This is probably thanks to the fact that apart from the entry and exit points, each equation is a linear. Therefore scalability, should we ever consider it, should not be a problem.

Our first objective was to tweak our system parameters to obtain something resembling actual traffic. This is important, because otherwise there is not really any use in trying to get a prediction out of the model. From our first simulations, the main difficulty was evident: traffic seemed to jam at the entry points, covering the rest of the distance to Ottawa with relative ease. Instead, what we should see is traffic begin to slow downtown, especially on the bridges, and then the effect should cascade up through the road network, reaching the entry points last. The two most obvious ways this might happen is if the numbers of workers is overestimated and if the maximum output of the key streets is underestimated. Incidentally, those are the two factors which were the least precisely determined.

However, “tweaking” wasn't a solution, as it turns out that our model is extremely robust with regard to parameter changes. Which in a sense is not entirely bad, as it means that whether your neighbour leaves at 6:50 or 7:20 won't really affect traffic as a whole. But in our case, removing a neighbourhood did not seem to make much more of a difference than the neighbour leaving late. (Unfortunately, we lacked the foresight to record one of these earlier versions, so you cannot see for yourself how little things change.) We cut populations down by ¼, then ⅓ for the problematic sectors (for a total of ½). We also multiplied some maximum output by factors of between 2 and 4, some entry points reaching 135 cars per minute. This is pretty much the limit of what we consider to be believable numbers for these parameters, and yet some streets still contain cars until 1 p.m. (down from about 5 p.m. originally). So rather than continue making up numbers, we will concentrate at what else there is to say about and from the model in it's current form.

Analysing our data poses a slight problem because there are 42 different street populations to follow. We can plot a subset of these populations, as we did to monitor the populations on the entry points (Figure 4 is such a plot). To see the whole picture, we built a visual representation of our map that shows the car concentrations change through time.

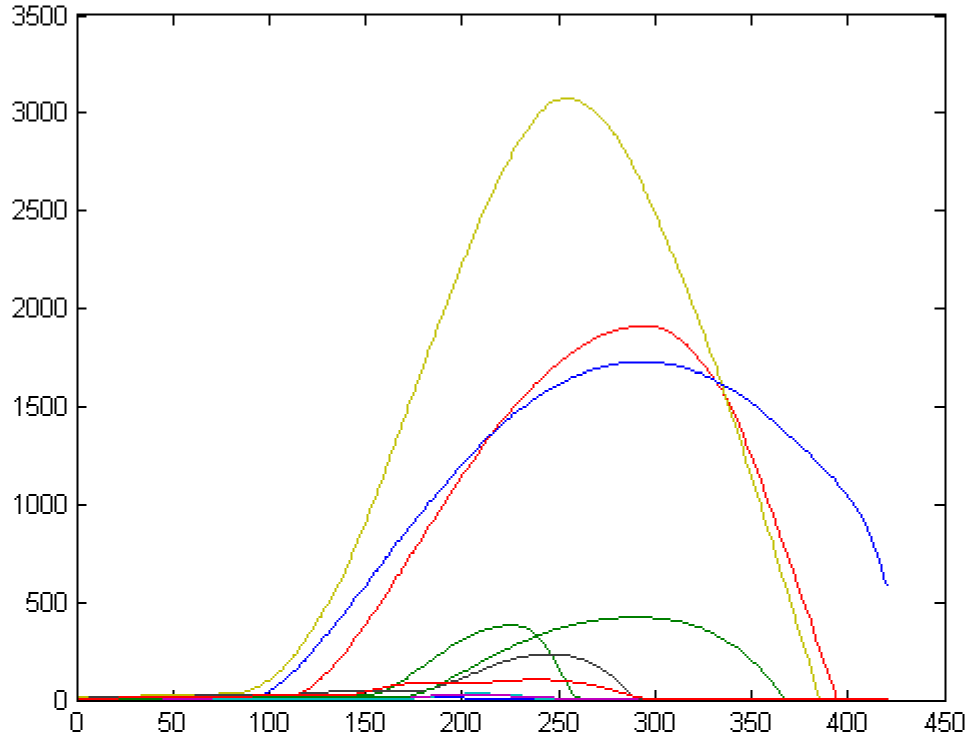

Figure 4: Car population on the entry points as a function of time. Here σ_N = 50. Time is in minutes

We recorded two such simulations using our lightly improved parameters for your own viewing pleasure. The closer a street's colour is to red, the denser the cars are packed. Time is indicated at the bottom, in hours:minutes:seconds. These two simulations were produced to compare the effect of changing the value of sigma from 50 to 90 for the entry points' equations, which turned out to be very small. We produced simulations to compare the other factors, but like the plots, it really is just more of the same.

