Chapter 4
SIMULATION EXPERIMENTS

This chapter discusses the numerical experiments performed to study the performance of representative networks under information and provides the conclusions and insights derived from them along with other results that verify the validity of the program components. The chapter begins with a discussion of the kind of information supply strategy studied, in section 4.1. The experimental designs are explained along with the different networks and information scenarios studied, in section 4.2. Section 4.3 provides the results on system performance, and section 4.4 provides results verifying the model integrity as well as some computational results.

4.1. INFORMATION SUPPLY STRATEGY

Traffic network information supply strategies can be largely of two kinds: descriptive and prescriptive. The former refers to the case where the drivers receive the information on network conditions (route trip times, for instance) and make the decisions themselves on which route to follow. The latter refers to the case when the information provided is in the form of guidance advice, such as the "best" path available (the guidance-control center selecting the "best" path based on certain criteria, for instance). Thus mandatory vehicle routing in a network is a case of prescriptive information supply. However, even when the information provided is in the form of *the "best" path(s) to follow*, if the drivers are allowed to make their on decisions, we consider it descriptive information supply.

The kind of information supply strategy studied here is largely descriptive in nature, in that drivers are provided with information on current trip times, which they use in selecting their routes, which may not always be the "best" routes displayed by the information system. The various simulations performed are for different levels of driver route switching propensity and various levels of market penetration, thus making it a study of alternative information scenarios, and not exactly of alternative information supply schemes. However, in some of the cases studied, the driver behavior assumed is such that the best available route is always selected, which would result in the same system performance as in the case of compulsory vehicle routing with prescriptive information about the best route. Thus, the cases with a value of zero for the route switch threshold fraction (η in eqn 2.11, section 2.3.3) effectively simulates a particular type of prescriptive information supply strategy.

Another important aspect of the information supply strategy here is that the provided information is always based on *current* conditions in the network, as opposed to expected or predicted conditions. This means that the route trip times used to make the route decisions are calculated with the existing link travel times and queue lengths. No attempt is made to carry out simulations with a trip time prediction algorithm (which is mainly due to the unavailability of comprehensive prediction techniques in the literature). This is not unrealistic, because many, if not most, of the existing or proposed traffic information systems are not expected to have reliable prediction algorithms in the near future.

4.2. EXPERIMENTAL DESIGN

This section discusses the design of the simulation experiments carried out. The different networks simulated, the different parameters values used for the driver decision, the different levels of information supply and the different traffic loading patterns are explained in that order. One base-case of no information supply is modelled along with 25 other cases (five different behavioral scenarios, each for five different information levels) for each network type and for each loading pattern.

4.2.1. Network Types

Two types of networks are simulated: (1) the three-highway, single destination corridor network and (2) the general network based on the core area of Austin, Texas.

4.2.4.1. Corridor network

The corridor network selected for the experiments is shown in figure 4.1. There are three uni-directional highways, each of 9 mile length leading to a single destination, equivalent to a central business district (CBD). The highways are of three different free-flow speeds: 55, 45 and 35 mph, thus representing a freeway, a fast arterial street and a slower arterial street. The highways are considered to be parallel to each other with half a mile distance from the middle highway (faster arterial) to each of the other two. Each highway is divided into 9 segments, each of one mile length. The segments within a distance of three miles from the CBD are assumed not to receive any traffic generation, while the farthest 6 miles receive traffic generation. The traffic generation profiles assumed are discussed in section 4.2.4.

Cross-over facilities are assumed to exist between these highways at certain points. These are at distances 3, 4, 5 and 6 miles from the CBD. There are separate and independent cross-over streets to go from each highway to each other

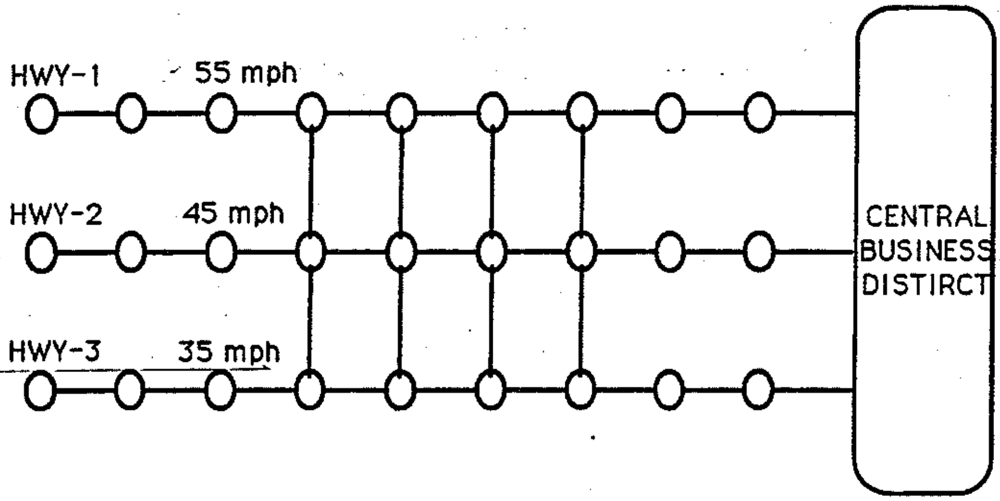

Fig. 4.1. The Three-highway Corridor

highway (thus forming 6 cross-over streets between the three highways at each of the 4 cross-over points). The cross-over streets between the freeway and the slower arterial are of one mile length and the others are of half a mile length. The drivers equipped for information are assumed to make route switch decisions based on the trip time information provided on each route, whenever they reach the cross-over points. The implicit assumption here is that only the three independent and non-overlapping highways are considered by the drivers while switching routes, meaning that complex paths with back and forth switching are not considered to be in their choice set. Of course, the drivers who do not receive information, do not switch highways and continue on the highway that they start their trip on.

4.2.1.2. Austin core general network

The network includes all the streets in the core region of Austin, Texas. The region is bounded by the Mo-Pac freeway on the west, the Interstate Highway 35 on the east, the Colorado river on the south and the 26th Street on the north. This region includes the downtown area, the University of Texas and the largely residential area to the west of Lamar street. There are 28 traffic demand zones in this area, as defined by the Planning division of the City of Austin. Due to the high density of businesses and offices in the downtown area and the presence of the University of Texas, parts of this region experience the worst traffic congestion in Austin during the peak periods. Traffic demand from outside this region and the inter-zonal demand within this region are used for the simulations as explained in section 4.2.4.2.

Due to the lack of easy availability of a reliable database of all the street characteristics in this network, the number of lanes assumed in the numerical illustration may not correspond to the actual Austin situation. As such the experiments are intended not to obtain conclusions specific to Austin, but rather to illustrate the capabilities of the simulation-assignment methodology. While the network is large enough and has many characteristics which were desirable to demonstrate the capabilities of the simulation framework in modelling realistic large networks, the lack of high-speed highways *through* the network has to be considered a disadvantage in studying the effects of route-switching under information. For this reasons some of the highways are assumed to be high-speed (55 mph) 3-lane arterials (Lamar street in the North-South direction and MLK street and Enfield street in the East-West direction).

The area studied is shown in figure 4.2 along with 7 sections of it (A to G). The network details (node numbers and connectivity) of these seven sections are given in figures 4.3-4.9. These figures show the extensions assumed for the two freeways, Mo-Pac and I-35, which are necessary to load the heavy traffic from outside the study area without entry queue congestion (see section 4.2.4.2. for more details on the demand pattern). All the arcs representing the city streets are considered straight but their lengths are the actual lengths.

No signalized intersections are explicitly modelled in the network, as the framework does not simulate signals (this is a future capability that is currently being developed). Under sufficiently congested conditions this aspect is not expected to significantly affect the results and conclusions on route-switching and the consequent system performance under information, as the model does simulate queuing at the nodes and thus capture the additional delay at the network intersections.

Fig. 4.2 The Austin core network studied (shown within the broken line). The sections A to G shown here are shown in detail in figures 4.3 to 4.9. The four residential zone-clusters around the study area are also shown

