RECURRENT AND NON-RECURRENT CONGESTION ESTIMATION USING TRAFFIC SIMULATIONS

A. de Palma[a], F. Marchal[b]
[a]The Institut Universitaire de France. THEMA, Université de Cergy-Pontoise,
33 bd. du Port, F-95011 CERGY-PONTOISE, France.
Phone: +33-1-34-25-61-81, Fax: +33-1-34-25-62-33,
[b]CoLab, Swiss Federal Institute of Technology
ETH HRS Zentrum, Hirschengraben 84, CH-8092 ZURICH, Switzerland.
Phone: +41-1-632-56-79
e-mail: depalma@eco.u-cergy.fr; fmarchal@inf.ethz.ch

Abstract

Non-recurrent congestion in transportation networks occur because of the stochasticity of the demand and of the supply. Intelligent Transportation Systems such as Advanced Traveler Information Systems (ATIS) and Advanced Traffic Management Systems (ATMS) intend to reduce non-recurrent congestion by providing information to the users or by controlling the unexpected traffic flows. For these reasons, the design of ATIS and ATMS requires to be able to forecast non-recurrent congestion. This paper proposes to measure non-recurrent congestion using dynamic traffic simulations performed with METROPOLIS. Stochastic events are generated that reduce the capacity of the network randomly. Users can adapt to the unexpected conditions from one day to another. Travel time increases are measured at the metropolitan level using large-scale realistic simulations. Assuming that most users are risk averse, any uncertainty corresponds to a utility loss. This utility loss is computed from the simulation results for several level of occurrences of incidents.

Keywords: Traffic simulation; Risk-aversion; Uncertainty; Variable supply; ITS
Topic Area: D3 Integrated Supply / Demand Modelling

1. Introduction

In real transportation networks, the travel time for a given journey often exhibits variability due to various sources of stochasticity. Indeed, the travel demand may vary from one day to the next because of external events such as sport events but also because the efficiency of the infrastructure undergoes unpredictable changes (e.g. weather condition and accidents). A substantial amount of the congestion observed in transportation services is believed to be due to those unexpected events. This is often referred to as non-recurrent congestion. A distinction is made between recurrent and non-recurrent congestion: recurrent congestion occurs on a regular basis and is caused by the fact that a given infrastructure has a lower capacity than the demand it needs to accommodate during peak periods. Before a peak period, no recurrent congestion occurs. As soon as the incoming flow exceeds the capacity of the infrastructure, recurrent congestion builds up. After the peak period, queues discharge and recurrent congestion disappears. Non-recurrent congestion is different since it can occur haphazardly at any moment of the day. From the point of view of the traveler, travel times are thus uncertain. The consequence of non-recurrent congestion is that the user does not know which route to select if he has not received relevant information. This first impact of a potential piece of information is the fact that an individual realizes *ex post* that he did not select the fastest route (which is not to say that the total cost is necessarily larger, but just that some users may unilaterally reduce their travel costs). The second impact of information is that it reduces the uncertainty. The fact that for example a user knows that he will arrive later than expected implies that he can, for example, re-plan his appointments or call his wife (or husband) to postpone the meal (and therefore hope to eat a warm dish). This reduction of uncertainty is widely described in the literature, but very little work has been done in order to measure how much this can be valued by the users. We propose to evaluate how much users value the certainty of their travel alternatives. This evaluation is crucial to the introduction of Intelligent Transportation Systems (ITS). Indeed, Advanced Traveler Information Systems (ATIS) and Advanced Traffic Management Systems (ATMS) are designed to reduce non-recurrent congestion by different means such as the provision of traffic information and route guidance. ATIS and ATMS cannot reduce recurrent congestion except for occasional drivers that ignore the network properties. Therefore, the evaluation *ex ante* of the amount of non-recurrent congestion of a given system provides an upper bound for the gain that can be expected from the real implementation of an ATIS or ATMS.

It should be noted that the frontier between recurrent and non-recurrent congestion is blurred: if some incidents are well-known to occur too frequently at a given location, users may well anticipate them. Users react to recurrent congestion by departing at a time that is compatible with their own schedule constraints and their willingness to incur congestion delays. Typically, this implies a trade-off between facing long travel times and arriving close to the desired arrival time. Users react to non-recurrent delays by altering their pre-trip decisions (such as mode shift, departure shift and path selection) as well as their en-route decision (i.e. route diversion, parking choice) by selecting travel alternatives that are compatible with the level of risk that they are willing to bear and based on the expectations they have from the transportation system. Obviously, a broad panel of reactions are possible, depending on the risk aversion of users and on their knowledge of the network. The microscopic agent-based simulation approach is an ideal tool to capture the heterogeneity of travelers.

The remainder of this paper details how o measure the impact of incidents on the level of nonrecurring congestion and the corresponding loss of utility for travelers. The impacts on travel time delays are measured by introducing random incidents in a microscopic traffic simulation of the peak period. The corresponding loss of utility is measured by combining the results from an empirical survey with the results of the simulation.

