NISS

Quality of Information Given By Advanced Travel Information Systems

Piyushimita Thakuriah and Ashish Sen

Technical Report Number 51
October, 1996

National Institute of Statistical Sciences
19 T. W. Alexander Drive
PO Box 14006
Research Triangle Park, NC 27709-4006
www.niss.org

QUALITY OF INFORMATION GIVEN BY ADVANCED TRAVELER INFORMATION SYSTEMS*

Piyushimita (Vonu) Thakuriah and Ashish Sen

Abstract

A number of studies have evaluated the services provided by Advanced Traveler Information Systems [ATIS] under the assumption that information supplied to drivers would be in some sense perfect. However, lack of sufficiently useful data and system design constraints can lead to information that is less than useful to the ATIS user. This paper examines the effects of such imperfection through a simulation-based model that was applied over a part of a large metropolitan area. The model has four basic components: (i) an ATIS structure (that specifies the information-gathering, processing and disseminating aspects of the system) (ii) traveler behavior (iii) network characteristics and (iv) vehicle movement logic. Using a 'yoked driver' concept, a number of different route guidance strategies are examined. The results indicate that some strategies that would appear to be desirable are not so. On the hand, under high-congestion situations, strategies can be constructed that come close to 'rectifying' completely the effects of information imperfection. Overall the paper reiterates the potential of ATIS if information-giving strategies are designed carefully.

1 Introduction and Overview of the Study

A number of studies have shown substantial benefits in travel times incurred by users of Advanced Traveler Information Systems [ATIS]. An ATIS is a functional area within Intelligent Transportation Systems [ITS] that offers a variety of services to users including navigational aids and yello-page type information. A principal component of ATIS is route guidance, or the capability of providing drivers with desirable routes from their trip origins to destinations, based on historical, current or predicted travel conditions. Most studies of ATIS benefits have assumed that drivers would receive perfect information from the system. However, the information provided by an ATIS could be far from perfect.

In an earlier paper, Koutsopolous and Xin (1993) pointed out that the effectiveness of dynamic route guidance systems depend to a great extent on the ability of system designers to address issues of information frequency and reliability. Our basic premise is along the same lines and the results presented in this paper serves to reiterate the fact that the information provided by an ATIS may not be useful to drivers unless adequate steps are taken to counter the effects of several factors that can be particularly damaging to the quality of travel time estimates.

The study, motivated by the authors' involvement with the design of a short-term travel time prediction sub-system within an operational ATIS, (ADVANCE, see Boyce *et al.*, 1994) considers one class of ATIS architectures — that which uses vehicles as floating probes to monitor traffic conditions and to gather link travel time data. Estimates of average travel times for every link and time period within a day are computed and placed on CD-ROMS within each vehicle. These estimates are called *static estimates* and route guidance based only on them is called *static guidance.* In addition, current reports from probes

*This work was partially supported by the National Science Foundation's Division of Mathematical Sciences Grant Numbers 9313013 and 9208758.

are used by a centralized traffic information center to compute current and predicted link travel times according to some prediction rules. These *dynamic* reports are sent via radio to equipped vehicles and the static estimates over-written. If, owing to lack of data or equipment malfunction, dynamic travel times are not available for a link, the default static estimates are used. Routes based on this combination — dynamic estimates where available and static estimates as defaults — are called *dynamic routes* and vehicles receiving them are said to receive *dynamic guidance.* The route is planned in-vehicle.

The use of probes to obtain input data for predictions raises complex sampling issues in computing the 'best' route for the user. What is inferred about the state of the network from these sampled observations, may not necessarily reflect actual conditions to the fullest extent. Moreover, in the highly stochastic conditions prevailing on the road network system, the best (smallest cost) route based on such partial observations may not necessarily be the actual best route for an origin-destination pair at a particular start time. But even if it were indeed the best route for those conditions, it may cease to be the best route as the guided driver proceeds toward the destination, due to changing network conditions.

In this study, we evaluate the quality of routes provided to ATIS users (in terms of route travel times) by explicitly considering the effects of sampling variations in the information supplied by a probe-based ATIS. In order to do this, we used a Monte Carlo simulation-based model that allowed us to study the effects of the following on the information supplied by ATIS: (i) the inherent variability in network flows and travel times even during recurrent congestion, due to complex traveler behavior and choices, which is further compounded by network control factors; (ii) partial observation of the network and sparseness of observed information during short intervals of time because of limited data gathering mechanisms; and (iii) the inability of an ATIS to adequately process and distribute the available information at the relevant clock time due to system design constraints, which further exacerbates the partial observation and data sparseness problem.

The results indicate that under recurrent congestion, irrespective of deployment levels, if dynamic estimates of link travel times are broadcast for links selected in accordance with certain carefully developed criteria, dynamic guidance can not only be superior to static guidance, but indeed provide very good guidance when congestion levels are high. However, dynamic guidance based on less well conceived criteria, such as developing routes based on all mean travel times, can lead to estimates *inferior to* static guidance, even for fairly high levels of probe deployment [i.e., even if most vehicles were probes]. This occurs because variances of dynamic travel time estimates remain high even at high deployment levels as compared to static estimates which are based on data for several days.

It needs to be pointed out that while the authors have drawn extensively from their experience with ADVANCE, we believe that the major issues addressed in this paper would hold for virtually all types of real-time probe-based route guidance systems, where minimum travel time routes are to be given to guided drivers.

A concise description of the method used to simulate route guidance follows. The day-to-day variability in link volumes, for similar day-types, exhibit a definite affinity to a mean volume level. On the other hand, the same-day over-time origin-destination demand creates a pattern of link volumes that can be characterized by a certain probability distribution. The knowledge of this distribution allows us to devise a sampling procedure that retains the temporal and spatial dependence of network volumes. We simulate travel times by means of an iterative two step procedure within each iteration. In the first step, volumes on links in the network are simulated for each time interval; in the second step, these volumes are used to calculate the travel time of each vehicle, but not at the detail of many available vehicle microsimulation models, because the latter were deemed as computationally prohibitive for the current exercise. The travel times of some of these vehicles are sampled. From these sampled travel times, an estimate of link travel times is obtained and supplied to guided vehicles by means of one of several *information-giving* strategies, that include corrections for sparseness of inbound real-time and historical information, time-lag effects, transmission constraints in distributing information and so on. The number of vehicles sampled is determined by the deployment rate, that is, the percentage of vehicles that are participants of the ATIS. Several different deployment rates are considered. The advice obtained under each estimation strategy is given to a single vehicle. In several ways, some of the evaluation designs are reminiscent of the yoked driver experiment used in PATHFINDER [Santa Monica Smart Corridor, Los Angeles (JHK & Associates, 1993)], where a number of drivers are asked to travel from a predetermined origin to a destination on the basis of information supplied in different ways.

We have evaluated route travel time savings under recurrent congestion. While it is true that an adequately designed ATIS would allow users to incur substantial savings during periods of non-recurrent congestion, the more typical case in road networks is one of recurrent congestion (Tarko, 1994).

The paper is organized as follows: Section 2 describes the model that we used to answer route quality evaluation questions. In Sections 2.1 through 2.4, we present the theoretical underpinnings, implementation considerations and design features of the model. We then present some illustrative examples of route guidance information-giving strategies, as discussed above, in Section 2.5, given various factors that affect the quality of travel time estimates. The results and implications for route guidance are discussed in Section 3. A summary of the major findings of the evaluation experiments are presented in Section 4 along with some of their implications.

2 Description of the Simulation

The main purpose of this study are twofold:

1. to compare travel times on routes recommended to drivers by the ATIS versus the actual best route prevailing during the relevant time period;
2. to compare the travel times taken by vehicles on routes formed under each of several information-giving strategy, alluded to in Section 1. These strategies have implications for the development of a short-term travel time prediction sub-system within an ATIS.

Since our focus is on the comparison of routes which are optimal according to different sets of travel times, these evaluations are virtually impossible to conduct by any means other than simulation. The following sections describe the simulation-based model that was devised for this purpose.

2.1 Network and Simulated Data

The evaluation experiments (of the 'yoked' driver nature) that we implemented required us to consider one predetermined and fairly widely separated origin-destination pair. Therefore, we considered a 329 link sub-network, which included such a pair, from a larger network in the northwestern part of the Chicago Area [the ADVANCE test area].

The subnetwork included all links that comprised the 70 best routes based on travel times yielded by a model described below [Section 2.2]. This was necessitated by the fact that routes between this O-D pair would be compared. However, traffic volumes on these 329 links were composed of volumes between all O-D pairs within the larger network that used these links. All links are two-directional. We did not have any information on signal control at the time the experiment was conducted, except that all 329 links are signalized. The signal control at each link was assumed to be pretimed, with two phases.

2.2 Network Volumes Processes

The temporal variations in link volumes are the result of changes in demand for travel between origins and destinations of trips and also in shifts in routes and departure times. It is reasonable to assume that people do not choose routes based on actual congestion levels but on expected levels (unless they are somehow informed about traffic conditions). Hence it is appropriate to have route selection procedures based on perceived travel times, which would be based on expected travel times.