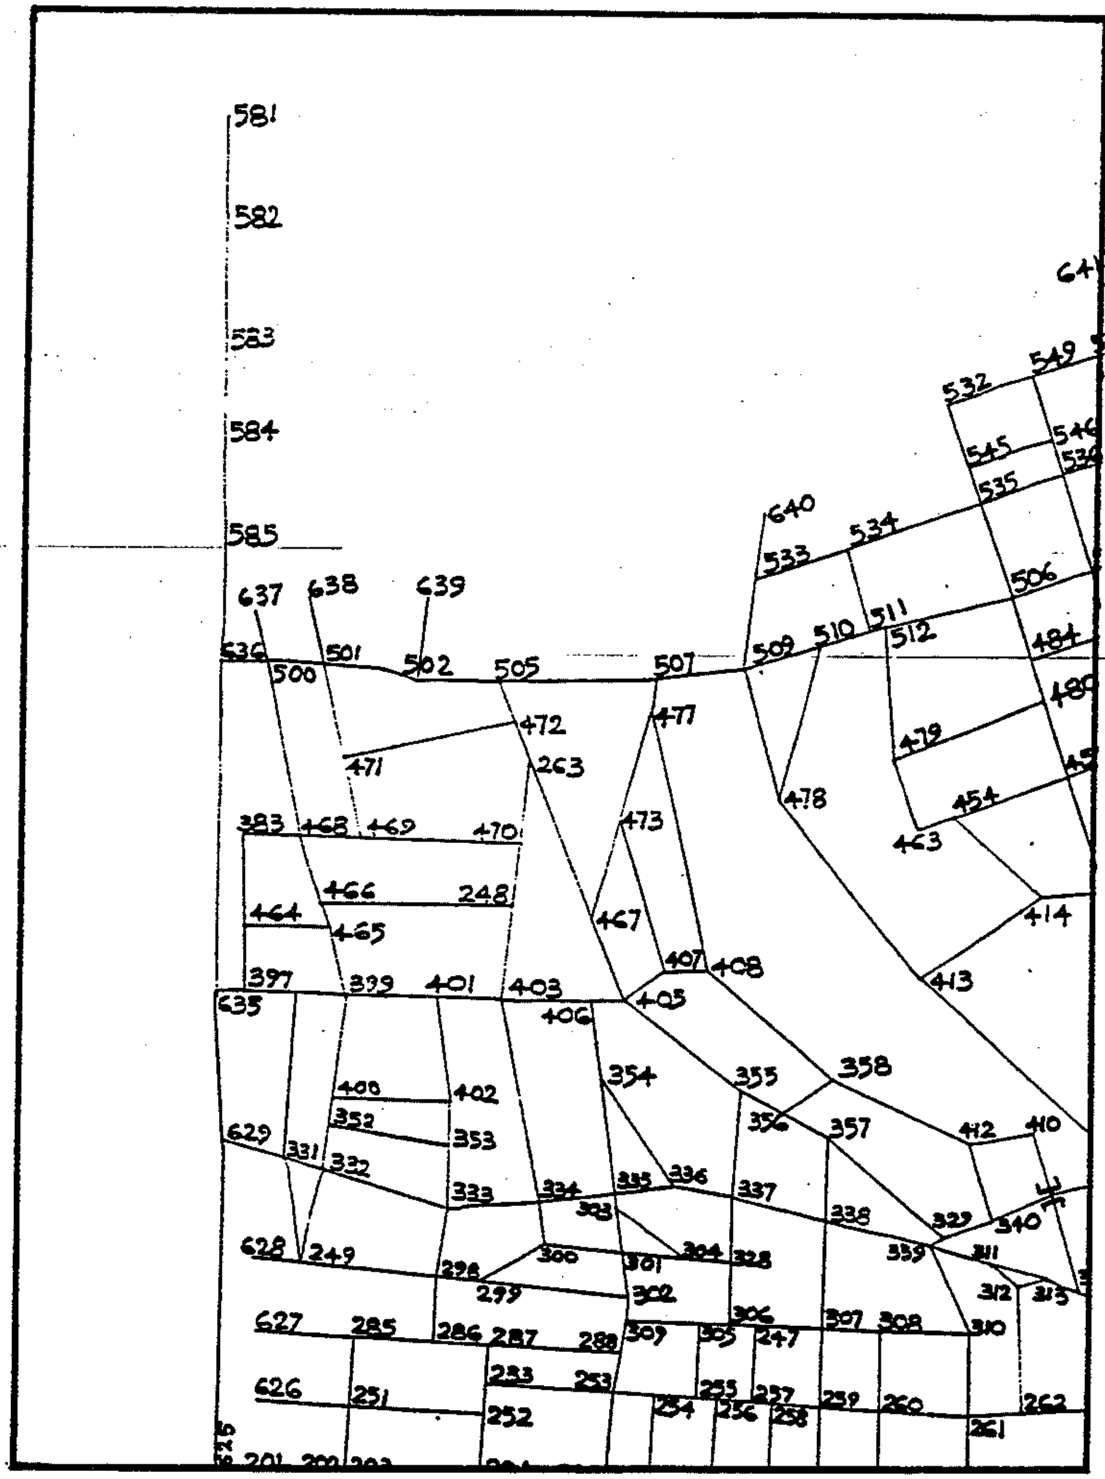

Fig. 4.3. Austin Network (portion A)

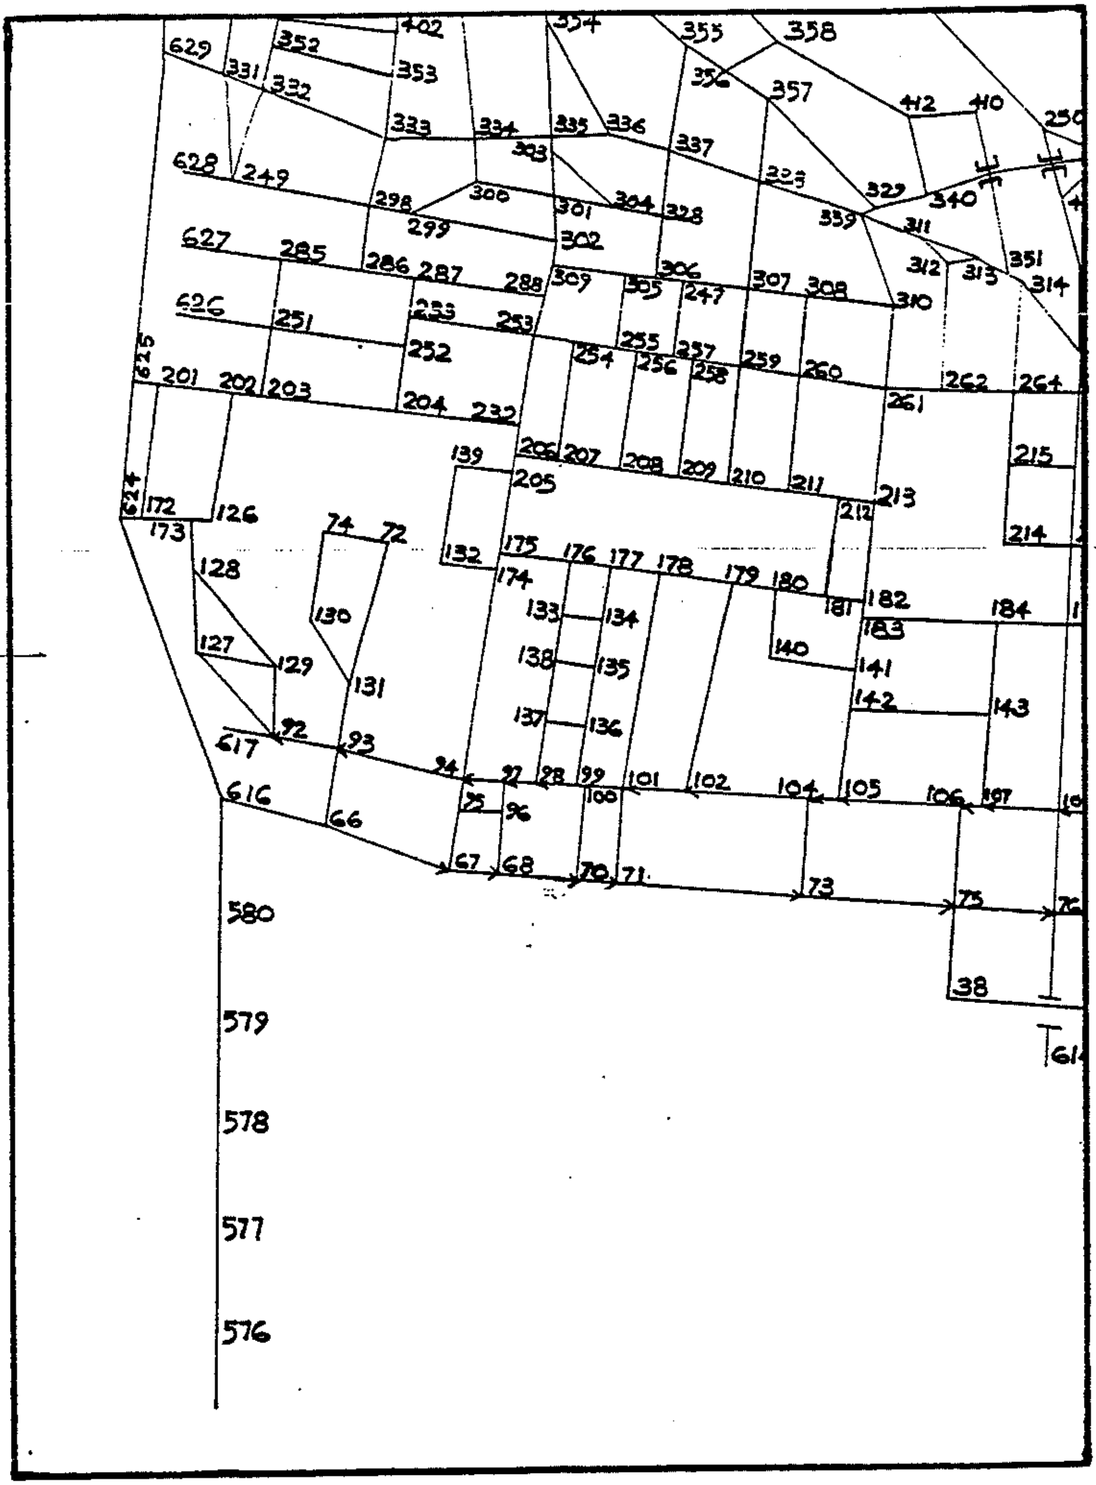

Fig. 4.4. Austin Network (portion B)

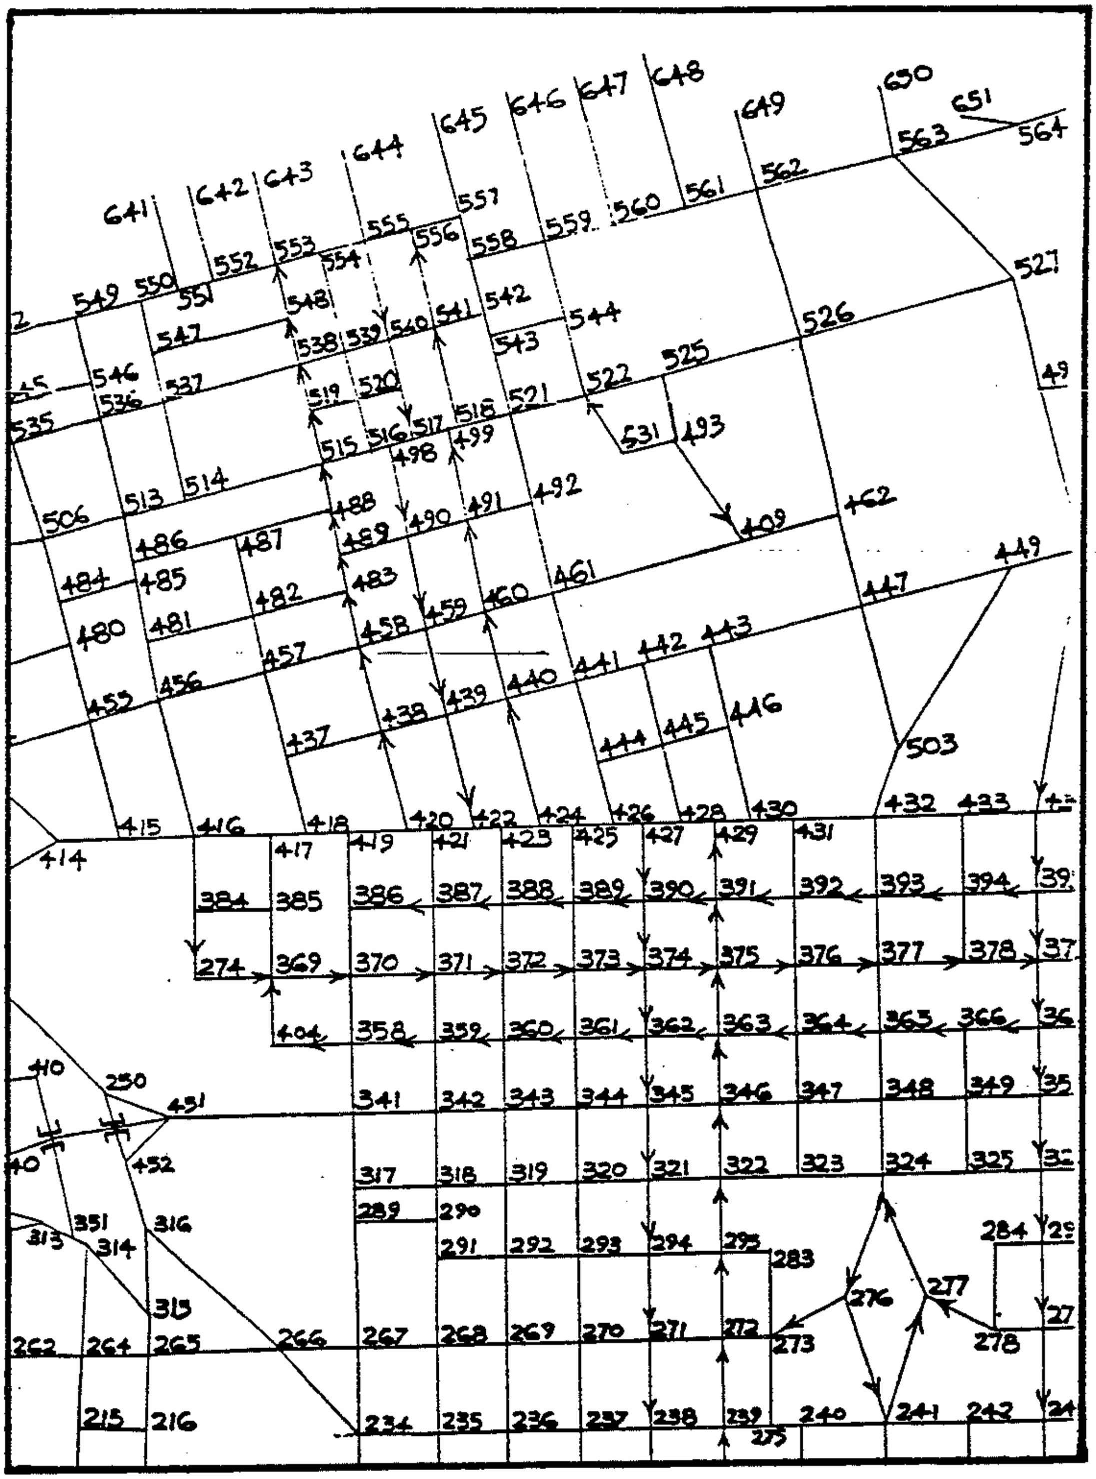

Fig. 4.5. Austin Network (portion C)