2. Measurement of utility loss

The most widely used model to capture risk behavior in Economics is that of expected utility theory. It relies on two separate assumptions. Firstly, it assumes that preferences can be described by a utility function that is known to the modeler. Secondly, it assumes that the attitude toward risk can be rationalized by an expected utility function. This latter function depends on a parameter ϑ called the risk aversion. A survey was administered by one of the author to estimate the level of risk aversion of drivers in Ile-de-France. Respondents are asked lottery-type questions where they have to rank different lotteries assigned to different level of risk. The method has been used also to compute the risk aversion of private investors (see the web site `www.RiskDynaMetrics.com`). Of course in the transportation area, the users are asked to compare route with different variability of travel times, while in the finance application, respondents are asked to compare different financial products which differ according to their level of risk and return. In both case the alternative are naturally ranked and we can use an ordered discrete choice model to estimate the level of risk aversion of the users and to determine the different factors and socio-economic characteristics, which influence the level of risk aversion. We found out, for example, that men are less risk averse than women, and that blue collars are more risk averse than white collars. Basically, the socio-economic characteristics and the purpose of the trip do influence the level of risk aversion. The major outcome of the survey (see Figure 7 of [4]) is the provision of the distribution of risk aversion in the population of drivers. In this paper, travelers are assumed to be mostly risk averse (i.e. to have a positive degree of risk aversion). Therefore, the variability of the driving conditions corresponds to a loss of utility. Different expected utility functions have been proposed in the literature. We select the Constant Absolute Risk Aversion (CARA) utility function. The utility of an journey with travel time τ is given by:

$$U_\theta(\tau) = \frac{1 - e^{\theta\tau}}{\theta}$$

where θ denotes the index of absolute risk aversion , which is constant in the case of the CARA utility. The limit case $\theta \to 0$ corresponds to the risk neutral individual with $U(\tau) = -\tau$. Travel times are assumed to be stochastic during some time period *T*, that can span a whole morning peak period or several days. That variability is described by a probability distribution $f(\tau(t))$. τ_1 According to the expected utility theory developed by von Neumann et Morgernstern, the expected utility of an individual who uses the system at time *t* is given by:

$$E\{U_\theta(t)\} = \frac{1}{\theta} - \frac{1}{\theta}\int_0^T e^{\theta\tau(u)} f(\tau(u))\, du$$

while the utility of the expected outcome is given by:

$$U_\theta(\langle\tau\rangle) = \frac{1 - e^{\theta\langle\tau\rangle}}{\theta} \chi$$

where $\langle\tau\rangle\chi$ is the expected travel time:

$$\langle\tau\rangle = E\{\tau(t)\} = \int_0^T \tau(u) f(\tau(u))\, du$$

For instance, if we have two evenly likely outcomes τ_1 and τ_2, we have

$E\{U_\theta\} = \frac{\chi}{\theta} - \frac{1}{2\theta}\left[e^{\theta\tau_1}{}^{\chi} + e^{\theta\tau_2}{}^{\chi}\right]$ while the utility of the risk-free choice is $\frac{1}{\theta}\left(1 - e^{\frac{\theta}{2}\overset{\chi}{(\tau_1+\tau_2)}}\right)$.

The time compensation, denoted by stochastic system described by $f(\tau(t))$ or a risk-free system with an average travel time $+\langle\tau\rangle$.

Therefore,

$$U_\theta\left(\langle\tau\rangle + \chi\right) = E\left\{U_\theta\left(t\right)\right\}$$

$$\int_0^T e^{\theta\tau(u)} f\left(\tau\left(u\right)\right) du = e^{\theta(\tau+\chi)}$$

$$\chi = \frac{1}{\theta}\ln\left\{\int_0^T e^{\theta\tau(u)} f\left(\tau\left(u\right)\right) du\right\} - \langle\tau\rangle \tag{1}$$

So far we have considered that the utility was only dependent on the travel time. In practice it is often desirable to use a travel cost specification C(t) that depends on other components: boarding costs, egress and access costs, schedule delay costs, etc. In that case, the *monetarization* of the utility loss can be computed using eq. (1) and has then the dimension of money. A meaningful parameter to measure the relative monetary impact is $\hat{\chi} = \frac{\chi}{\langle C\rangle}$ which has no units.

3. Simulation of non-recurrent congestion

So far, nothing has been assumed concerning the probability distribution of travel times (or travel costs) $f(\tau(t))$. We propose to compute $f(\tau(t))$ by performing explicit Monte-Carlo simulations in which the stochasticity is caused by random incidents. The incidents disrupt the capacities of some road sections. The corresponding travel time delays are computed using the simulation tool METROPOLIS.