The distribution of link volume $v_{\ell,t}$ for a given link ℓ and time-period t is crucial for Monte Carlo simulation. Let the travel demand between origin $i \in \mathcal{I}$ and destination $j \in \mathcal{J}$ starting from i during interval $I_{t'}$ be $N_{ij}(I_{t'})$ and let $\tau_{ij}(I_{t'})$ be its expectation, that is, $\tau_{ij}(I_{t'}) = \mathrm{E}[N_{ij}(I_{t'})]$. We drop the subscripts i and j from our notation with the understanding that routes relate to ij pairs. Under mild conditions, the aggregate flows between i and j are known to have independent Poisson distributions (Smith, 1987). Consider a long time interval I [e.g., a day] and partition it into non-overlapping subintervals $I = \cup_{t'} I_{t'}$. If $N(I')$ were Poisson and the trips were randomly assigned to the subintervals with known expectations $\mathrm{E}[N(I_{t'})]$, then the $N(I_{t'})$'s would be independent Poisson. Thus, this model includes departure time choice, as long as it is random with fixed probabilities.

Suppose K routes are used by travelers from i to j. Let $p^k(I_{t'})$ be the probability that a traveler uses route k and starts his or her trip during interval $I_{t'}$. Let the number of travelers who depart from i within $I_{t'}$ and choose to travel by route $k \in \mathcal{K}$ be $N^k(I_{t'})$. If each trip between i and j on route k is independent of others and $N(I_{t'})$ is fixed, then $\boldsymbol{N^K}(I_{t'}) = \left[N^1(I_{t'}), N^2(I_{t'}), \ldots, N^{K_{ij}}(I_{t'})\right]$ is a multinomial random variable. Then, since each $N(I_{t'})$ varies as Poisson, the $N^k(I_{t'})$'s can be shown to have independent Poisson distributions. This is shown in Thakuriah (1994). An important consequence of the Poisson characterization of route flows is that the only statistical parameters which govern route flow patterns are mean flows $\mathrm{E}[N^k]$.

Let $v^{\ell,k}(I_t, I_{t'})$ be the number of vehicles on route $k \in \mathcal{K}$, departing origin $i \in \mathcal{I}$ during interval $I_{t'}$, and entering link ℓ during time interval I_t. That is, we consider a partitioning of all vehicles departing i during time period $I_{t'}$ — those that are on link ℓ during interval I_t and those that are not. Then, under the independence assumptions already made, $v^{\ell,k}(I_t, I_{t'})$ is binomial, given $N^k(I_{t'})$. Since $N^k(I_{t'})$ is Poisson, it follows that $v^{\ell,k}(I_t, I_{t'})$ is Poisson. Since, total link flows are sums of these $v^{\ell,k}(I_t, I_{t'})$, it follows in addition that link flows are Poisson, as indeed is often assumed in traffic studies.

To empirically verify the Poisson link flow conjecture, we analyzed volume data that were available from nine road detectors along Roselle Road in the northwestern suburbs of Chicago. The volumes had been aggregated over 15 minute intervals, starting each midnight and extending to the following midnight. We analyzed the data from all nine detectors (Thakuriah, 1994) and found that, after accounting for (i) temporal aggregation that occurs by forming discrete-time estimates of the continuous link volume stochastic process and (ii) effects of oversaturation, the Poisson conjecture is a reasonable approximation of the time-varying link volume process.

Unguided travelers, for the most part, choose routes before they start on their trips. As already mentioned, travelers probably use some estimates of the expected route travel times in making their route selection decisions. We assume, therefore, that these choice probabilities are independent of the random variable $v_{\ell,t}$, that actually determines travelers' travel times on their trips, on the day d of travel. Instead, they are assumed to be dependent on average conditions prevailing over several days.

If expected flows between each origin, i and destination, j are known then link volumes can be computed using a model that gives a Stochastic User Equilibrium [SUE] solution, (Daganzo and Sheffi, 1977, Sheffi, 1985, see also Hicks, *et al*, 1992). These hourly link volumes were divided by 16 to obtain volumes for each of the 3.75 minute time-slices that we used in the simulation. The resultant volumes were then treated as the link volumes to obtain travel times as described in the next section. The travel times thus obtained were used to revise mean link volumes. These mean volumes were then used to generate independent Poisson 'realizations' which were treated as initial link volumes for every link, time-period and day. These initial volumes were modified as described below.

Volumes on contiguous links are correlated because, within a time slice, many of the vehicles on them are in fact the same. Ignoring this correlation would have been convenient but inappropriate. If link volumes and hence link travel times were to be independent, route travel times would merely be the sum of independent link travel times. The Law of Large Numbers would then reduce the variance in travel times for long routes to the point where there would be very little need for guidance, except under situations of non-recurrent congestion!

Our model uses a link correlation structure that preserves the correlation between volumes on contiguous links without having to track the routes of individual or groups of vehicles, by using the following approach. Suppose volumes v_m are being discharged from links ℓ_m on a simulated day d for a specific time-slice where each $\ell_m : m \in \mathcal{M}$ is immediately upstream of link ℓ with volume v_ℓ. Let $\mathcal{L}$ be the set of links which belong to at least one route in $\mathcal{K}$ and let $\mathcal{R} = \mathcal{M} \cap \mathcal{L}$. Suppose also that a vehicle from link ℓ_m, $m = 1, \ldots, M$, goes into link ℓ with probability P_m. Then $\mathrm{E}[v_\ell] = \sum_m P_m \, \mathrm{E}[v_m]$. For each vehicle being discharged from link $\ell_r : r \in \mathcal{R}$ during a time-slice, Bernoulli trials were used to determine which ones would enter link ℓ. Let the number of such vehicles be $v'_r : r \in \mathcal{R}$.

Let $\nu = \mathrm{E}[v_\ell] + \sum_{r \in \mathcal{R}} (v'_r - \mathrm{E}[v'_r])$. Notice that

$$\mathrm{E}[v'_r] = P_r \, \mathrm{E}[v_r]$$

and

$$\mathrm{E}[\nu] = \mathrm{E}[v_\ell].$$

Therefore, if we use a Monte Carlo method to generate a realization of a Poisson random variable with mean

$$\omega = \mathrm{E}[v_\ell] - \sum_{r \in \mathcal{R}} \mathrm{E}[v'_r]$$

and add it to $\sum_{r \in \mathcal{R}} v'_r$, then the resultant would be Poisson with mean $\mathrm{E}[v_\ell]$. This is precisely what we did. We generated a realization of the Poisson random variables with mean ω and added to it $\sum_{r \in \mathcal{R}} v'_r$ to get v_ℓ for each time interval.

Moreover, the volume level for a link in one time period is strongly related to those in contiguous time intervals. This temporal correlation over successive time intervals were ensured by the congestion build-up and decay due to oversaturated and overflow-queuing effects simulated by the travel time component of the model, as discussed in Section 2.3.

2.3 Link Travel Times

In this section, we will discuss the method used for simulating link travel times. Since link travel time varies over clock-time, we need to index it. We have chosen to index it by the time of entry of the vehicle into the link. Since it is specific to a given day d, we will denote it by $T_{t,d,\ell}$, where t is the time, on a 24-hour clock, of entry into link ℓ. We simulated $n = 35$ days and set $T_{t,\ell} = n^{-1} m_{I_t}^{-1} \sum_{d=1}^{n} \sum_{t \in I_t} T_{t,d,\ell}$ as our static estimates, while $\sum_{t \in I_t} w_t T_{t,d,\ell}$ would be the dynamic estimates for day d, where m_{I_t} is the number probes entering link ℓ during interval I_t on day d and w_t is a weight which when it is $m_{I_t}^{-1}$ yields a simple mean. However, we shall also consider cases where w_{I_t} increases with t as means of addressing information aging.

The rest of the section is devoted to the computation of the individual $T_{t,d,\ell}$'s. The method used is exactly equivalent to the time-based simulation described below — in the sense that the actual numbers produced would be identical. However, the actual method used is not identical to the description given for two reasons: for computational efficiency some of the steps were combined into larger steps and some modifications were necessitated in order to make use of portions of programs that had been previously written.

We divided every traffic-light cycle into 2-second intervals. We also assumed that the time s taken for a single vehicle to discharge from an intersection is two seconds. At the end of every 2-second interval, the simulation performs the following steps:

Step 1. Zero or more vehicles enter the link. For immediately upstream links belonging to a path in $\mathcal{K}$, it is easy to determine whether a vehicle from that link is entering the link ℓ currently under consideration during the 2-second interval [see Step 3, below]. For other immediately preceding links, whether a vehicle enters link ℓ during a particular 2-second interval is determined on the basis of the following assumption: vehicles are assumed to enter link ℓ uniformly within their respective green times, the spacing being determined by the appropriate volume. This assumption is not entirely realistic; but its effects are minimal.

Step 2. The position of every vehicle not in the queue at the end of the link is adjusted. Each vehicle is assumed to travel at a constant speed which could vary from vehicle to vehicle. The cruise speed is

$$cs_v = sp + ps_1 + \rho_v ps_2 \tag{1}$$

where sp is the posted speed limit, ps_1 and ps_2 are positive numbers that vary with congestion levels and ρ_v is a random number in the [0,1] interval pertaining to the vth vehicle. Therefore, the adjustment in vehicle position consists of moving the vehicle forward by $2 \times cs_v \times 5280/3600$ feet [the last factors converting miles per hour into feet per second] unless

- the vehicle reaches the end of the queue where it is stopped, and the queue length appropriately augmented;
- the vehicle reaches the intersection in the absence of a queue where it enters the intersection if the signal phase is green, or stops [and starts a queue] if the phase is red.