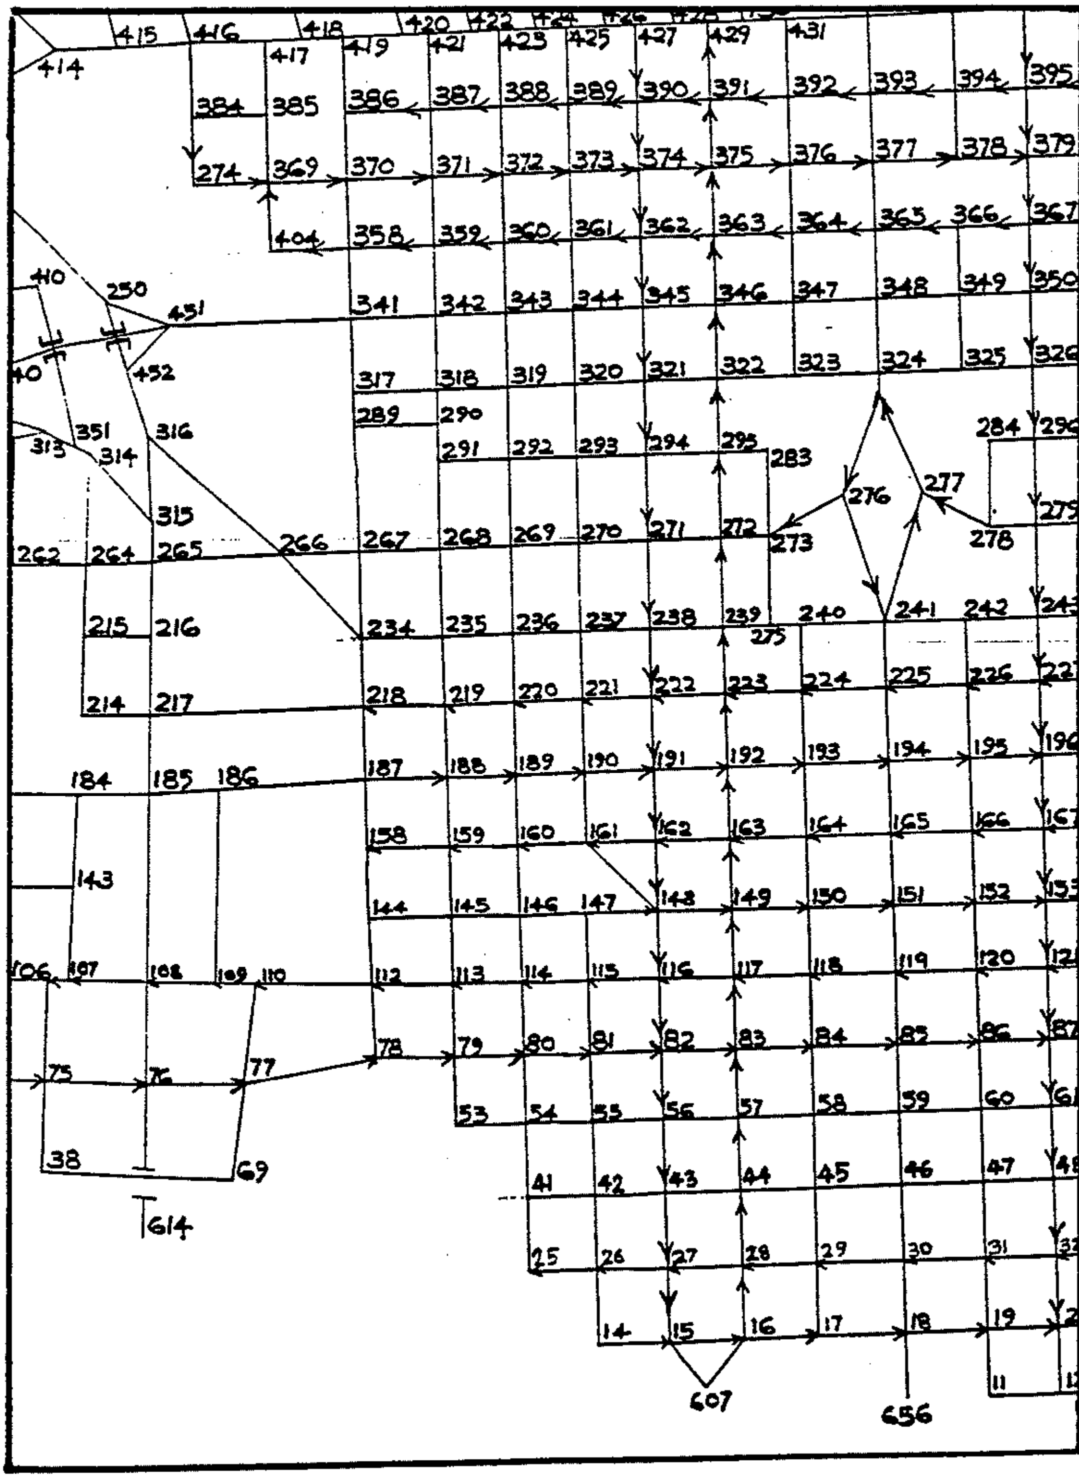

Fig. 4.6. Austin Network (portion D)

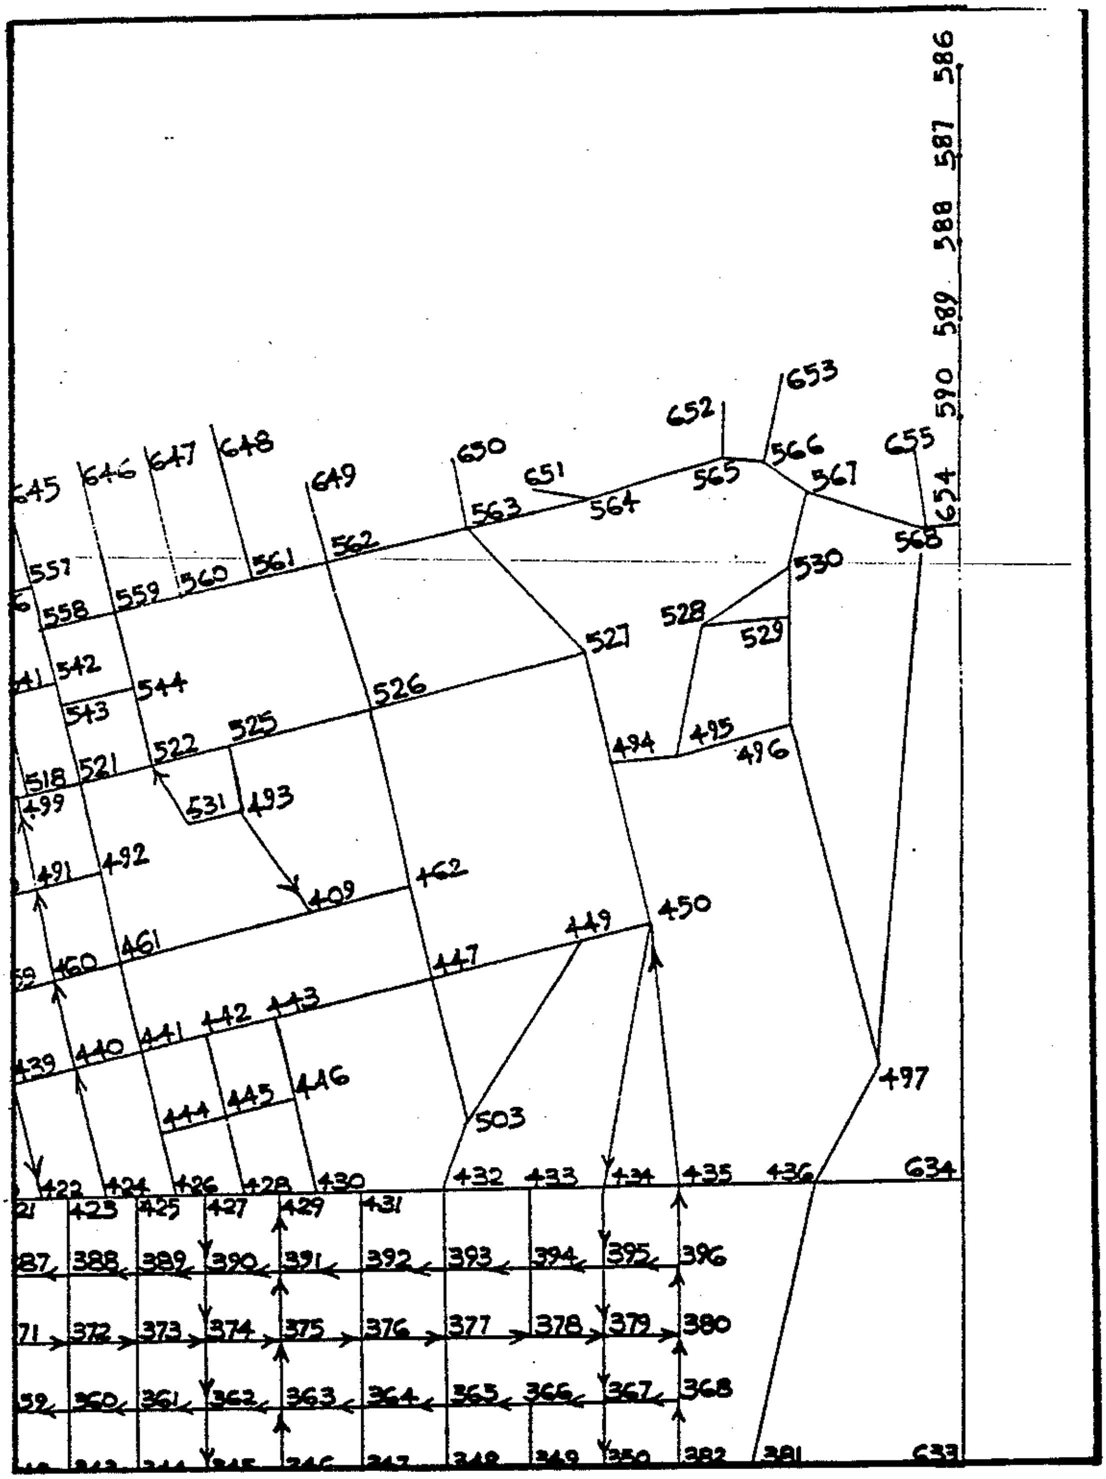

Fig. 4.7. Austin Network (portion E)