3.1. Information-based traffic simulation

METROPOLIS ([10, 12, 11]) is a simulation tool that has been developed by the authors. It is intended to be a fully dynamic model that features within-day as well as day-to-day traffic dynamics. We recall here some of its properties that are relevant to our current study. Traffic models have usually two main components: (a) a supply model that describes how the traffic conditions evolve in the road network given the users' driving choices, sometimes called DTA (Dynamic Traffic Assignment) and (b) a demand model that describes the users' behavior given their driving environment and other drivers' decisions. The architecture of METROPOLIS considers information as a third component of transportation models (see Fig. 1). Information means, in the broad sense, any piece of knowledge that can play a role in the users' travel decisions: experienced travel times, shortest routes, congestion levels, radio messages, road guidance advice, etc. Information is user specific since it corresponds to the users' perception of traffic conditions. In the case of the introduction of ATIS or even radio broadcasts, we can consider that information has a physical layer and is a separate entity. For the supply side, we assume that dynamic congestion laws provide some aggregate data on the traffic conditions. All this data about the network state is embedded in the *information* block. Some external information can also take place in that box, like information provided by ATIS devices, radio broadcast or other technologies. In turn, users are assumed to be disaggregated and make use of the available information (which can be user specific like in the case of ATIS equipment) to perform their travel decisions. The supply consists of the coded car network: zones, intersections and links. The congestion is located on links and modeled by dynamic congested laws. Both supply and demand data interact in the simulator that computes sequentially the mode choice, the departure time choice and eventually, the route choice for each simulated user. Destination choice is not modeled. We distinguish three types of information data:

- *historical* or *pre-trip* information is the users' perception of the traffic congestion from previous days,
- *instantaneous* information is acquired en-route and interacts in the route choice: users perform direction choice at each intersection and can be diverted from their original paths due to changes in the information structure,
- *external* information consists of any information provided to the users exogenously by any devices or technology (ATIS, variable messages signs, radio broadcasts, etc.).

The traffic assignment procedure uses a mesoscopic event-based approach. As the traffic simulator proceeds, the *within-day* time evolves. Supply outputs are collected as dynamic level of services: time-dependent travel times and time-dependent traffic flow patterns on each road sections. On the demand side, users choices are collected: departure time, route and mode taken. These results may be aggregated to compute Measures Of Effectiveness (MOEs) or performance indexes. When the simulation of the considered period (e.g. morning peak) is over, the learning process uses these results to update the historical information of previous days. The overall process consists in looping sequentially over the traffic simulation procedure and the learning process. This corresponds to the *day-to-day* dynamics of the simulation. Consequently, the framework model *within-day* dynamics, mainly with the departure time model and the DTA process as well as *day-to-day* dynamics captured by the learning process. If everything remains the same in term of input data and in the absence of exogenous shock, the system is assumed to converge toward a stationary state. Indeed, users improve continuously their knowledge of the traffic condition. At one point, they will not be able to improve anymore their travel choices in order to minimize their travel cost. The simulator computes a generalization of Wardrop's first principle: at equilibrium, no user can strictly decrease it generalized by changing either his mode, his route or his departure time.

3.2. Learning process

The learning process can be seen as a black box that models day-to-day dynamics: it captures the fact that users take into account their past driving experiences in future travel decisions. The data processed here is the *information* that must be understood in a broad sense. It groups together historical drivers' information and instantaneous information. It consists in any data relevant to user travel choices, like data required to estimate time-dependent shortest routes, but also the expected variability of travel times under non-recurrent situations or data relevant to road tolls.

Input *Historical* information denotes the accumulated knowledge that users gain from using the transportation system days after days. Should traffic conditions never change (which is never the case in real life), this historical information would reach a stationary state and remain constant in a well behaved model. Instantaneous information refers to what each user may learn during the within-day simulation: perception of actual traffic conditions, forecasts or exogenously simulated information provided to the drivers, etc. *Instantaneous* information is continuously updated as the dynamic assignment proceeds, while historical information is only updated daily. Information is partially shared among the drivers (public information) and partially specific to individual drivers (private information). Information determines all the travel choices modeled in the simulator: Mode, departure time and route choices. The flexibility of this information handling allows, for instance, to provide different users segments with different sets of instantaneous information based on their equipment (e.g. navigation systems). *External* information consists in any information provided from the outside of the transportation system by means of information technology like radio, variable messages signs, on-board computers, etc.

Output Once a simulation day (or morning rush) is completed, the historical information is updated according to the driving conditions experienced by the simulated drivers. The updating process is based on a Bayesian mechanism that combines the historical information and the last day changes in the network usage according to the user characteristics (i.e. perception and cognitive abilities). Since these mechanisms are typically very expensive to implement literally at the numerical and computing memory point of view (see [13]),

we decided to resort to heuristic laws. The historical information available to users on day $\omega^{\chi}+1$ is the output of the learning process on day ω, that is, the accumulated knowledge of the ω previous days (i.e. a Markov process of order 1). For a given O-D pair, the expected travel time when departing at time t_d on day $\omega^{\chi}+1$ is computed as follows:

$$E^{\omega+1}\left(\tau\left(t_d\right)\right) = \left(1-\lambda\right) E^{\omega}\left(\tau\left(t_d\right)\right) + \lambda \tau^{\omega}\left(t_d\right)$$

Note that this process is considerably different from MSA (Moving Successive Average) since it does not converge necessarily. A typical value used in practice is $\lambda = 0{,}1$.