Step 3. If a queue is present and the signal phase is green, one vehicle is released into the intersection and the position of all other vehicles in the queue appropriately adjusted [by slightly over one vehicle length, which is randomly generated with mean of 20 feet]. If the phase is red, no change occurs. For every vehicle entering an intersection, the choice of the next link is determined by means of a Bernoulli trial as described in the last section and the required input created for the next link in the route [$\in \mathcal{K}$].

Step 4. After appropriate numbers of these 2-second intervals have been considered, the signal phase is changed from green to red or *vice versa*. The cycle length, the length of the red phase and effective green phase are integer multiples of this basic interval.

As already mentioned, we only consider links with signal controls (as opposed to freeway links). In such links, a vehicle's link travel time has two components: a cruise time and an intersection delay. Cruise time is that taken by a vehicle to traverse the link from the time it enters the link up to the time it joins the queue, if any, at the downstream signalized approach. Intersection delays due to signalization are typically of two types: (i) uniform delay, which is incurred by a vehicle that arrives during the red time of the signal cycle and has to wait until the next green time to depart; (ii) overflow delay, which is caused when demand exceeds signal capacity, and some vehicles have to remain at the intersection at the end of the green time following the vehicle's arrival at the intersection. Overflow delay could occur due to sustained periods of oversaturation or due to random fluctuations in the arrival flow rate. As described above, the link travel time generated in the simulation for each vehicle is the sum of a cruise time and an intersection delay which can be of either type.

The $T_{t,d,\ell}$'s used in the analysis were generated for only a short period of time every day. We call this period the study window. However, the simulation was run for two simulated hours before the beginning of the study window. This was done to reduce the effects of any initial assumptions and to get realistic conditions on the road system [for example, initial queue lengths], before starting to 'gather' data for creating static and dynamic estimates.

The program for the simulation was written by the first author. Previously available traffic simulation models were not found to be suitable for this analysis for several reasons. Our study required the simulation of a fairly large network for a total of 140 'days.' A traditional vehicle microsimulation would have required exorbitant resources and traffic macrosimulation models would not have allowed us to obtain link travel times of individual vehicles. Also, we needed to track individual guided vehicles along pre-assigned routes, that are determined by the information system. The simulation used in this study, while containing all relevant aspects of traffic behavior that were crucial to our study, was still economical in terms of computer time needs to provide the type of platform which enabled us to carry out the desired investigation.

2.4 The Design of the Simulation

In the previous two sections, we described the method used to generate traffic volumes on each link and the conversion of these volumes into travel times for each individual vehicle. In this section, we describe how these times were used to compute static and dynamic routes, as well as routes we call omniscient driver routes. Finally, we describe the manner in which we track guided vehicles across the network in simulated time and space.

2.4.1 Static and Dynamic Estimates

Static travel times were generated for every day and time of day during the study window. This was done as described earlier — by averaging average link travel times for each 15-minute interval over 35 days. Two congestion levels (see Section 2.5 below) were simulated for each of these 35 'days' yielding a set of static estimates for high congestion levels and a set for medium congestion levels. In the rare instances where no observations were available for any day for the relevant time period, the SUE estimate was used as a default. The optimal route on the basis of these static estimates is called a static route.

Although in our simulation we did not distinguish between type of day (our 'simulated' days are all of the same day type), various strategies can be used to enhance the quality of static guidance in an actual operational ATIS, by considering historical information of type-of-day profiles, presence/absence of construction and repair, special events, etc [see Sen and Thakuriah, 1995].

Dynamic estimates were, of course, for single days. In the results presented here, each of these days were distinct from the ones from which static estimates were computed. The method for computing dynamic estimates has been partially described already. However, several different strategies for providing dynamic estimates were examined, which we describe in Section 2.5. Dynamic estimates were generated for each strategy within each of the 35 simulated 'days' during which the comparisons were drawn and evaluations carried out. As mentioned earlier, for links not possessing such dynamic estimates during a simulated time period on a simulated day, static estimates were used as defaults. The optimal route obtained on the basis of these times are called the dynamic route for that strategy.

Naturally, the 35 days for which dynamic estimates were obtained were also the days for which routes given by each strategy described in Section 2.5 were compared. Two congestion levels were examined for each of the 35 days. Thus a dynamic route according to each strategy and during each congestion level was computed for each of 35 days. For each of these 35 days, each type of dynamic route and the static route was supplied to a single driver, whose travel time on the route that they were advised to take was then compared. Our comparisons in Section 3 were based on these times. Since we considered two levels of congestion, we simulated our study window for a total of 140 days.

2.4.2 Distinction between guided vehicles and probes

Notice also that in this simulation, we distinguish between probes and guided vehicles. Probes only gather data on link [not route] travel times and receive no guidance while the guided vehicle gets the benefit of this information but is not a probe. This distinction, along with the use of a single guided vehicle per strategy per day avoids a number of complications. For example, if probes were also guided, we could have situations in which no reports would be forthcoming from particularly congested routes and also that of creating congestion simply by guiding a number of vehicles on to the same link. Thus our approach allows us to focus solely on information quality aspects of ATIS, which is the purpose of this study.

We assume that every guided driver follows the route information that he or she has been given (that is, we did not introduce issues of non-compliance into the simulation).

2.4.3 The Omniscient Driver

We compare the travel times on static and dynamic routes with the travel time on the actual best route for the time that the guidance was provided. The actual best route is that predicted by a 'perfect' information system, which disseminates route information on the basis of link travel times that would be exactly the same as that experienced by a driver, if the realization of the link travel time random variable were known to the system before the route is given.

Hence, we use the concept of an omniscient driver who, at the time he/she chooses the route, has perfect knowledge of the conditions prevailing in the entire network at every time for the entire duration of the trip. In the simulated ATIS, link travel time estimates are created based on the observations that are gathered by the data-gathering elements of the ATIS architecture, which is then processed by other components of the ATIS. Hence, the guidance received by the guided driver may be far from perfect. On the other hand, the omniscient driver chooses a route that is given by a perfect ATIS, which means that he is able to choose his route on the basis of link travel time estimates that reflect conditions exactly as he *would experience them*. The omniscient driver needs to be clairvoyant and is hence fictional and such choices cannot be known except by simulation. The route travel time of an omniscient driver, who departs at the same time that a guided driver does, provides a benchmark for evaluating how far the system-provided route information is from perfection.

2.4.4 The Study Window

In order to track vehicles in the simulation, we used a sliding window view of the network. This concept is not unlike the 'rolling horizon' used by The Center for Transportation Research, University of Texas at Austin (1994), except in that case, the window was used for the assignment of vehicles to routes.

The ultimate purpose of the simulation is to provide travel time information to guided vehicles, each of which receives information under a different simulated ATIS information-giving strategy. Since different guidance strategies are being used, there is, for one iteration of the simulation procedure, one guided vehicle receiving guidance under each strategy. To 'gather' data for the purpose of guidance, probe observations over a window (of time) are needed; observations before this period are not used. The guided vehicles are in different links at different times, as they move from the origin to the destination, but whatever the simulated clock time of the guided vehicles' entry into the link, the probe observations used for guidance, are those for the travel time estimate updating interval preceding it. This leads to the concept of a sliding window, as illustrated in Figure 1[A]. In Figure 1[A], the expected location of a guided vehicle at time t' is k on route r. The size of the sliding window is $[t, t+\delta]$. Figure 1[B] gives a two-dimensional view of the sliding window over a route. The part $[t, t']$ is the simulated clock time interval for which we make use of simulated data for this guided vehicle. The part $[t', t+\delta]$ is the clock time interval for which we make predictions for that guided vehicle. The sliding window is essentially bounded in time and space by the time-location of the first probe and that of guided vehicle.

2.5 Factors considered in the Simulation

The purpose of the simulation was to compare static guidance routes, dynamic guidance routes and the omniscient driver route under different conditions. In this section we describe these different conditions.