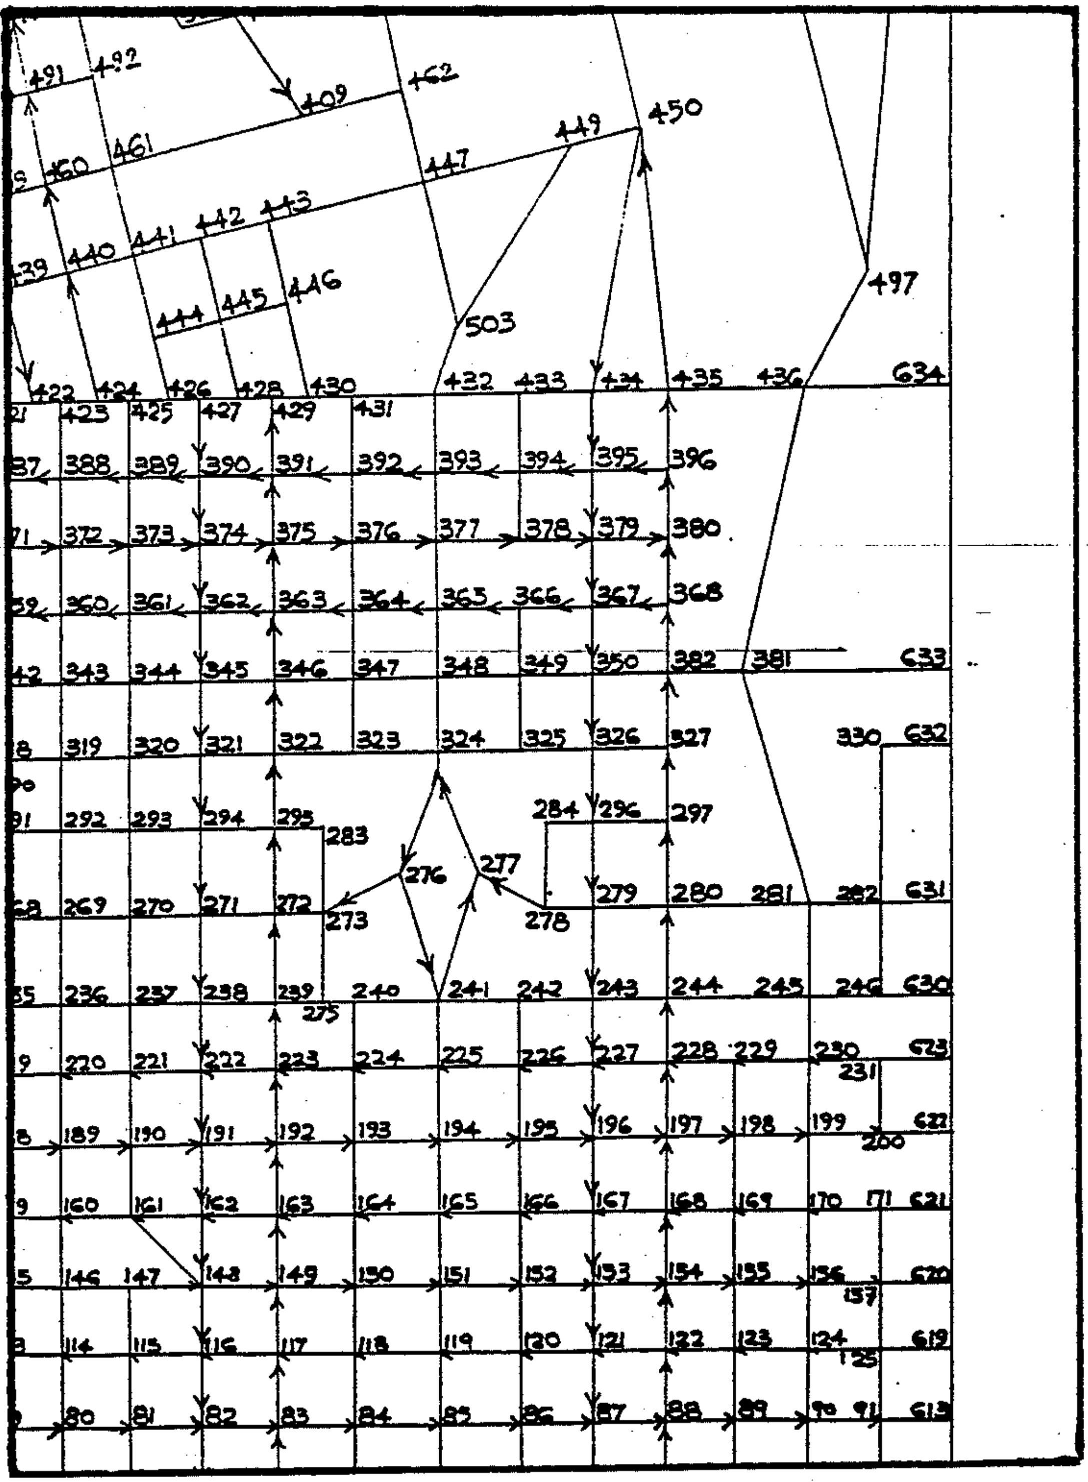

Fig. 4.8. Austin Network (portion F)

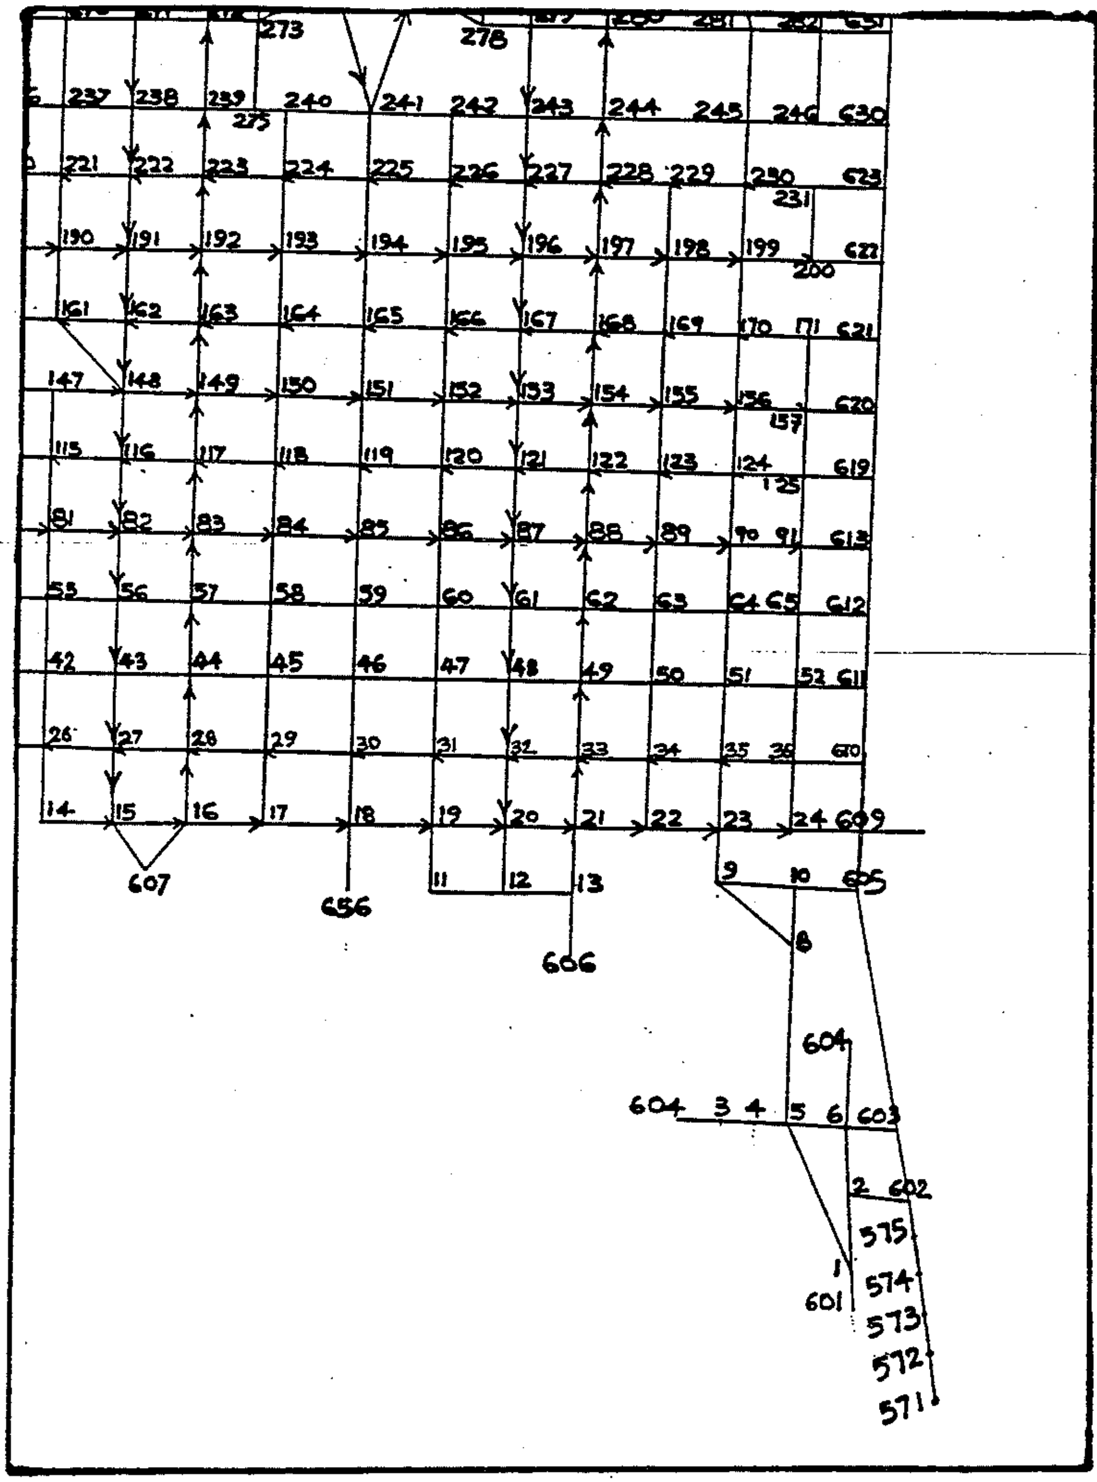

Fig. 4.9. Austin Network (portion G)

4.2.2. **Behavioral model parameters**

The behavioral model parameter that is changed for the different simulations is the mean relative route switching indifference threshold fraction, $\bar{\eta}$ which is the mean of the η_j values (see route switching model explained in section 2.3.3) of all the drivers. Each of the above-mentioned networks are simulated for five different values of this parameter: 0, 0.1, 0.2, 0.3, and 0.5. This covers the spectrum of reasonable driver responses, from the case of myopic route switching when $\bar{\eta} = 0$, to the case of conservative route-switching when $\bar{\eta} = 0.5$. The values above 0.5 were not simulated as they do not appear to be realistic, in that there is very little chance that a driver will not switch to a route that provides 50 % advantage in travel time over the current route. The η_j values are assumed to have a triangular distribution over the driver population with a range between $0.75\bar{\eta}$ and $1.25\bar{\eta}$ for all the simulation cases.