3.3. Generation of incidents

A possible solution to evaluate the impacts of non-recurrent traffic incidents would be to use a static traffic assignment approach. Given the distribution of the probability of occurrence of the incidents $f_I(t)$, links could be characterized by a stochastic capacity. Fully stochastic assignment algorithms can then be used to solve the problem (see [14]). This latter approach lacks however an important aspect which is within-day departure time choice adjustment. Indeed, if users encounter non-recurrent traffic congestion on certain part of the network at certain moment of the day, they might decide to schedule their trip at another time. Also, the impacts measured by the static approach are those of a supposedly long-term situation where the users have some-how discovered $f_I(t)$ by experiencing traffic conditions. But the transient situations to get to that situation might yield much bigger impacts that cannot be measured without taking into account explicit day-to-day dynamics. Note that a sophisticated dynamic traffic assignment is not sufficient since some mechanism has to be responsible for the adaptation of (exogenous) departure time profiles.

Exogenous traffic hazards are straightforward to introduce in event-based models such as METROPOLIS. At the beginning of each day (or morning peak), a random number is drawn for each link where a potential incident can happen. If $\mathbf{R} < p$ then an incident happens on that link ($f_I(t)$ is uniform). The probability of occurrence p is the same for a selected subset of important links of the network (e.g. main arterials with more than two lanes) and $p = 0$ for the rest of the network. The incidents are characterized by a capacity drop of 50% that lasts for the whole morning peak. The linear bottleneck congestion model is applied throughout the whole simulated period. Note that incidents could also be defined with a specific duration (e.g. half an hour for a stopped vehicle blocking a lane). We select only incidents that lasts for the whole simulated period (i.e. the morning peak) for the sake of simplicity. Nevertheless, these incidents still affect the departure time choice since roads are usually under-used (i.e. below capacity) at the beginning (and end) of the peak period. If the capacity drops recurrently during the peak period, some users will consider departing earlier or later to avoid congestion. Vehicles that reach an intersection are informed if an incident has happened on a downstream link so that en-route diversion is possible.

4. Simulation results

4.1. Control runs

A first set of experiments is performed on the well-known toy network of Sioux-Falls to test the day-to-day adjustment process. About 50,000 individual trips are simulated. Incidents are introduced on four links that are located on the north-south corridor identified on Fig. 2. Three simulations are performed:

1. base case scenario without incidents,
2. introduction of incidents in the corridors from day #20 on,
3. introduction of incidents in the corridor between day #20 and day #50.

Fig. 3 presents the travel costs for the 100 first iterations of the three simulations. It can be seen that the impacts of the incidents are quite large on day #20 since users are uninformed and did not yet had the opportunity to adapt their travel habit. This would be the impact reported by a dynamic traffic assignment model *without* feedback on departure time choice. Note that even the base case system exhibits some oscillations. This is due to the fact that the simulation model is stochastic and that the exponential smoothing process does not lead here to a unique situation but rather to a set of stationary states. We will show below that this does not hinder the evaluation of non-recurrent congestion. Two things are worth noting on this example: the adaptation in the case of the standing incidents is such that the travel cost follows a decreasing trend. Also in the case of the third scenario, the decrease in travel cost does not happen immediately because of the inertia of the system. Note that during the first twenty iterations the three curves are, of course, overlapping. The peak at iteration #10 shows that the system is not yet completely stabilized.

4.2. Impacts of incidents

The second set of experiments is performed on a real-world example for the Paris area. The coded network consists of about 17,000 links and more than 3,000,000 individual trips are simulated for each morning period (the computation of a single iteration takes around 20 minutes on an Apple G5 with a clock speed of 2Ghz). Both commuting trips and non-commuting trips are simulated to take off-peak congestion into account. The incidents are introduced on the major roads of the area, defined as the roads that have at least three lanes (see Fig. 4). Several simulations are run for different probabilities of occurrence p ranking from 0 to 1. In each case, 50 iterations (days) are performed. The corresponding global indicators are reported in Tab. 1. The impacts are important because each traveler uses on average at least one road sections that belong to the major roads (there are about 1,000 major-road links and the average trip length is 16 links). The schedule delay costs measures the penalties incurred by travelers arriving too early or too late at their destination. Obviously, drivers arrive later than expected when the level of incidents increases. Note that the indicators show that the case p=0.5 is worse than the risk-free case p=1, even if the average capacity of the overall system is higher. This discrepancy between the stochastic and the risk-free situation is evaluated in the next section.

Table 1: Impacts of non-recurrent congestion.

p	0	0.25	0.5	1
Travel cost [$]	7.6	11.3	12.0	11.9
Schedule delay cost [$]	2.2	3.2	3.4	3.3
Travel time [min.]	25.0	37.4	40.0	39.8
Change in consumer surplus [$]	0	-1.9	-2.3	-2.1
Congestion index [%]	28.9	55.4	60.8	65.0
Mileage [10^6km]	56.5	59.8	60.2	59.8
Early arrivals [%]	49.5	48.2	48.1	46.1
Late arrivals [%]	29.7	34.7	35.0	36.5

4.3. Utility loss

The compensation χ introduced above for travel *time* variability is measured on the Sioux Falls example. This compensation is extended here to the compensation of travel *cost* variability instead. This value can then be interpreted as the utility loss due to uncertainty (or *monetarization* of uncertainty). Total travel costs includes free flow costs, queuing costs and schedule delay costs. The variability $f(C(t))$ is computed by recording the total travel costs for T = 300 iterations. Therefore, the monetarization of uncertainty is computed as follows:

$$\chi_\theta = \frac{1}{\theta} \ln \left\{ \frac{1}{T} \sum_{i=1}^{i=T} e^{\theta C_i} \right\} - \frac{1}{T} \sum_{i=1}^{i=T} C_i$$