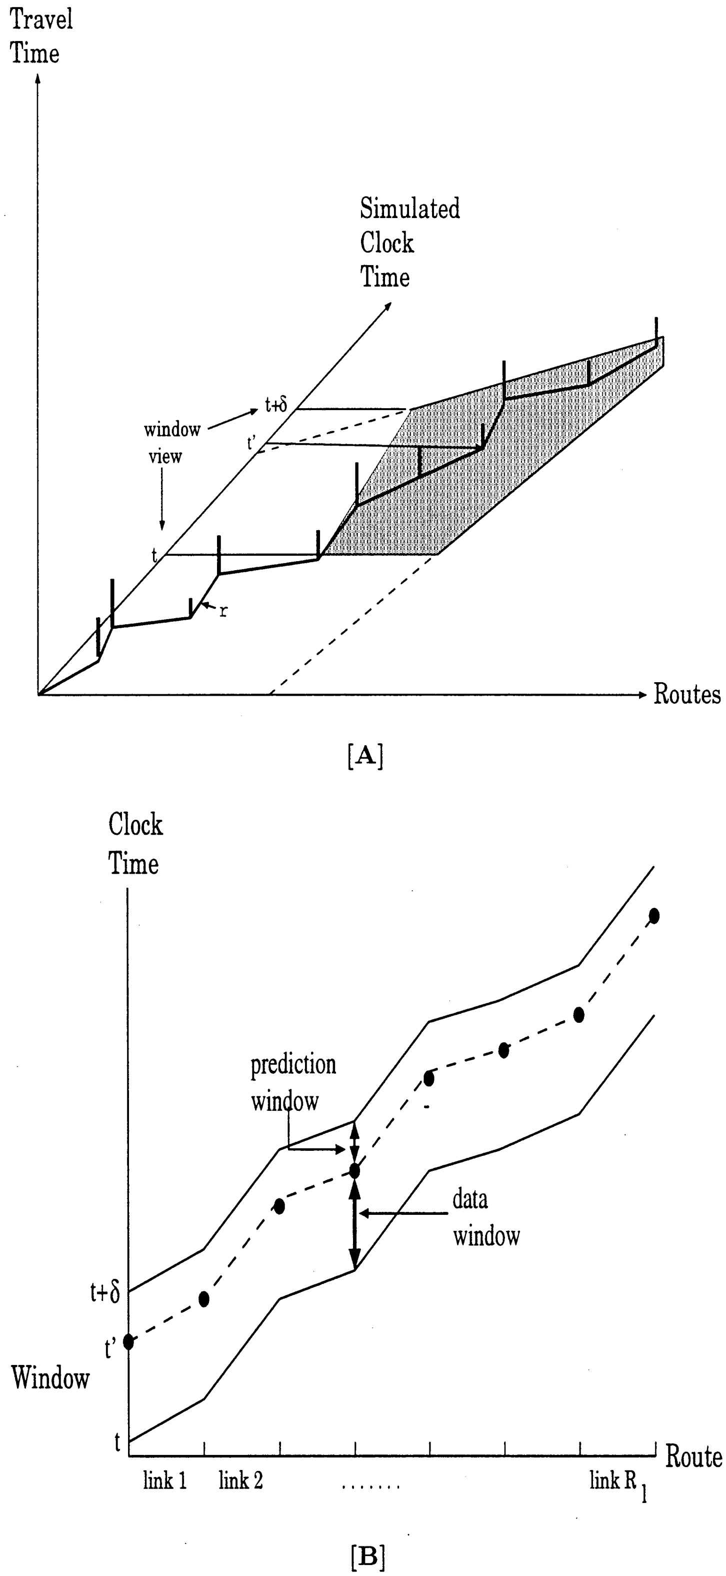

Figure 1: A sliding window view of the simulation along the route of the guided vehicle

Level of Congestion	Average v/c ratio	Minimum v/c	Maximum v/c
High	0.82	0	2.13
Medium	0.49	0	1.81

Table 1: Average volume-to-capacity ratio for each congestion level simulated for a sample of days and all types of links

They can be classified into 4 factors: two of them have 3 levels each, one has 2 levels and one 10 levels. The first two factors described below reflect conditions over which the designer of the ATIS has little control, while the last two factors consist of guidance strategies.

Factor 1: Congestion Level. The congestion level in the network can contribute substantially to the variability of network behavior. Therefore, we studied ATIS user benefits under each of the guidance strategies presented below in Section 2.5, for two levels of congestion — a medium level and a high level. Summary information on the link volume levels, on a random sample of 'days' that were simulated are presented in Table 1. The v/c ratios presented in the table are ratios of arrival flow rate per cycle to cycle capacity. The maximum values indicate the 'peak' points of this ratio within these sampled values and are not typical of conditions that prevailed for most simulated cycles. We did not consider low levels of congestion, because route guidance would not appear to be of much value when only recurrent congestion under free-flow conditions is considered.

Factor 2: Deployment Level. By level of deployment in this study, we mean the proportion of vehicles that are probes. Clearly, other things being the same, the higher the level of deployment the higher the number of probes on a given link over a given interval of time. Thus as deployment levels increase, the higher the sample sizes on which estimates are based and the lower the variance of the estimates. This is true for both static and dynamic estimates. We have considered 10 levels of deployment in which 1, 5, 10, 20, 30, 50, 75, 80, 85, 100 per cent of vehicles are assumed to be probes.

Factor 3: Forecast computation. Road conditions change with time. Issues relating to information-aging and route guidance have been presented in Watling and Van Vuren (1993). Therefore, one might expect that most recent observations would be most useful in providing dynamic guidance. On the other hand, if we confine ourselves to very recent observations, the sample size becomes small. An in-between approach might be to consider observations over a medium-sized interval but to weight the observations towards the end more highly than those towards the beginning. Of course, this too can have detrimental effects on the variance of the estimate — the greater the difference between earlier and later weights, the worse the variance.

We considered three methods of computing dynamic estimates, in each case using a 15 minute updating interval. In one we weighted all observations uniformly, i.e., we considered a simple mean. We call this case *uniformly weighted [UW]*. In the other two cases we partitioned the 15-minute updating interval I_t into four equal sub-intervals $I_{(t,1)}$, $I_{(t,2)}$, $I_{(t,3)}$ and $I_{(t,4)}$ where $I_{(t,4)}$ is the 'current' sub-interval. Let $\hat{t}_\ell(i_{(t,j)})$, $j = 1, \ldots, 4$ be the mean of probe reported travel times from each of the sub-intervals. A model of the form

$$\hat{t}^d_\ell(I_t) = \sum_{i=1}^{4} \omega_i t_\ell(I_{(t,i)})$$

was used to obtain the estimate $\hat{t}^d_\ell$ of the dynamic travel time on link ℓ during time I_t. Two sets of ω_i's were used and they are shown in Table 2. We call the two strategies *Time Lag factor 1 (TL1)* and *Time Lag factor 2 (TL2)*. Notice that for TL1, the ω_i's are steeply increasing. For TL2, the ω_i's are less steep. These latter sets of weights were actually obtained by least squares from some preliminary simulated data. Thus for this factor, we have the three levels: UW, TL1 and TL2.

Table 2: Values assigned to ω_1, ω_2, ω_3 and ω_4 under TL1 and TL2.

Weight	TL1	TL2
ω_4	0.90	0.49
ω_3	0.05	0.33
ω_2	0.03	0.14
ω_1	0.02	0.04

Factor 4: Message Throttling. In preliminary analyses of the results of the simulation model and earlier versions of it, it became very clear that dynamic travel time estimates had very large variances, especially at low deployment levels. These estimates resulted in routes which were inferior to routes resulting from static estimates. In an attempt to remedy this difficulty, we considered a strategy of providing dynamic estimates to vehicles only when they differed 'adequately' from the corresponding static estimate.

One way of viewing such a strategy is as follows: dynamic estimates would be provided only if we reject the hypothesis that the conditions are well represented by the static estimate. However, the concept dates to the early days of the design of the ADVANCE project, when the same strategy was proposed because it was believed that radio channel capacity would not allow the broadcast of all dynamic estimates. The term 'message throttling' was applied at that time — a word we have retained in this discussion. The ultimate implementation of ADVANCE does use message throttling, broadcasting a dynamic estimate only if it differs from the corresponding static estimate by more than 20 seconds — a number not too dissimilar from those that emerge when we use the relaxed criterion described below.

We used two criteria in our simulations:

- A *Relaxed Criterion [RC]* when we broadcast dynamic estimates if they differed from the corresponding static estimate by 1 'standard error.'
- A *Strict Criterion [SC]* when we required a difference of 2 'standard errors' for broadcast.

In each case the 'standard error' was obtained using the usual sample standard deviation of probe reports divided by $\sqrt{n}$, where n is the number of probe reports used in constructing a dynamic estimate. We have enclosed the term 'standard error' within quotes for the following two reasons:

1. Probe reports are most likely correlated and, therefore, the estimate is used would be an underestimate.
2. Even when we used weights within the estimate, we used the same criteria.

Thus, these criteria should be viewed as simply a relaxed and a strict criteria without any attempt to assign probabilities. Thus the message throttling factor has three levels:

1. A base case where all available dynamic estimates are broadcast,
2. An RC case
3. An SC case.

Dynamic and static routes were compared for each level of each factor and each combination of levels, [as will be described shortly]. Notice that except for deployment levels and congestion levels, other factors do not affect static routes. None of the factors, other than congestion levels affect the omniscient driver routes.

Notice also that different congestion levels required separate runs of the simulation models. However, the other factors did not require separate runs. Since travel times of all vehicles were known after each simulation, various deployment levels simply meant that we sampled different proportions of vehicles as probes. Different methods of computing dynamic estimates were just that: the estimates were computed when all probe travel times were known.

Each run of the model simulates one day and within each day, one driver is given guidance based on each strategy and for each deployment level. Hence, on each day, we have one driver who receives static guidance under each deployment level and one driver for each dynamic guidance strategy/deployment level combination. The number of repetitions of the simulation (thirty-five days) was determined more by computer time needed to generate the runs than any other reason.