The value assumed in each simulation for the other parameter of the behavioral model, which is the absolute minimum travel time advantage τ (see section 2.3.3) that the drivers seek, depends on the network case and also the switching propensity. For all the simulations with myopic route switching (i.e., $\mu = 0.0$), a value of 0.0 was assumed for τ, to represent the extreme case of perfectly myopic route switching which is equivalent to the case with a compulsory best path routing strategy. For all the other simulation runs, the value of τ depends on the network. For the corridor network simulations, a value of one minute was assumed for this parameter, while for the Austin network simulations, a value of 15 seconds was assumed. In an actual corridor network of the type studied, there are fewer switching opportunities, and the drivers generally understand that they do not have many opportunities to switch back if the alternative route turns out to be worse than expected, especially because the last three miles do not offer any cross-overs. Their trip distances are relatively larger too, as all of them travel for at least 3.5 miles. On the other hand, for the general network, the drivers know that they can make switches till the last few blocks near their destination and may do so especially on shorter trips between zones which are close by. Thus a lower τ value is reasonable for the general network simulations. However, observations of actual behavior are necessary in order to calibrate this and other behavioral parameters.

4.2.3. Level of information

Six different information levels are simulated for each network, for each of the five behavioral scenarios and under each loading pattern. This is accomplished by assuming values of 0.0, 0.1, 0.25, 0.50, 0.75 and 1.00 for the fraction of the driver population that is equipped for information. Of course, the case of zero fraction of the drivers receiving information is not repeated for different behavioral scenarios as all the drivers continue on the initially assigned route throughout their journeys. More levels of the fraction of the drivers receiving information are simulated at the low end compared to the high end, because sharper variation in the system performance are expected at the lower levels of information supply, which is confirmed by the results (see section 4.3).

4.2.4. Traffic loading patterns

The following two sections explain the traffic loading patterns assumed for the two networks. Two loading patterns are considered for the corridor simulations; an arbitrary uniform departure pattern and a dynamic stochastic user equilibrium pattern. The loading patterns used for the general network are obtained by modifying the interzonal daily traffic demand for Austin zones in 1985, the modifications being done to achieve O-D demands with time-dependent variation during the peak period.

4.2.4.1. Corridor traffic loading

The area surrounding the three highways in the corridor (see fig. 4.1) is subdivided into one-mile sectors, which contain the corresponding parallel links of the three facilities. Only sectors 1 through 6 (with sector 1 denoting the most distant from the CBD) are assumed to be residential sectors which generate commuting traffic. The traffic loading is the result of the time dependent driver departure functions on each of the three highways. No vehicle generation is assumed onto the crossover links.

Two different loading patterns are considered in these experiments. The first is an arbitrary departure pattern simply referred to as loading pattern-1. Under this pattern, commuters in each sector split equally among the three highways and depart uniformly over a 20-minute period, at a rate of 26.67 vehicles per minute for each facility. The loading periods for each sector are staggered with a time lag of five minutes between adjacent sectors with sector one starting first, which effectively models the earlier departures from sectors farther away from the CBD. See figure 4.10 for the cumulative departure pattern for two sectors (sectors 2 and 5), at 7 and 4 miles respectively from the CBD. All three highways receive traffic generation at the same rate.

The second loading pattern satisfies the stochastic version of the dynamic

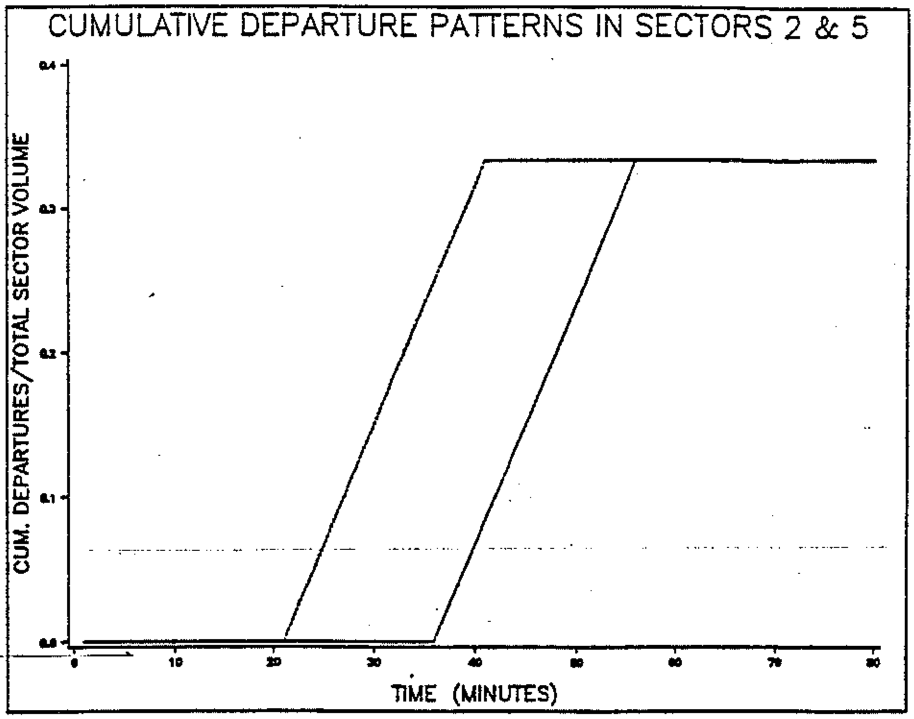

Fig. 4.10. Cumulative departures under Loading pattern-1 for the three highways in sectors 7 and 4 miles away from the CBD. Simulation case: Route-switch threshold fraction = 0.2, equipped for information = 0.5)

user equilibrium (DUE) conditions, which means that no driver can unilaterally improve his perceived (random) utility by unilaterally switching either departure time or route. The departure pattern is obtained by iterative simulations till an equilibrium state is reached under a particular utility function, starting from utilities under free-flow conditions. To explain simply, these iterations are needed to achieve consistence between the departure pattern (itself based on the utilities that drivers receive for a given system performance) and the system performance (which results from the particular departure pattern). Chapter 5 provides a review of dynamic user equilibrium concepts and for further details on the iterative equilibration technique and the calibrated utility function assumed for it. The departure pattern here is obtained with the same utility function as the function U1 explained in section 5.3. One important aspect of this departure pattern is that it is distributed over the whole 80 minute period. This is a direct result of the utility function assumed and the randomness of the perceived utility which causes a certain probability for departures at all the departure time and route alternatives (again, see chapter 5 for details of the definition of departure alternatives etc.). Figures 4.11 and 4.12 show the cumulative departure patterns for two representative sectors, at 7 and 4 miles respectively. The difference in the traffic departure pattern here from the case of loading pattern-1 can be seen on comparison of these figures with figure 4.10. The DUE departure pattern involves unequal traffic generation rates into the three highways, and results in a much more spread out traffic loading.

4.2.4.2. Austin network traffic pattern

The traffic pattern assumed for the simulations with the network of Austin is derived from 1985 data on the average daily interzonal travel demand. Two

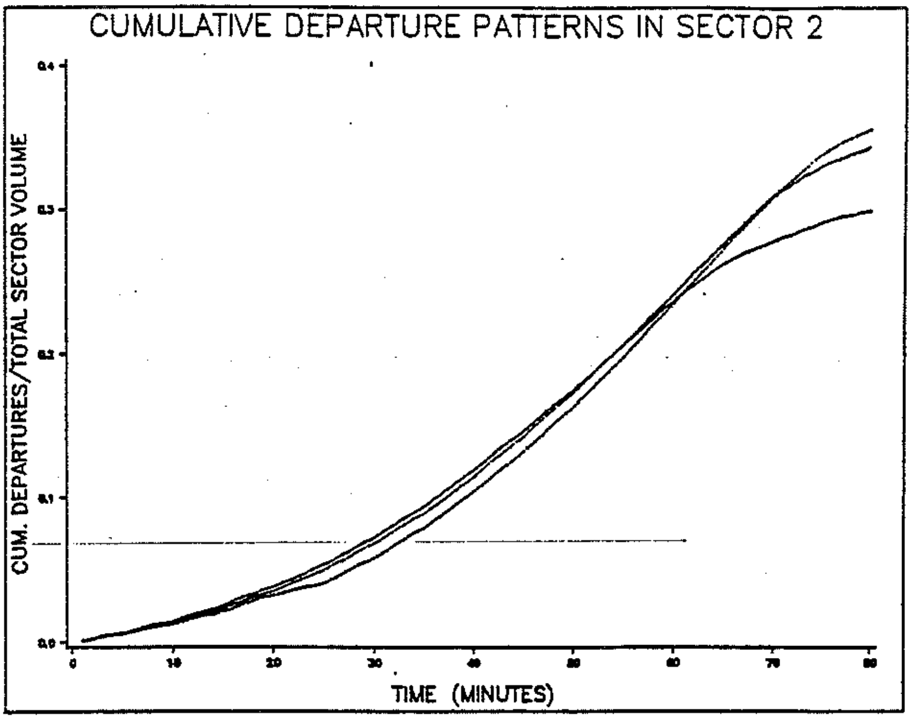

Fig. 4.11. Cumulative departures under Dynamic User Equilibrium loading for the three highways in the sector 7 miles away from CBD. Simulation case: Route-switch threshold fraction = 0.2. Fraction of drivers equipped for information = 0.5)

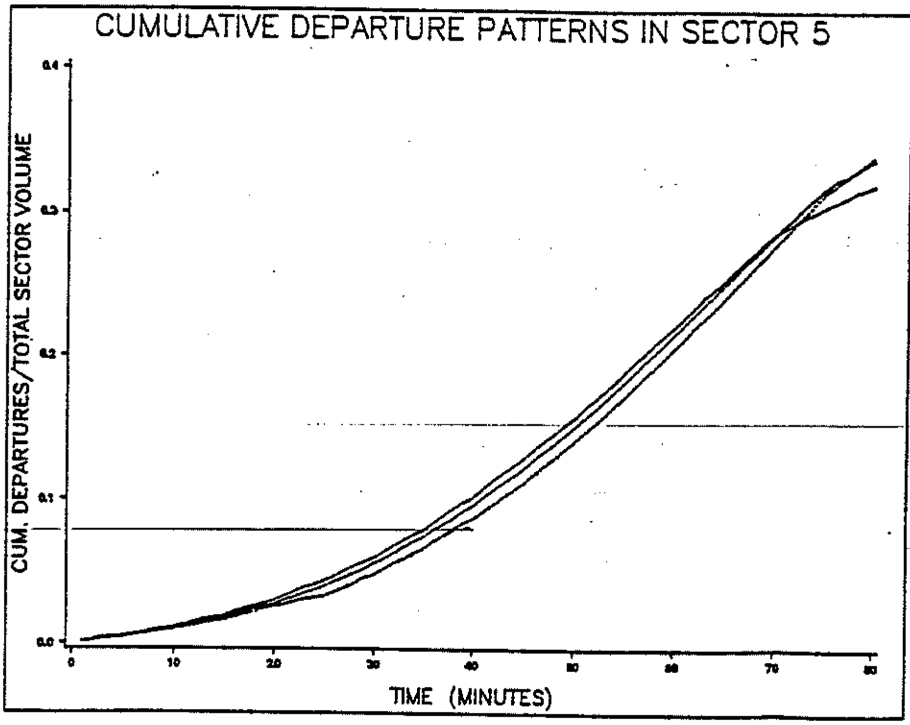

Fig. 4.12. Cumulative departures under Dynamic User Equilibrium loading for the three highways in the sector 4 miles away from CBD. Simulation case: Route-switch threshold fraction = 0.2, Frcation of drivers equipped for information = 0.5)

different loading patterns were used in this study. One is a uniform loading pattern at 0.1 peak hour rate (which is the same rate as in the case of 0.1 fraction of the total daily demand being generated over an hour) and a dynamic 7-interval demand pattern, both derived from the interzonal data. The seven intervals for the dynamic pattern (called 'peaked loading pattern' here) are the network fill-up period of ten minutes during which the demand generation is at 0.075 peak hour rate, the rate changing to 0.15, 0.225, 0.3, 0.225 and 0.15 over the five 5-minute intervals after that with a rate of 0.0375 after the simulation period of interest (the vehicles that entered during the 25 minute period of interest clear out during this period). Thus there is a demand peaking assumed for the simulation period of interest which is the 5 intervals after the fillup (start-up) period. Only the drivers entering the system after the fill-up period till the end of the 6th period are used for calculation of system performance measures. Both the traffic loading patterns are shown in figure 4.13.