Incidents are introduced during the T iterations with same level of probability p. Several runs are performed for values of p ranking from 0 to 1. Fig. 5 presents the results of the evaluation of χ_θ as a function of the level of risk aversion θ. As shown in the control runs (Fig. 3), travel costs can oscillate even without incidents because of the stochastic nature of the traffic simulation. Nevertheless, those oscillations are rather small, which explains why the risk-free situations (p=0 and p=1) correspond to the bottom curves associated with an (almost) null cost for uncertainty. As for the Paris example, the case p=0.5 yields the highest costs of uncertainty. Obviously it is not a linear dependence (see the discrepancies between (p=0.4 and p=0.6).If we assume θ = 10$ (a value compatible with the survey results of [4]), we get the relative impacts is $\phi = \frac{\chi}{\langle C \rangle}$ reported in Tab. 2. The travel time equivalent of the compensation, τ_χ, is also computed by assuming a value of time of 10$/$h$.

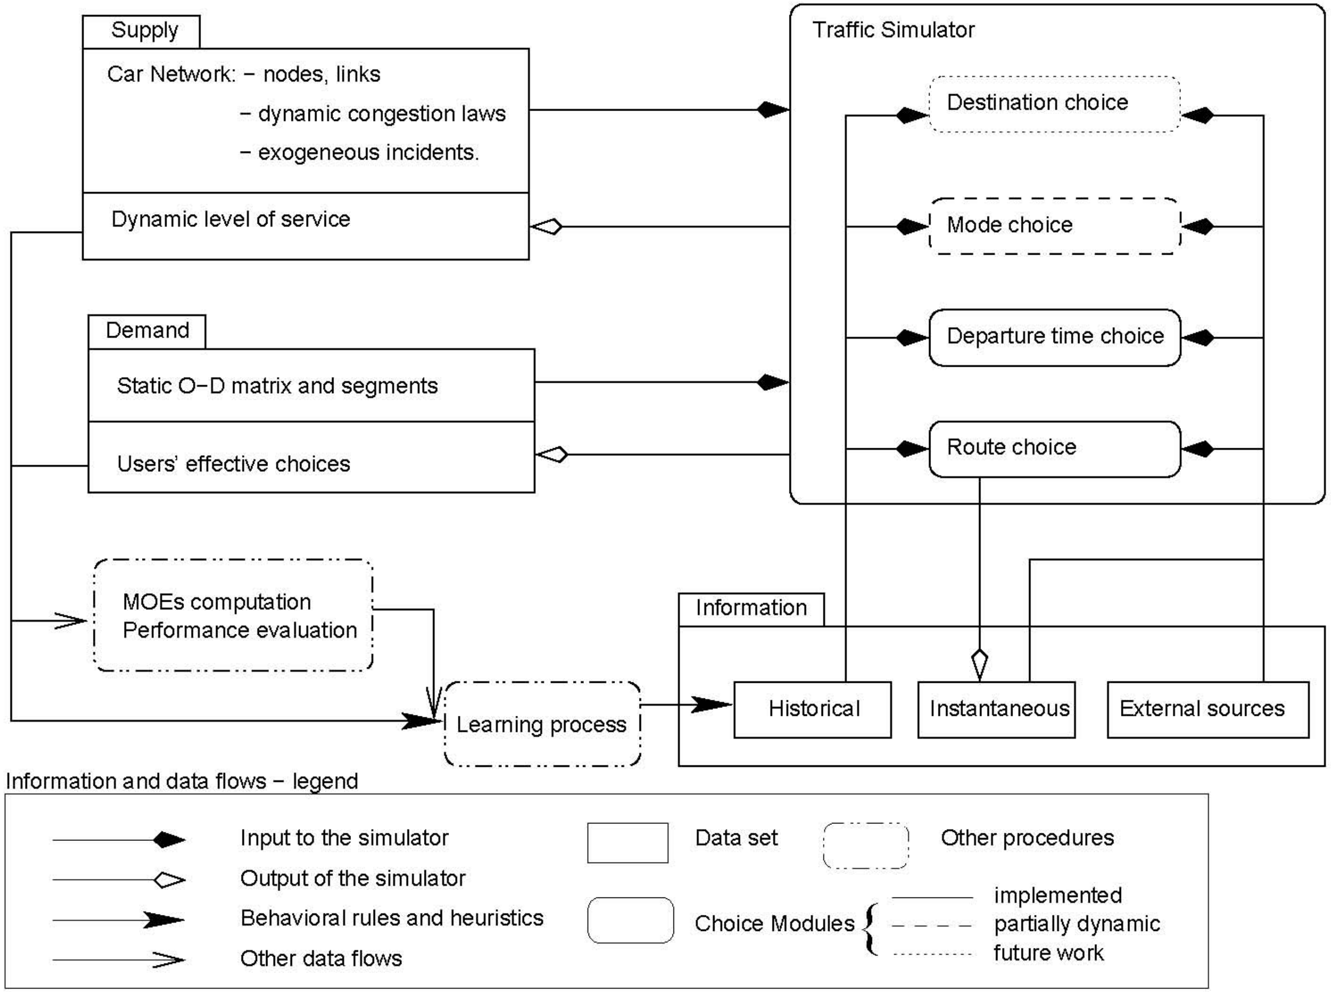

Figure 1: Architecture of the day-to-day learning process implemented in the METROPOLIS simulation environment.