3 Simulation Results

In this section, we examine the effects of the factors described in Section 2.5 on the quality of route guidance. In the discussion of each case, comparisons are repeatedly drawn with static guidance, which, as we pointed out earlier, is the same for all dynamic cases considered, and is affected only by congestion levels.

We present results in terms both of average travel times and proportion of times one strategy yields better routes than another. Let p be the proportion of times dynamic routes have lower travel times than static routes. Then

$$\begin{array}{c} H : \mathrm{E}[p] = .5 \\ \text{against the alternative} \\ A : \mathrm{E}[p] \neq .5 \end{array} \tag{2}$$

allows us to test the hypothesis that the quality of dynamic guidance is about the same for static guidance against the alternative that the quality is not the same. A test statistic that could be used for this test is

$$t = \frac{p - .5}{\sqrt{\frac{(.5)\times(.5)}{n}}} \tag{3}$$

and H would be rejected at the 5% level if

$$|\, p - .5 \,| > 1.96 \times (.5) \times n^{-\frac{1}{2}}. \tag{4}$$

Since $n = 35$, this yields $|\, p - .5 \,| > .17$ and therefore, the rejection of the hypothesis will occur whenever $p < .33$ or $p > .67$.

3.1 Basic Cases

The first case we discuss is the Basic Case where no message throttling occurs and dynamic times are computed by the UW strategy. In this section we confine ourselves to the high congestion situation, which is more interesting than the medium congestion situation considered in Section 3.5.

Figures 2(A) and 2(B) show the results from providing guidance during simulated periods of high congestion. The average Basic guidance route time shows that, at low levels of probe deployment, the driver receiving static guidance does better on the average than the driver receiving Basic dynamic information.

While dynamic estimates of travel times are likely to reflect current conditions better, they would typically have higher variances. This is especially true at lower levels of deployment. The issue of high variance is

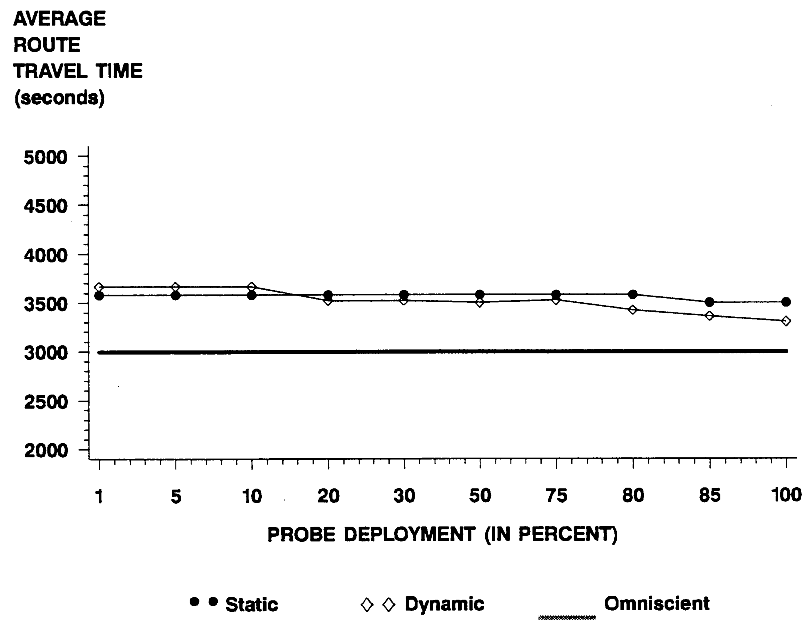

Figure 2: (A). Average route travel times for the Uniformaly Weighted Case.

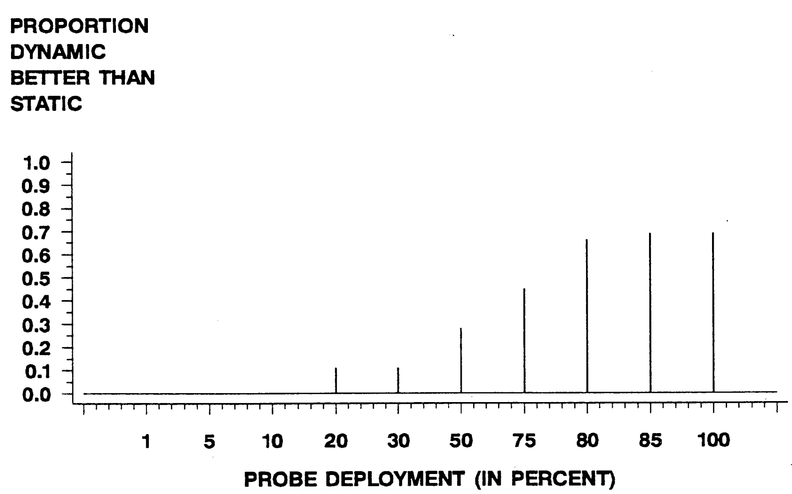

Figure 2: (B). Proportion of UW dynamic users with route travel times better than static.

critical in travel time estimation, since the conditions encountered by the guided vehicle can be substantially different from that predicted for it. Hence, not only can the route travel time be different from that predicted by the system, the best route may be also quite different. Therefore, in an ATIS design like the one we considered, there is a trade-off between the two important considerations of low variance versus information currentness.

The average dynamic evaluation route time improves with more probes. An interesting feature of the results is that even the static route improves with deployment level. With very few probes gathering static information, the period in which the simulated information for the static estimates was gathered (35 days) did not yield data on many links, and consequently retained the estimates obtained from the SUE model. Up to an 80% deployment level, static guidance leads the user, in fact, to the very same route.

The average travel time on the static route is 501.98 seconds greater than the true best route average time [the average travel time on routes taken by omniscient drivers].

These results offer some insight into driver behavior if the route based on static guidance is considered similar to the choice made by unguided but knowledgeable drivers. Drivers explore route possibilities in their travel and tend to "settle" on one route that they somehow favor. This is probably because the very short-term effects that affect link travel times are not known to them except by up-to-the-minute information to which they do not have access. This suggests the possibility of an upper limit to the quality of route improvement that static guidance offers.

Figure 2(B) presents the proportion of cases where Basic dynamic guidance give users lower cost routes than static guidance for high congestion levels. For the three lowest levels of deployment that were studied (1%, 5% and 10%), dynamic routes were always poorer than or tied with static routes. The proportion of times when dynamic routes were better than static routes start to increase with a 20% level of deployment. At 50% and 75% levels of deployment, the quality of guidance is about the same (that is, the proportion remains between .33 and .67, which are our benchmarks). At 80% and up, a plateau is reached where in 69% of the cases, dynamic guidance is clearly better.

We point out that at low levels of deployment of a probe-based ATIS, the use of historical data in dynamic route guidance is not just a useful strategy; it is a necessity. This is because there is not enough dynamic data to even compute a logical route, even if one were to totally ignore issues of giving good route guidance. In the simulation, most link travel time estimates that form the dynamic route estimate were, in fact, the default historical or static estimates. At low levels of deployment, the few Basic dynamic link travel time estimates for the relevant time period have much higher variance, compared to the static estimates, so as to lead the Basic dynamic guided driver to experience a route travel time that is worse than that experienced by a driver using static guidance.

3.2 Message Throttling

We now consider the Strict Criteria and Relaxed Criteria (see Figure 3(A), (B) and (C)) under the UW and a high congestion levels.

At low deployment levels (1% and 5%), RC and SC yield about the same average route time for the dynamic driver. Notice, however, that at these levels of probe deployment, both strategies are better than the average static route time (refer to Figure 2(A)) and much better than the average Basic dynamic route time. With more than 10% of the vehicles in the network sending dynamic information, RC shows greater average improvement than SC. This trend holds true till the highest level of deployment but the improvement RC gives dynamic drivers levels off after 50% deployment level. Overall, RC appears to be superior to SC.

From Figure 3(B), it can be seen that the proportion of cases in which dynamic guidance under SC is better than static, for all deployment levels, is within a rather narrow range of .43 to .71. Although users receiving guidance under deployment levels of 30% and higher do better than static users in terms of travel time (with proportions consistently over .67), this trend is seen intermittently at lower levels of deployment as well.

When an RC throttling strategy is used, the situation improves still further (see Figure 3(C)). In fact, with deployment levels of 50% and higher, RC yields better route times than static route guidance 91% of the time.

3.3 Forecasts Using Non-uniform Weights

In this section, we consider the case of non-uniform weighting schemes [TL1 and TL2] without message throttling, under high congestion. An immediate observation from Figure 4(A) for low levels of deployment, is that with both weighting schemes, dynamic routes based on the TL1 strategy are *far worse* than static routes as well as the corresponding UW routes. When 1% of the vehicles in the network transmit dynamic travel time information, static drivers spend 3577.21 seconds on the average on ATIS recommended routes, the driver receiving UW dynamic recommendation spends 3667.25 seconds on the average, dynamic drivers under TL2 spend 4067.33 seconds on the average, drivers under TL1 spend 4365.64 seconds, whereas the actual best route allows the omniscient driver to travel on an average of only 2997.54 seconds. This trend holds till 10% deployment level. At higher levels of deployment, the average dynamic route times under TL1 and TL2 are *far better* than the average times on static best route and UW dynamic best route for comparable deployment levels.