As the original interzonal trip interchange data of Austin is a matrix among about 600 zone in the larger Austin area, the data from the zones outside of the 28 zone region being simulated was aggregated to 8 external zone-clusters (which are specified just like any regular zone as part of the input data) based on certain geometric and other considerations. Four of these zones (outside the residential areas shown in fig 4.2) were assumed to feed all their demand to the study area through the corresponding freeway extensions (Mo-Pac and I-35, each from North and South; see figures 4.2-4.9) while the other 4 zone-clusters which surround the study area (these are groups of zones shown in fig 4.2 to the North, South, East and West of the network, which are predominantly residential) are assumed to feed traffic through the entry links around the periphery of the

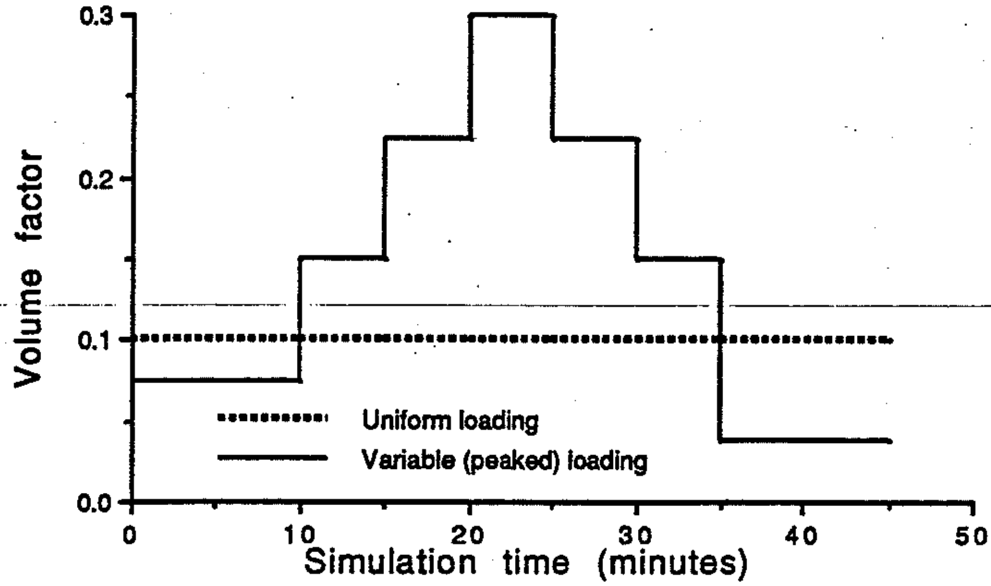

Fig. 4.13 The two loading patterns used for the Austin network. Volume factor = 0.1 results in a vehicle generation rate equal to the rate when 10 % of the average daily demand is generated over an hour.

network on the corresponding sides of the roughly rectangular study area.

The four residential zone-clusters are assumed not to have any demand towards them during the peak period due to their residential nature. Each of the other 4 zone-clusters generate traffic towards the zone-clusters that the corresponding freeway (i.e., the freeway that is extended to that zone. See figs. 4.3-4.9) leads to, in addition to the traffic to the 28 zones within the study area. Thus a reasonable level of through traffic is generated on to the Mo-Pac and I-35 freeways. While this demand pattern is not an exact representation of the actual traffic, it is expected to be somewhat close to reality and sufficient for the current purposes.

All the vehicles that are generated during the fill-up period are assumed not to receive information and reasonable level of traffic is expected in the network at the end of the fill-up period. When a vehicle is generated, it is assigned to an initial path to its destinations, which is randomly selected from the 10 best paths (between its origin and destination) at the end of the fill-up time. It should be mentioned here that the assignment of the initial paths could have significant impacts on the resulting system performance, because the drivers without information stay on these paths throughout their trips. This could cause unrealistic congestion on some of the paths, especially when these 10 paths are not very dissimilar and share common stretches. An alternative would have been a dynamic user equilibrium assignment of the drivers to their initial paths, which is too involved to have been attempted within the scope of these simulations.

4.3. SIMULATION RESULTS

This section presents the results from the various simulation experiments.

The main performance measure is the travel time, the reduction of which is commonly assumed as the single most important reason behind the introduction of information systems in urban traffic networks. Statistics on route-switching are also examined as they provide significant insights into the network dynamics. It should be remembered that other performance measures are possible such as the link-level congestion statistics in the network, variance of the travel times over different groups of drivers and the statistics regarding the average values of earliness and lateness based on assumed work start times. Such measures are valuable in the evaluation of specific cases of information supply design in specific networks. The simulation program produces such output measures, but they are not discussed here because of their site-specific nature and the absence of more general insights in this regard. The following two sections describe the results from the corridor network simulations and the general Austin network simulations. The last section summarizes the conclusions.

4.3.1. Corridor simulations.

An examination of the variation of vehicle concentrations on the highways for the two different departure patterns is helpful in understanding the reasons behind the system performance summary measure results. The vehicle concentration profiles for the three highways over the one mile stretch in sector 5 (at a distance of 4 miles from the CBD) are shown in figures 4.14 and 4.15, for the loading pattern-1 and the DUE loading pattern respectively (for a particular information scenario of half of the drivers with information and 0.2 as the mean route-switch threshold fraction). We can see that due to the considerably more distributed traffic loading, the vehicle concentrations are quite low in the case of

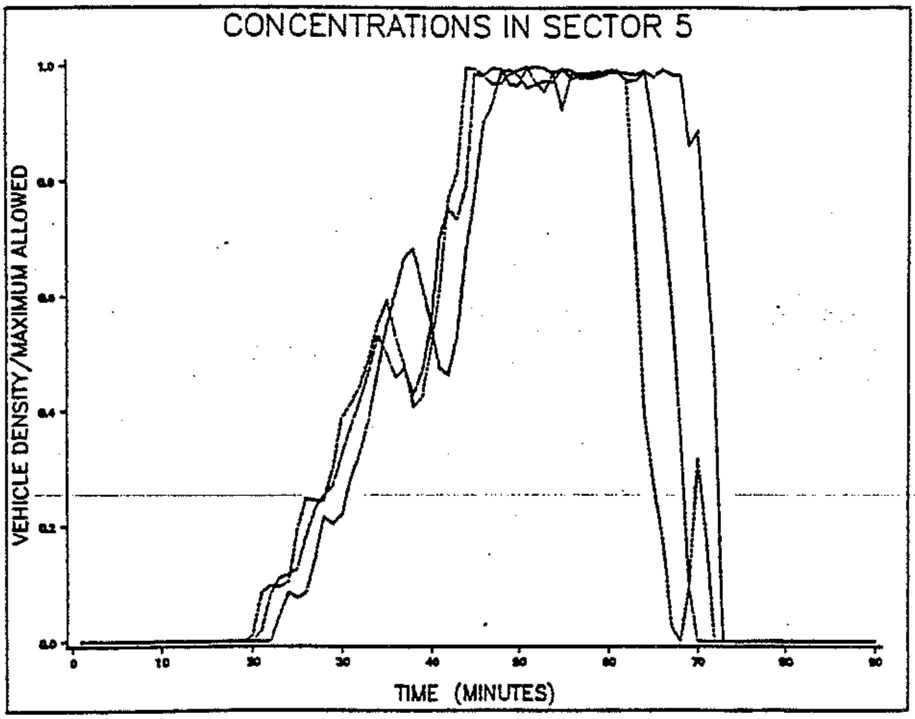

Fig. 4.14 Link vehicle densities under loading pattern-1 for the three highways in the sector 4 miles away from the CBD. Simulation case: Route-switch threshold fraction = 0.2, equipped for information = 0.5)

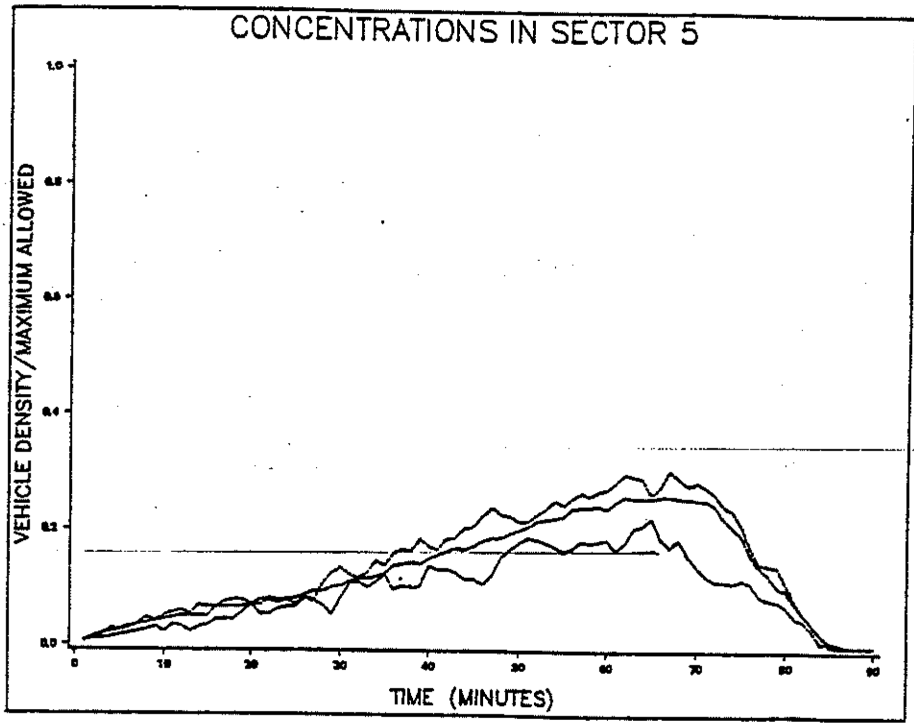

Fig. 4.15 Link vehicle densities under Dynamic User Equilibrium loading for the three highways in the sector 4 miles away from CBD. Simulation case: Route-switch threshold fraction = 0.2, Fraction of drivers equipped for information = 0.5)

the DUE pattern, as compared to the loading pattern-1. In the DUE case, the densities rarely reach congestion levels, while for the loading pattern-1, the highways are quite congested, with long periods with traffic moving at the minimum possible speed (6 mph) under maximum density conditions.