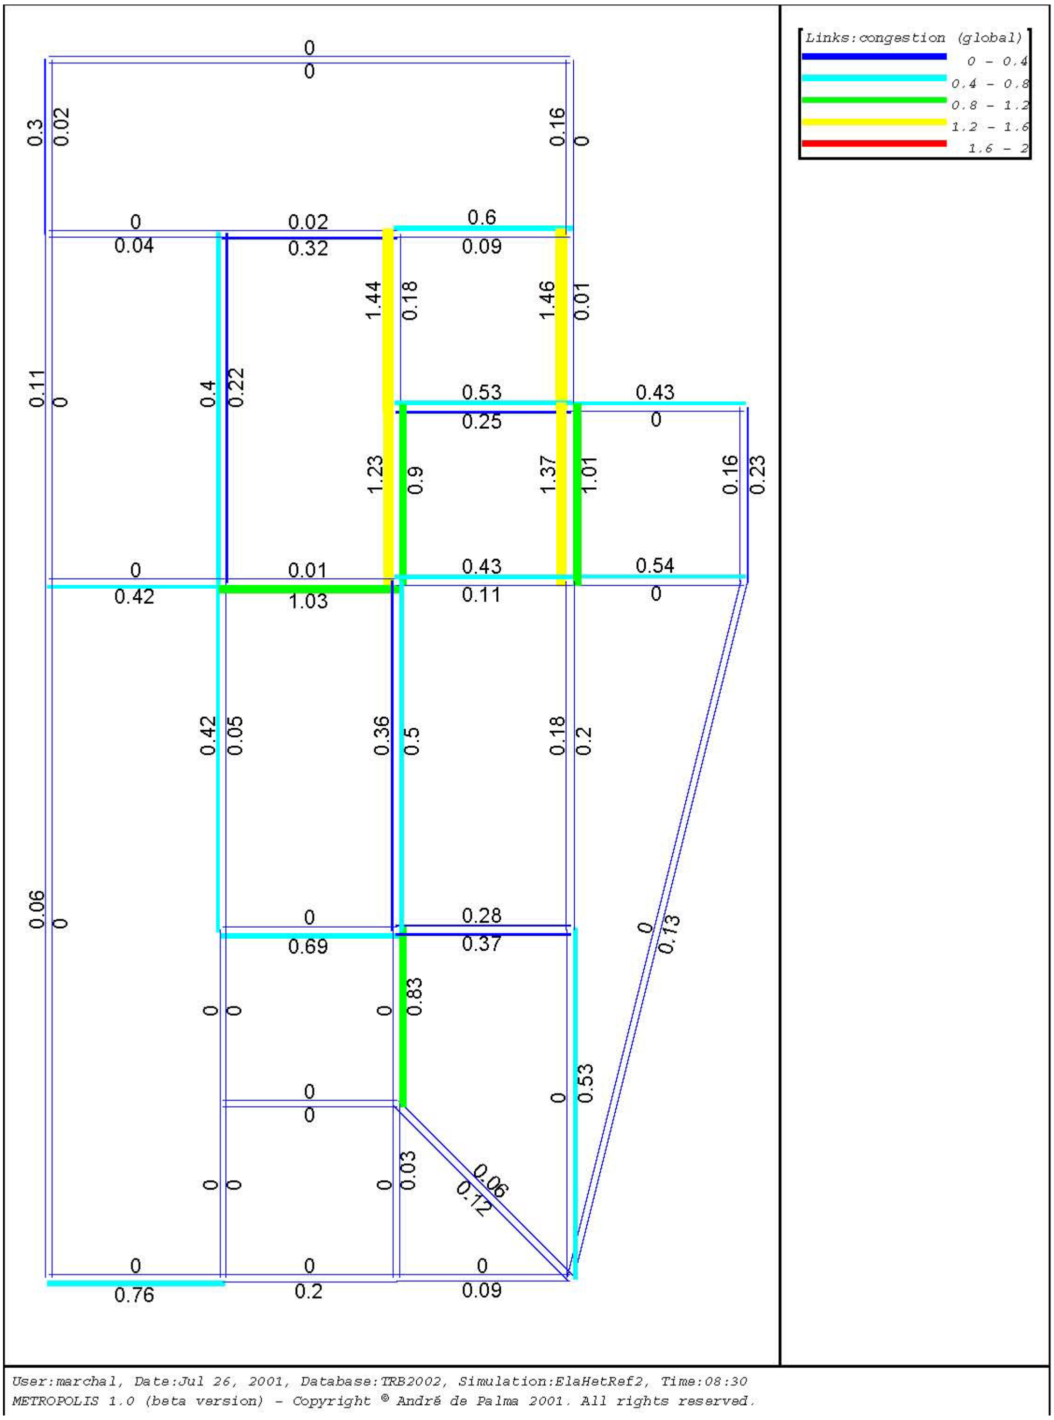

Figure 2: Congestion index on the Sioux-Falls network in the base case scenario (without incidents). Incidents are introduced on the most loaded north-south corridors (yellow links with index>1)