These trends occur due to the following reasons: recall that the updating interval is split into smaller sub-intervals. At low levels of deployment, the number of reports during each sub-interval from which the sub-interval dynamic link travel time estimate is formed is small and in some cases even zero. However, most of the total weight under both time lag factors is given to the estimate for the most recent sub-interval. Lack of observations within that sub-interval will cause the travel time estimate for the whole interval to have a fairly large variance.

Figures 4(B) and 4(C) give the proportion of drivers receiving dynamic route guidance based on TL1 and TL2 under high congestion situations, who experience lower travel times than static drivers. The proportions look about the same for both weighting schemes. The proportions also roughly increase with deployment level.

3.4 Non-uniform Weighting and Message Throttling

In this section, we present some examples of *mixed information-giving strategies* (i.e., various combinations of time lag factors and dynamic estimate screening approaches) for updating dynamic estimates of link travel times. Other possible combinations of cases have been considered in Thakuriah (1994). It needs to be kept in mind that in these cases weighting is performed *before* message throttling is applied.

Figures 5(A) and (B) show that dynamic route guidance with link estimates based on the TL2-RC strategy yield average travel times that are only about 55 seconds lower than the average static route times at the 1% and 5% levels. However, under intermediate (10%, 20% and 30%), levels of deployment, they provide dynamic users receiving information, route times (averaging about 3499.52 seconds) that are achieved by static guidance only at very high levels of deployment (85% and 100%). Considerable improvements occur after that and, at 3100.85 seconds at the 100% deployment level, TL2-RC yields the best average dynamic route times of all strategies considered and, in fact, the average TL2-RC time approaches the average travel times of the omniscient driver. Figure 5(A) also shows average route times for the TL2-SC strategy. Figure 5(B) shows that dynamic TL2-RC users are more likely to do better than static users at all deployment levels.

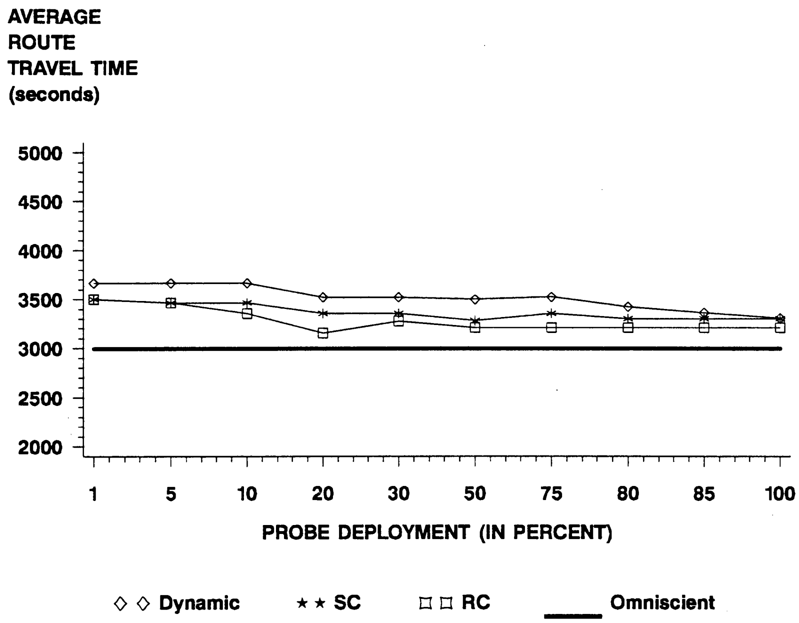

Figure 3: (A). Average route travel times under the Strict Criteria and Relaxed Criteria dynamic updating criteria.

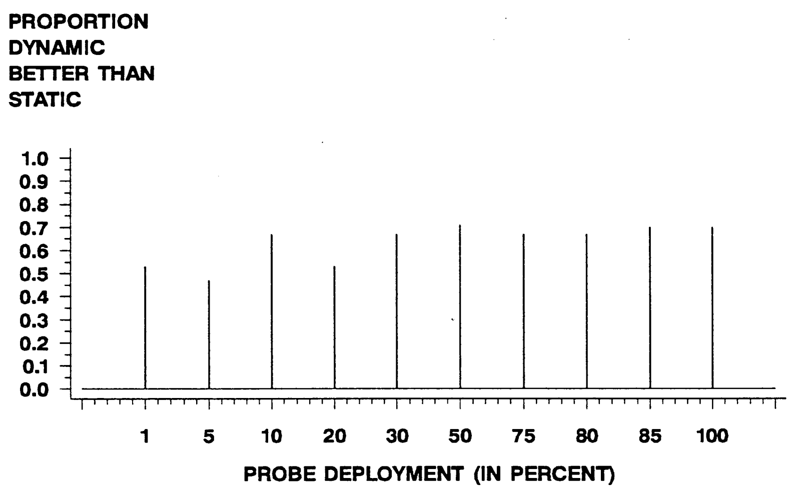

Figure 3: (B). Proportion of dynamic routes under Strict Criteria screening constraints that are better than autonomous routes.

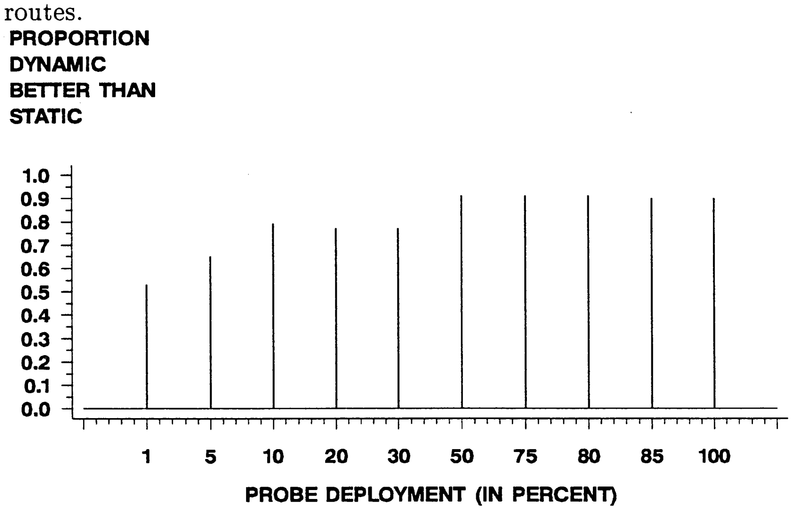

Figure 3: (C). Proportion of dynamic routes under Relaxed Criteria screening constraints that are better than autonomouss routes.

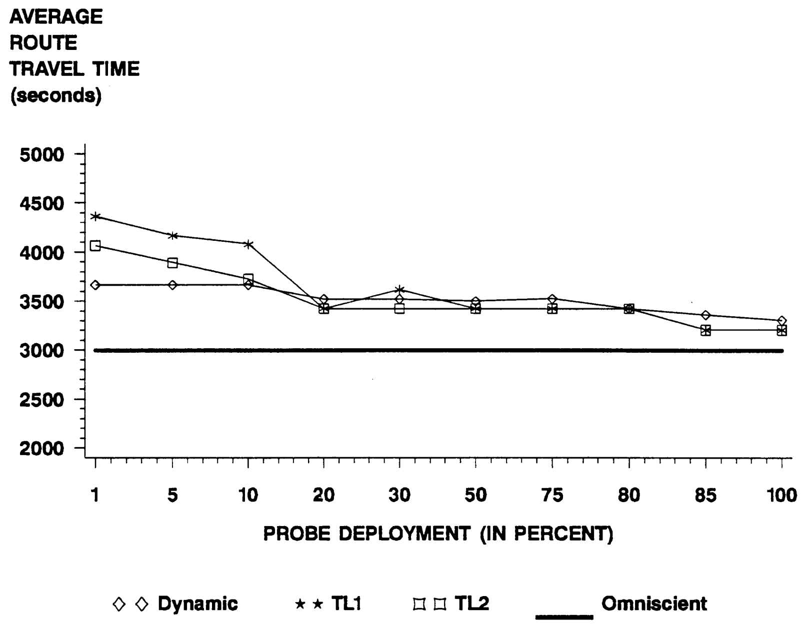

Figure 4: (A). Average route travel times under the Time Lag 1 and Time Lag 2 criteria.

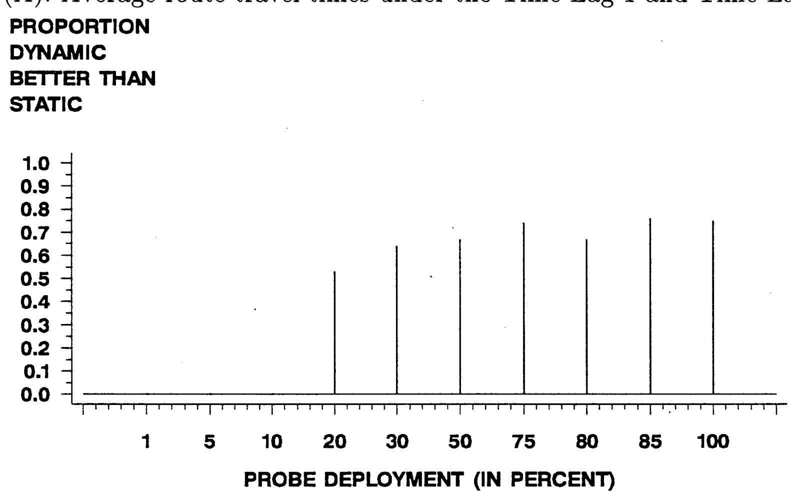

Figure 4: (B). Proportion of dynamic route travel times under the Time Lag 1 criteria that are better than autonomous route travel times.