One thing that is important to keep in mind when looking at these simulations is that the short entry streets do not have any meaningful length, since they represent the streets of entire residential sectors. As such, the density of cars on them can quickly become very high without it being a cause for concern (as noted previously though, the fact that this density is not zero at noon *is* problematic).

Now if we look at what happens around the bridges, in downtown Hull, we notice that traffic starts to form in areas where two major arteries meet. This corresponds to the actual traffic phenomenon, so at least the structure of the model seems sound. Playing the simulations slowly or frame by frame between 6 and 7 a.m. really shows cascading affect of traffic, as intended. However, there really isn't any congestion in downtown Hull before the Ottawa traffic kicks in at around 7:30. And as soon as it's over, downtown Hull begins to empty itself, followed by the rest. The fact that downtown empties itself before the incoming streets is unquestionable proof that the congestion around the bridges as it appears in our model is almost entirely due to the simulated Ottawa traffic, and not the one produced by our model. This tells us two important things :

- That traffic on the Quebec side is highly dependent upon traffic on the Ottawa side.
- Consequently, if adding a bridge on Kettle island reduces Quebec side traffic, it would be a side effect of reducing traffic in Ottawa.

The second observation is a disappointing, because it means that we cannot attain our main objective, which was to evaluate the effect of adding a bridge on Kettle Island.

There is another problem we realized in the last couple days that also prevents us from simply adding the Kettle Island bridge to the model to see what happens. The way we constructed the model, for every intersection we predetermine what proportion of the cars go in each direction. This is great because it completely avoids the problem of determining routes for each car. It works because we never care about *which* car is on a certain street but only *how many*. However, this also means that any time we want to close, reduce or add a street, we have to manually reroute the traffic. This pretty much renders the model useless for such things as modelling the effects of construction or accidents on traffic, since it would more or less regurgitate the new routes as determined by the human operator.

Conclusion

In conclusion, what can be said from our results ? With it's shortcomings, our model is not yet a viable candidate to answer specific questions about traffic flow in our region. However, it does exhibit characteristic traffic behaviour, even it isn't the one we want. Therefore, it is still a reasonable basis for more general conclusions. We feel that our model still makes a good case of demonstrating two such conclusions. First, that traffic flow is a very stable phenomenon. Whether our model amplified this or not, it still demonstrates how difficult it can be to influence traffic by changing parameters such as the maximum output. Although we did not attempt to model cases where streets were added or removed, but we suspect this has a much greater effect than the parameters we did play with. This would, among other things, demonstrate the value of reducing traffic by adding roads or lanes.

The second conclusion we reached was that modelling at least part of Ottawa's traffic is absolutely essential to getting an accurate model for Gatineau's traffic. Of course, this isn't exactly groundbreaking news, but we initially believed we could get a something at least resembling the true traffic by considering only cars in Gatineau. Our results show that this is not the case. In a sense, our model is just good enough to show that it isn't good enough.

We already discussed many of the shortcomings of our model, but this does not mean it's potential is limited to the above two conclusions. With some precise data on certain key points where we were unable to reproduce real traffic, we could probably correct much of the inconsistencies between reality and our results. Then we could write an algorithm to determine the routes different groups of cars would follow and from that obtain the $R(j,i)$ ratios used to calculate traffic. That would allow us to model the effects of modifying the road network. Finally, as noted, we would also need to consider at the least Ottawa's downtown traffic. This is nothing more than translating a map of Ottawa into data the model can manipulate. However, Ottawa's road network is much more intricate than Gatineau's, for which manually entering the data was quite error prone due to the sheer amount of it. Because of this, adding Ottawa's road network would also require cooking up some sort of interface to avoid manually everything into gigantic matrices of numbers. Of course, doing this would require much more time then we have for this time of project, but it does show that the work we have done is the first step towards a working, complete model.

As such, although we did not reach our objective, we nevertheless consider our model to be valuable in it's own right.