Table 2: Evaluations of the cost of uncertainty for an individual with a risk aversion level of θ = 10$ and a value of time of 10$/*h*

p	0.25	0.40	0.50	0.60	0.75	0.90
χ[$]	2.10	2.55	3.35	1.35	1.15	0.45
τ_χ[min.]	12	15	20	8	7	3
ϕ	0.37	0.45	0.57	0.23	0.20	0.08

Figure 3: Control runs. Three simulations with and without incidents. Day-to-day evolution of the total travel cost.

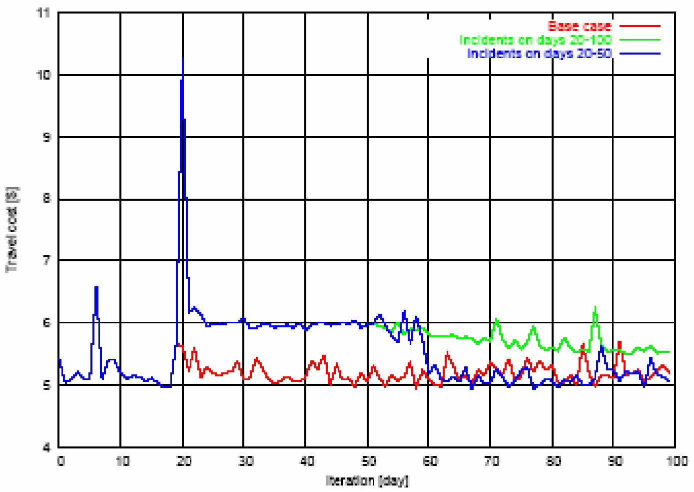

Figure 4: Ile-de-France area surrounding Paris. Incidents are introduced on the major roads (red).

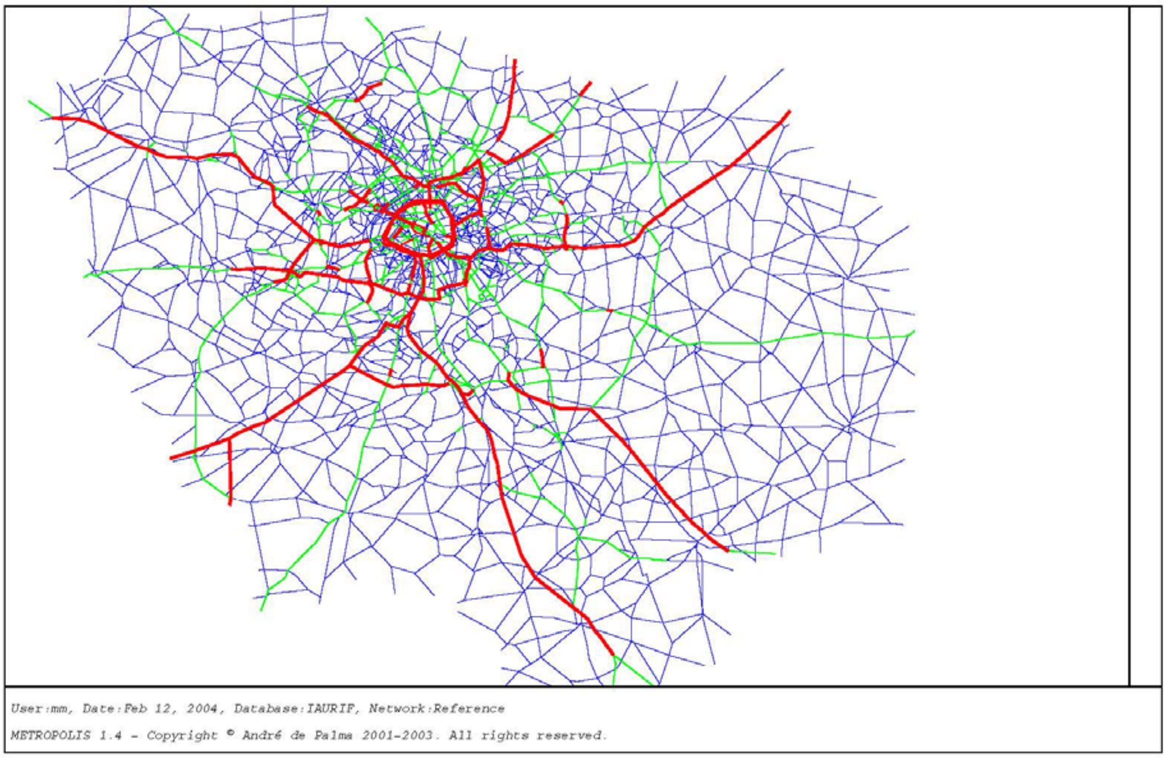

Figure 5: Monetarization of uncertainty. Measurements for different probability *p* of occurrence.