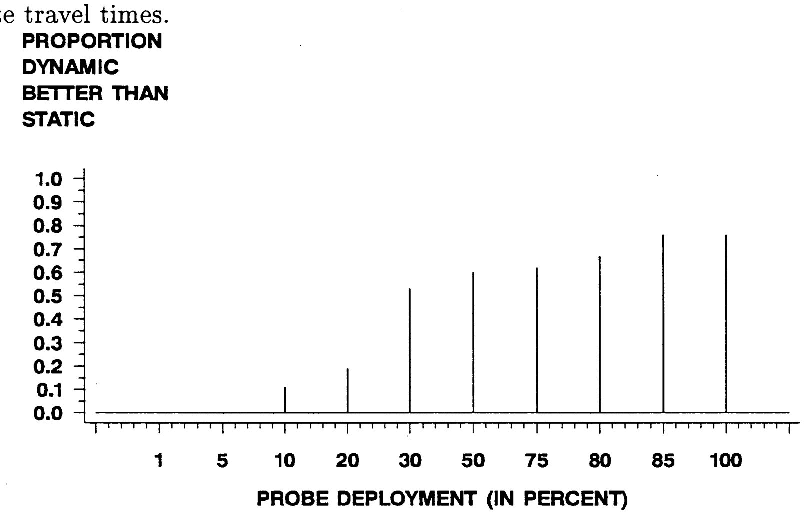

Figure 4: (C). Proportion of dynamic route travel times under the Time Lag 2 criteria that are better than autonomous route travel times.

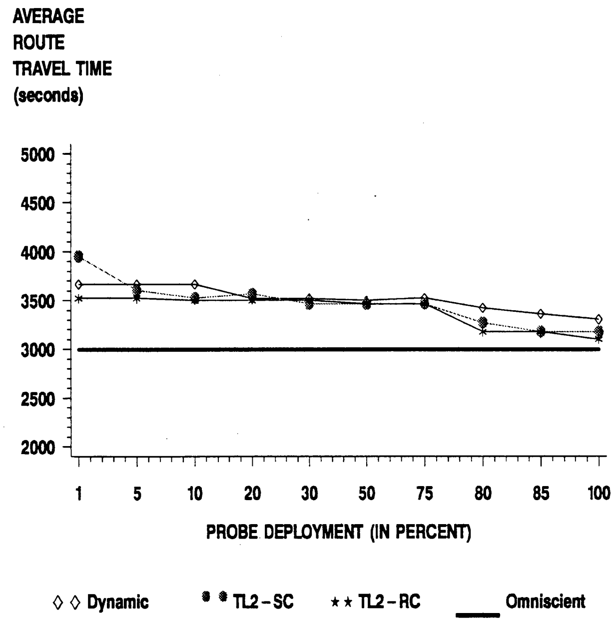

Figure 5: (A). Average route travel times for mixed cases.

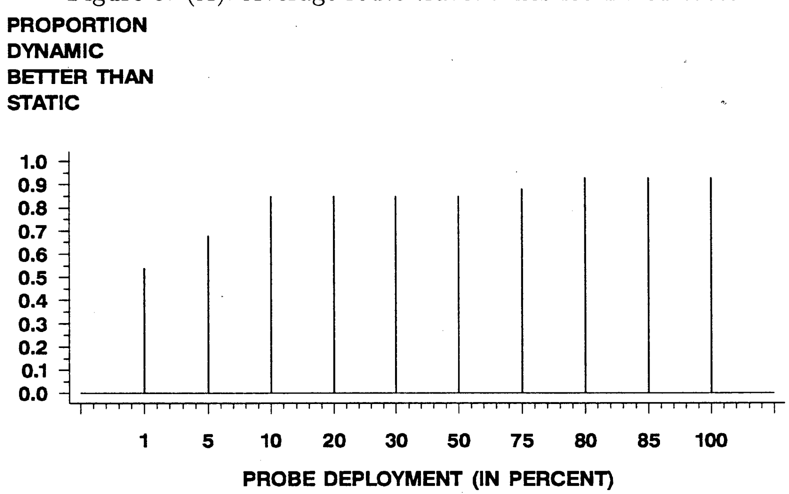

Figure 5: (B). Proportion of cases where dynamic route travel times under Time Lag 2 and Relaxed Criteria [TL2-RC] are better than autonomous route travel times.

3.5 Effect of Flow Levels

We now turn our attention to medium congestion levels. Figure 6(A) presents the results of the Relaxed Criteria for this case, in the context of other basic results. The Relaxed Criteria turns out to be the best possible dynamic route guidance strategy, even under medium flow levels and offers drivers receiving dynamic guidance an average of 125 seconds travel time savings over static drivers at high deployment levels.

Surprisingly, under medium congestion, the *best* (for all deployment levels) average Basic [i.e., with no message throttling and uniform weighting] dynamic route travel times are obtained under the lowest levels of deployment. The reason for this is that with medium levels of congestion and a small number of probes, most links in the dynamic route are actually based on static estimates of link travel times. With increased number of probes, a larger set of links receive dynamic updates. However, the greater number of probes are not enough to compensate for the variability in these reports. This effect is most pronounced at intermediate levels of deployment. The situation starts to improve at very high levels of deployment. When 85% and 100% of the vehicles are equipped, the difference between the average static route times and average dynamic route times is a matter of 44 seconds.

Another factor of interest was that with lack of congestion effects, there were two routes that dominated as best routes between the origin-destination pairs considered. All other routes were far inferior to these routes. Also, these two routes differed from one another by a matter of a few links.

The RC strategy (Figure 6(A) and (B)) helps to correct the disadvantages of overwriting of all static reports by simple averages of all dynamic travel time reports, irrespective of variability of reports, by acting as a "gatekeeper" against "unreliable" dynamic updates.

The effects of countering information-aging have a similar effect for medium and high congestion levels. The results are shown in Figure 6(A) and (C). At low levels of deployment, the screened average dynamic routes given under TL2 are far worse than the average static route and the Basic dynamic route. The same explanation that was given for the high congestion case applies here. With high weights given to the last sub-interval, the dynamic route that is estimated for the user is 'attracted' to links that receive no observations in the last interval because the estimated dynamic times are small for such links. However, the greater number of travel times that are observed with higher deployment levels, allow dynamic-TL2 to be a worthwhile information-giving strategy.

4 Conclusion

We first summarize some of the results of the study that we initially found surprising:

1. When one considers the stochasticity of link travel times, one gets a number of apparently counter-intuitive results. They are too numerous to mention here but have been described in the section above.

2. Many information providing strategies, that one might consider 'intuitively reasonable' and would perform well under semi-deterministic conditions, do very poorly when one takes information imperfection into account. For example, at low deployment levels, the variances of Basic dynamic travel time estimates [i.e, using simple means of travel times for all links] are so high as to have little value under recurrent conditions. Indeed, using this method of guidance often renders dynamic route-guided drivers worse off in terms of travel times than drivers receiving static information.

   During periods of recurrent moderate congestion, static guidance appears to be more effective than guidance given by the Basic dynamic strategy irrespective of deployment levels.

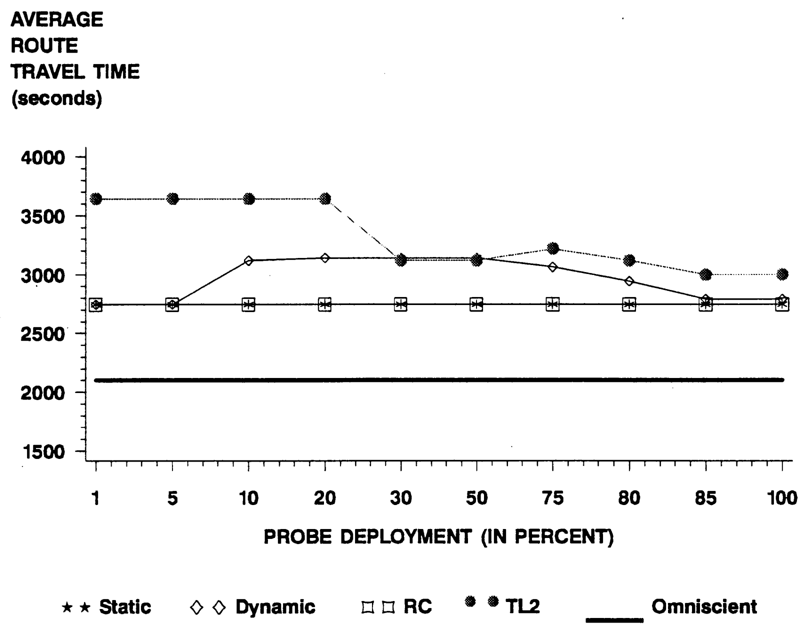

Figure 6: (A). Average route travel times during medium congestion levels.