In fact, it is the dynamics of density variations in the three highways that induces much route-switching behavior. As the information supply system assumed does not predict future trip time variation on alternative routes, which is a function of the speed and concentration levels along that route, drivers may select alternative highways which perform worse (often due to the heavy route-switching by drivers to that route). This aspect can be clearly seen in figure 4.14 where densities in highway-1 and highway-3 seem to vary in a negatively correlated fashion, indicating the improvement and worsening of the conditions due to switching activity between the highways. This also means that the drivers often encounter worse conditions in the alternative routes than they expected while switching, at least during the congestion build-up phase of the peak period.

Figure 4.16 depicts the variation of the systemwide trip time with the fraction of the population with access to information, under each of the five assumed levels of switching propensity. Note that the trip time is expressed as a percent of the total trip time under the no information case (i.e., no switching); thus the values in excess of 100 percent correspond to a worsening of systemwide performance under information. Such worsening occurs under the assumption of myopic switching, when users switch routes any time an alternative path offers some improvement in trip time, no matter how small, over the current path.

On the other hand, when an indifference threhold is assumed, there appears to be benefits in terms of trip time. The best systemwide improvement

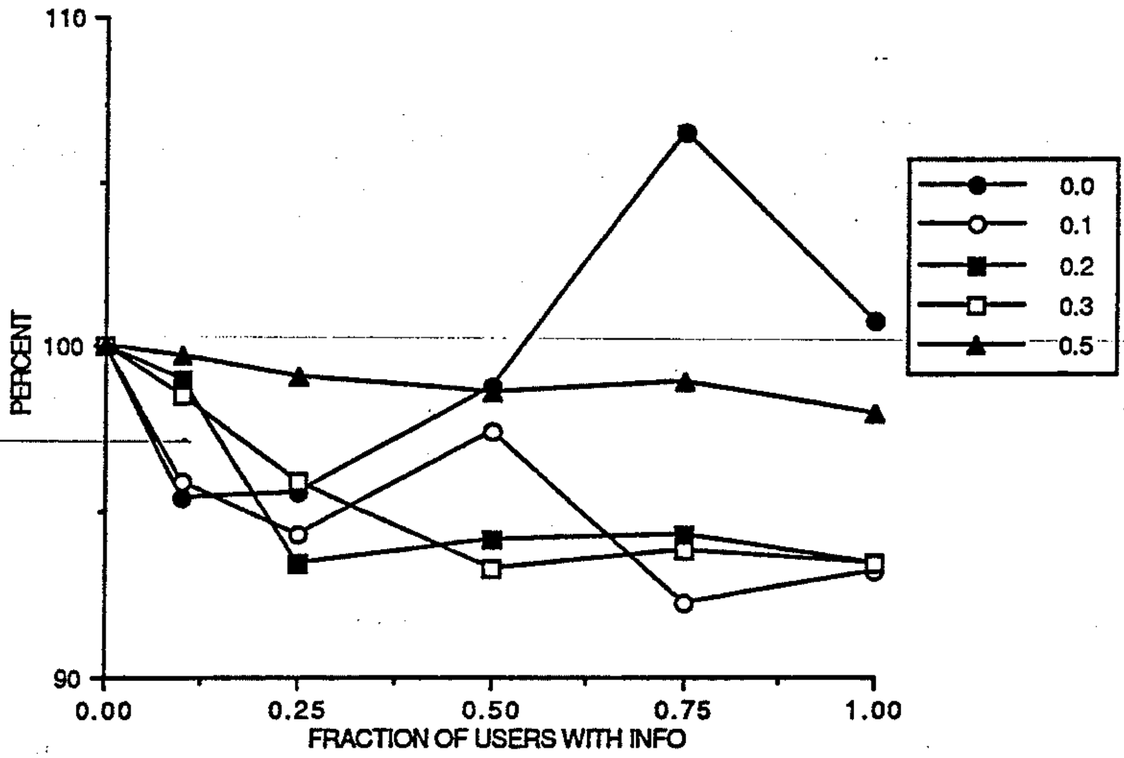

Fig. 4.16 The variation of the total system travel times in the 3-highway corridor as a percentage of the trip time in the base case of no information, under different propensities for driver route switching. Case: Loading pattern-1. (Note: One curve for each value of the mean route-switch threshold fraction. The values are shown in the legend.)

takes place under a mean relative route switch threshold fraction (η_j) of 0.2 to 0.3, i.e., when the equipped drivers on average do not switch routes at opportunities that offer less than 20 or 30 percent improvement in the remaining trip time. At the other extreme, with a η_j value of 0.5, the drivers miss too many opportunities, resulting in a systemwide improvement of no more than 2 percent. In general, the marginal system-wide improvements decrease significantly as the information supply reaches higher market penetration, especially after it reaches about 25 percent of the population. The maximum systemwide benefit obtained is about 7 to 8 percent reduction in overall trip time.

Figures 4.17 and 4.18 depict the average trip times (expressed as a percent of the corresponding average in the base case of no information supply) experienced by those who have access to information and those who do not, respectively[61]. Several important phenomena are illustrated by these graphs. First, the users *with* information could do worse than they would have in the no information case, when at the same time those *without* information experience a reduction in their average trip time relative to the base case. Second, this worsening is experienced by users when they switch too readily, as seen in the zero band for higher fractions with information. Third, benefits are incurred by those *without* information in most of the cases here. Fourth, the relative

[61] Please note that the curves are not connected to the 100 percent point on the ordinate in the case of average trip times of the group with information, which is because at very low fractions, this group may not be expected to have average trip times near that in the base case. Also note that there are no drivers without information when 1.0 fraction of drivers receive information, which explains why the average trip time graphs extend only to 0.75 for the group without information.

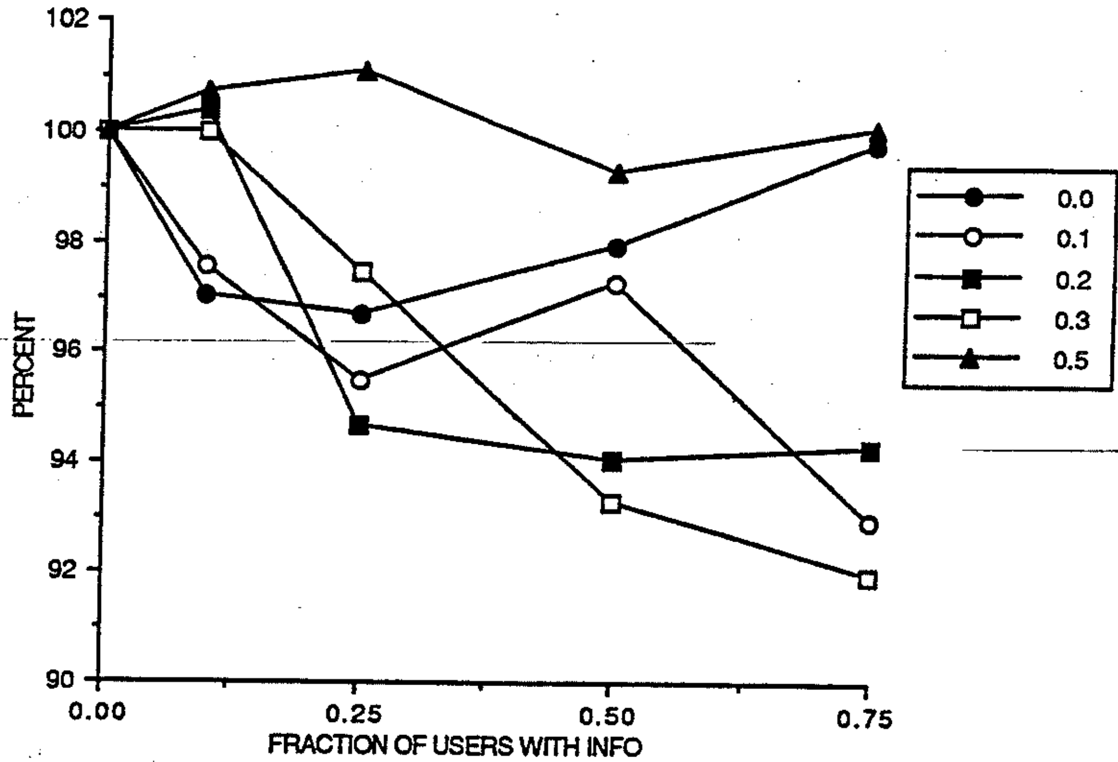

Fig. 4.17 Average trip times for drivers not receiving information in the 3-highway corridor as a percentage of the average trip time in the base case of no information, for different driver route-switching propensities. Case: loading pattern-1. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

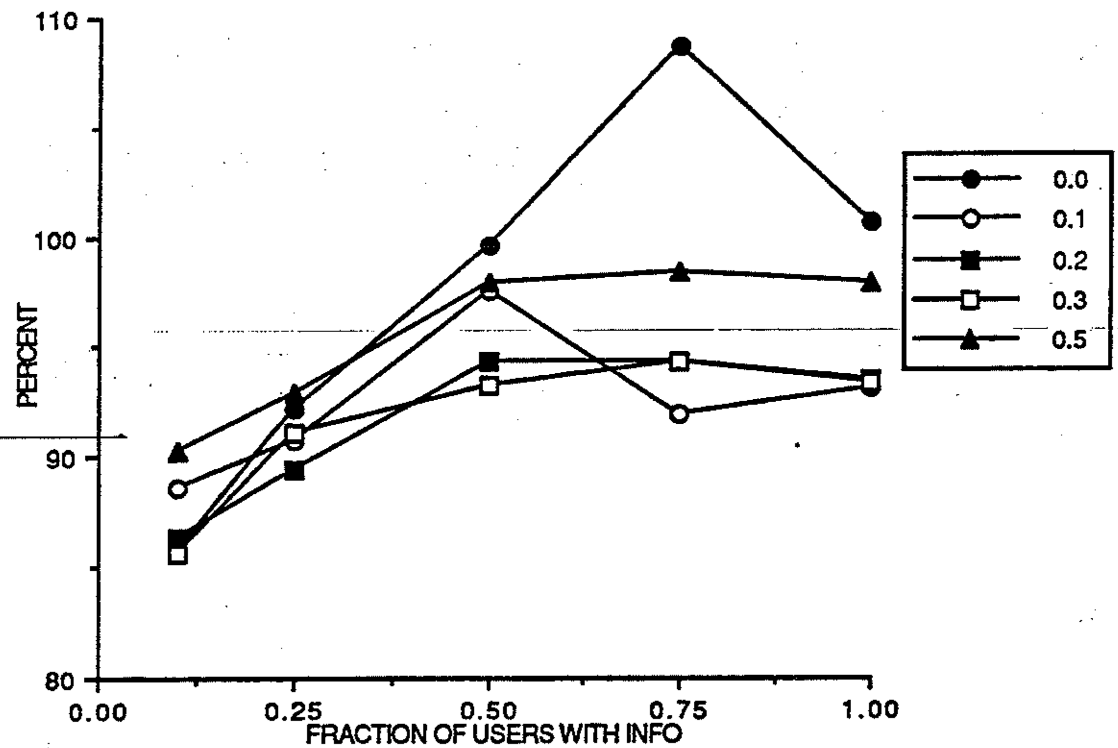

Fig. 4.18 Average trip times of the group receiving information in the 3-highway corridor as a percentage of the average trip time in the base case of no information, for different driver roue-switrching propensities. Case: loading pattern-1. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

performance of the two groups is strongly dependent on the underlying switching propensity, and/or fraction of information. The group with information shows a trend of decreasing benefits as the fraction with information increases. On the other hand, the group without information shows increasing benefits from information as the fraction with information increases.