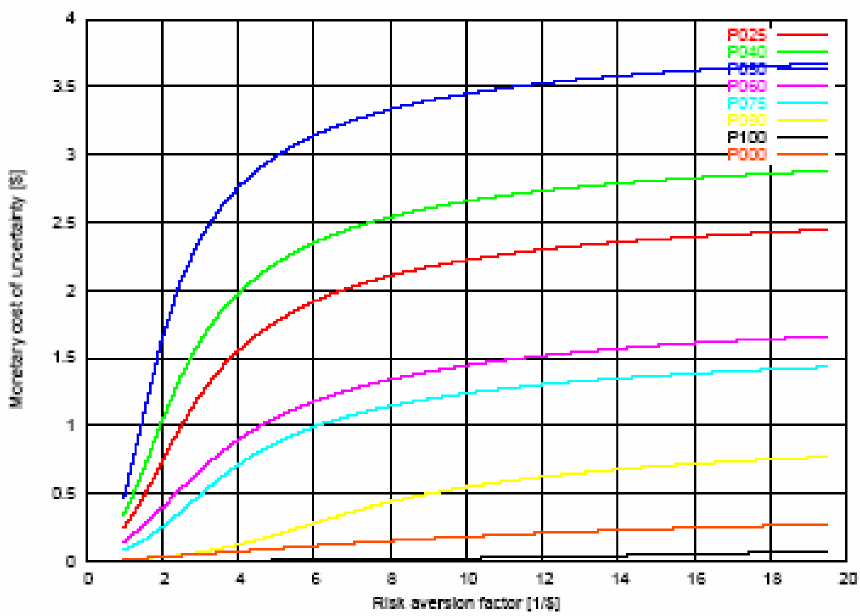

5. Concluding comments

We have developed in this paper a method to measure both recurrent and non-recurrent congestion in transportation systems. The methodology relies heavily on the usage of METROPOLIS, a dynamic traffic simulation tool that is able to handle very large realistic networks. Random incidents are introduced in the system and users react to them by adapting their departure time and route on a within-day and day-to-day basis. We have stressed out the importance of day-to-day adjustment process by showing that the distinction between recurrent and non-recurrent congestion is blurred: if the same incidents occur with a given probability on a long period, users might eventually learn how to adapt and take into account travel time variability as an additional cost. This additional cost or monetarization of uncertainty has been estimated using both empirical measures and simulation results. This evaluation is fundamental to the designers and managers of ATIS and ATMS since it gives a benchmark of the potential benefit that can be obtained with such technologies.

Acknowledgments

We would like to thank IAURIF and more particularly Joseph Berthet and Dany Nguyen who kindly provided us with the data for the simulations on Paris. Computing resources and support were provided by CoLab-ETHZ.

References

Cohen, M., Jaffray, J.-Y., and Tanios, S., 1987. Experimental Comparison of Individual Behavior under Risk and under Uncertainty for Gains and for Losses, Organizational Behavior and Human Processes, 39 1-22.

Cohen, M. and Jaffray, J.Y., 1988. Certainty Effect versus Probability Distortion: An experimental Analysis of Decision Making under Risk, Journal of Experimental Psychology, 14 (4) 554-560.

Demange, G., and Rochet, J.-C., 1992. Méthodes Mathématiques de la Finance, Economica.

De Palma, A., and Picard, N., 2003. Route Choice Under Uncertainty, Paper presented at IATBR, Luzern, Switzerland.

El Amri, W., 1999. L'analyse du choix de portefeuille: un modèle de choix discret. Master Thesis, University of Tunis III, Institute Supérieur de Gestion, Tunisie.

Kahneman, P., and Tversky, A., 1979. Prospect Theory: an Analysis of Decision under Risk, Econometrica, 47 263-291.

Machina, M. (1982). "A Stronger Characterization of Declining Risk Aversion", Econometrica, 50, 1069-1079.

Pratt, J., 1964. Risk Aversion in the Small and in the Large, Econometrica, 32 122-136.

Tversky, A., 1967. Additivity, Utility and Subjective Probability, Journal of Mathematical Psychology, 4 175-201.

de Palma, A., Marchal, F., and Nesterov, Y., 1997. METROPOLIS: A Modular System for Dynamic Traffic Simulation. Transportation Research Record 1607, 178-184, Washington, D.C.

de Palma, A., and Marchal, F., 2001. Dynamic Traffic Analysis with Static Data: Some Guidelines with an Application to Paris, Transportation Research Record 1756, 76-83, Washington, D.C.

de Palma, A. and Marchal, F., 2002. Real Case Applications of the Fully Dynamic METROPOLIS Toolbox: an Advocacy for Large-Scale Mesoscopic Transportation Systems, Networks and Spatial Economics 2, 347-369.

Arentze, T., and Timmermans, H., 2003. Representing Mental Maps and Cognitive Learning in Micro-Simulation Models of Activity-Travel Choice Dynamics, paper presented at the 10th International Conference on Travel Behaviour Research, Lucerne, August 2003.

Hazelton, M., 2003. Total Travel Cost in Stochastic Assignment Models, Networks and Spatial Economics, 3 457-466.