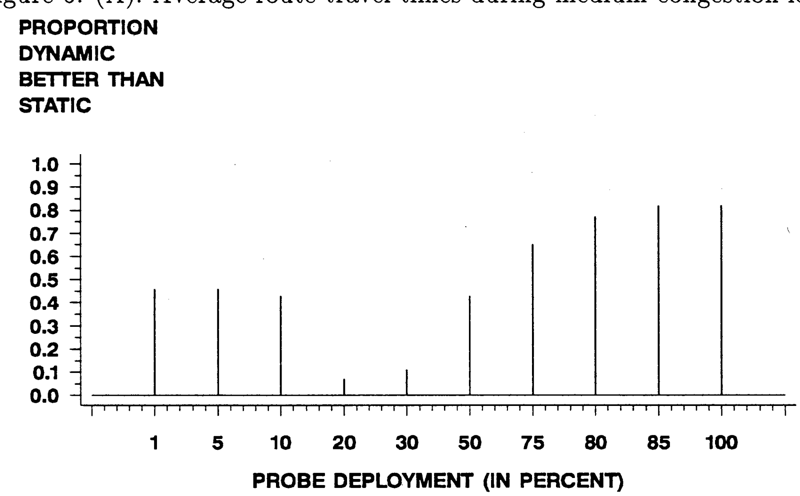

Figure 6: (B). Proportion of dynamic routes under RC screening constraints that are better than autonomous routes under medium levels of congestion.

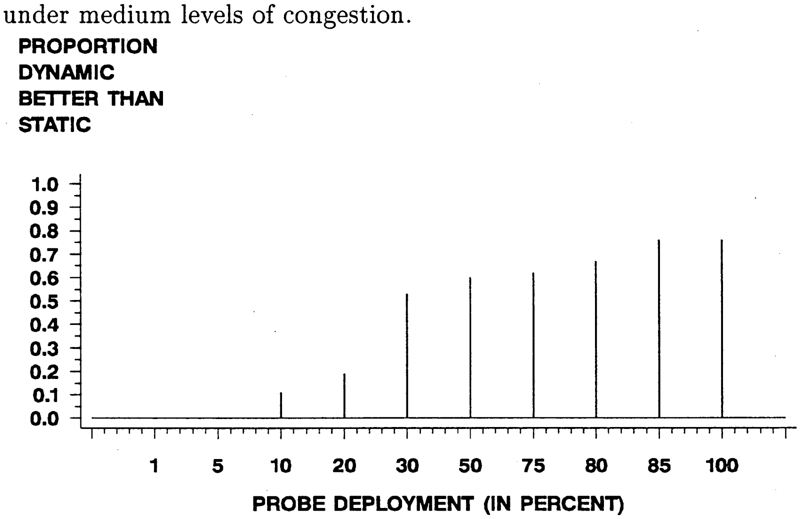

Figure 6: (C). Proportion of dynamic route travel times under the Time Lag 1 criteria that are better than autonomous route travel times during medium congestion.

3. However, the quality of dynamic routes can be improved by altering the guidance strategy employed. In fact, one of the strategies examined in this study — the TL2-RC strategy — behaves excellently under high congestion. Not only does it yield routes superior to routes based on static estimates, but it provides routes that, under high enough deployment levels, have travel times close to those taken by the omniscient driver [defined in Section 2.4.3 as the driver who knows the exact time he or she will actually encounter on any link no matter when he or she uses it].

   However, TL2-RC does not do nearly as well under moderate congestion although RC alone does provide routes that are better than static routes under some deployment levels. This is not entirely surprising, since if we identify the static guided driver as a highly experienced driver without guidance, all this means is that under moderate recurrent congestion, ATIS is not likely to be of enormous value to drivers knowledgeable about the area.

   We may summarize the above findings as follows:

   (a) Moderately screening outbound dynamic estimate strategies yield better route times than updating static estimates with dynamic information under all conditions.

   (b) At higher levels of congestion, placing emphasis on recent information by using empirically-derived weights within the same updating interval can lead to increased travel time savings as compared with the Basic dynamic case.

4. Dynamic updating information is likely to be available from a wider set of links during high congestion so that the number of route alternatives that are open for recommendation increases. Our results have indicated that the scope of dynamic route guidance was limited during medium congestion levels because of the competitiveness between a few dominantly good routes. Dynamic route guidance seem to be more effective where a larger number of reasonable route alternatives exist *and* the system is able to have information on those alternatives. This implies that route guidance may have different levels of effectiveness, depending on the type of network. Unless policy decisions are taken to 'spread' guided drivers among different routes, the presence of a dominantly superior route may defeat the purpose of route guidance.

5. Even static guidance can show improvement in route quality with greater deployment levels. However, *an upper limit seems to exist on the improvement offered by static guidance.*

From the remarks above it follows immediately that evaluation of ATIS strategies and decision-making regarding them must be performed bearing in mind the stochasticity of travel times and flows. Ignoring these factors can easily lead to serious mistakes.

A second conclusion, also important, may not be as apparent. This concerns the importance of static estimates, a subject that has not as of yet been given much prominence in the literature. The results presented in Section 2.5 indicate that static guidance can provide reasonable guidance to ATIS users under recurrent congestion. However, static estimates of travel time play a much more important role than simply providing static guidance because the quality of dynamic route guidance upto fairly high deployment levels depend critically on static estimates of travel times in a probe-based system. Indeed, the strategies that emerged as the best in our study [TL2-RC and RC] require static estimates *at all deployment levels.*

A final conclusion is that under high congestion levels, even for drivers quite familiar with the area and even under only recurrent congestion, ATIS has tremendous potential for time saving if information dispensing strategies are constructed carefully.

Acknowledgements — The authors would like to thank the ADVANCE project. This work was partially supported by the National Science Foundation's Division of Mathematical Sciences Grant Numbers 9313013 and 9208758.

References

1. Boyce, D. E., A. M. Kirson and J. L. Schofer (1994) ADVANCE — The Illinois Navigation and Route Guidance Demonstration Program. In *Advanced Technology for Road Transport*, Ian Catling (ed.), Artech House, Boston.

2. Center for Transportation Research, University of Texas at Austin (1994). *Development and Testing of Dynamic Traffic Assignment and Simulation Procedures for ATIS/ATMS Applications.* Technical Report DTFH61-90-R-00074-FG. Prepared for U.S. Department of Transportation, Federal Highway Administration, Office of Safety and Traffic Operations, Research and Development, 6300 Georgetown Pike, McLean, VA 22106. June 1993, Revised June 1994. pp. 75-76.

3. Daganzo, C. F. and Sheffi, Y. (1977) On Stochastic Models of Traffic Assignment. In *Transportation Science*, **11**(3).

4. Hicks, J. E., Boyce, D. E. and Sen, A. (1992) Static Network Equilibrium Models and Analyses for the Design of Dynamic Route Guidance Systems. A Technical Report in support of the design phase of the ADVANCE project.

5. JHK & Associates (1993) PATHFINDER Evaluation Report. Prepared for California Department of Transportation, District 7. February.

6. Koutsopoulos, H.N. and H. Xin (1993) An Information Discounting Routing Strategy for Advanced Traveler Information Systems. In *Transportation Research C*, **3**, 249-264.

7. Olszewski, P. (1990) Modeling of queue probability distribution at traffic signals. In *Transportation and Traffic Theory*, ed. Koshi, M., Elsevier, New York.

8. Powell, W. B. and Sheffi, Y. (1982) The Convergence of Equilibrium Algorithms with Predetermined Step Sizes. In *Transportation Science*, 16:1.

9. Sen, A., and Thakuriah, P. (1995) Estimation of Static Travel Times in a Dynamic Route Guidance System. In Network, Control, Communication and Computing Technologies for Intelligent Transportation Systems; a special issue of *Mathematical and Computer Modeling*, Pergamon Press, New York, 83-102.

10. Sen, A. and T. E. Smith (1995) *Gravity Models of Spatial Interaction Behavior.* Springer-Verlag, Heidelberg.

11. Sheffi, Y. (1985) *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods.* Prentice-Hall, Inc., Englewood Cliffs, New Jersey.

12. Smith, T. E. (1987) Poisson Gravity Models of Spatial Flows. *Journal of Regional Science*, **27**(3).

13. Tarko, A. and Rouphail, N. M. (1994) Traffic Data Integration in Advanced Traveler Information Systems using Fuzzy Operator Logic. A Technical Report Motivated by the ADVANCE Project to the Illinois Department of Transportation. Urban Transportation Center, University of Illinois of Chicago.

14. Thakuriah, P. (1994) Evaluation of Route Guidance in Advanced Traveler Information Systems. Ph.D. Dissertation, University of Illinois at Chicago, Fall.

15. Thakuriah, P., Li, J. and Sen, A. (1992) Development of a Short-Term Travel Time Prediction System for a Dynamic Route Guidance System. Paper presented in the *39th Regional Science Association International Annual Conference*, Chicago, November.

16. Watling, D. and Van Vuren, T. (1993) The Modeling of Dynamic Route Guidance Systems. In *Transportation Research C*, **1**(2), 159-182.