Similar phenomena can be detected in the case of the Dynamic User Equilibrium (DUE) departure pattern also. As this departure pattern results in considerably less congested conditions than loading pattern-1, it should be expected that the potential of information to reduce systemwide trip time would diminish. Figure 4.19 presents the variation of the system trip time improvement under information for the DUE traffic loading. The results differ from those under the loading pattern-1 in that while no cases are encountered where the supply of the information increases the total time, the reductions attained are generally smaller in relative magnitude than those under the more congested first loading pattern. The maximum improvement is only about 5 percent. The same general trends as in the case of loading pattern-1 are present here, in that the marginal effectiveness of information dramatically decreases, and in some cases actually reverses (as in the case of myopic switching with zero threshold) after a certain fraction of the population, between 25 to 50 percent, is equipped to receive information.

Figures 4.20 and 4.21 present similar results as figures 4.17 and 4.18 for the DUE pattern, reflecting the relative improvement of average trip time for users with and without information. The results differ from those under the first loading pattern in that the users with information seem to outperform those without information in virtually all cases. This appears to be a case of the system having

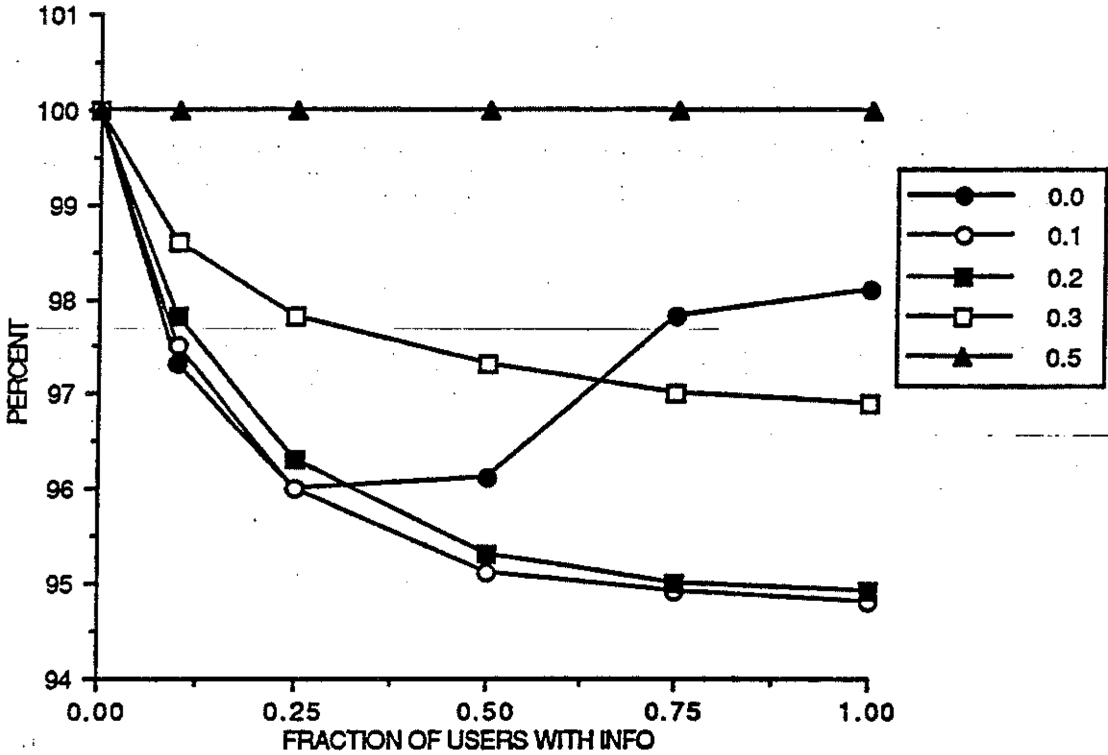

Fig. 4.19 The variation of the total system travel times in the 3-highway corridor as a percentage of the trip time in the base case of no information, under different propensities for driver route switching. Case: DUE loading pattern. (Note: One curve for each value of the mean route-switch threshold fraction. The values are shown in the legend.)

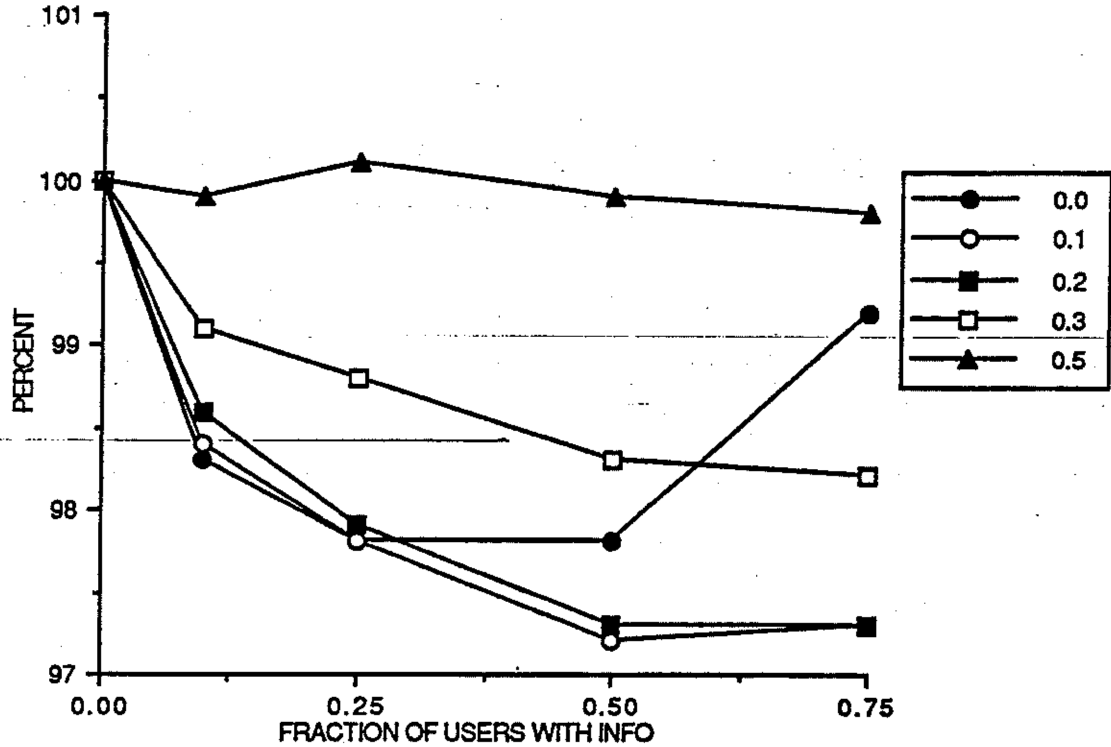

Fig. 4.20 Average trip times for drivers not receiving information in the 3-highway corridor as a percentage of the average trip time in the base case of no information, for different driver route-switching propensities. Case: DUE loading pattern-1. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

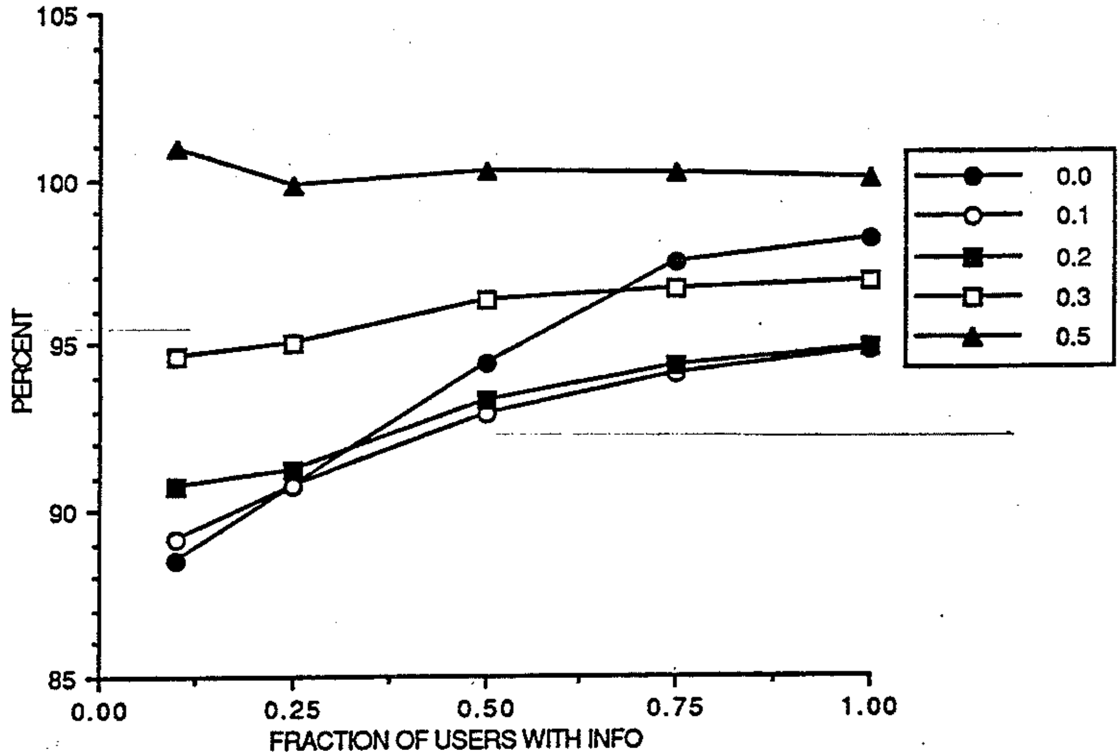

Fig. 4.21 Average trip times of the group receiving information in the 3-highway corridor as a percentage of the average trip time in the base case of no information, for different driver roue-switrching propensities. Case: loading pattern-1. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

a 'smoother' traffic pattern with greater peak spreading and less congestion, where information availability allows the drivers to fine-tune their path selection, thus causing their locally optimal decisions to contribute towards global improvements.

4.3.3. General Austin Network Simulations

Compared to the cases of the corridor simulations, the Austin network simulations show a few different performance trends. Some of these differences are expected due to the vastly different network structure as well as the assumed traffic loading. First, in this network each driver makes route-switch decisions with different sets of routes while all the drivers have the same three routes to select from in the case of the corridor. Second, the general route structure and the abundance of routes of comparable characteristics imply that the routes in the choice sets of the drivers may not be providing vastly different trip times (note that in the case of the corridor the choice set always consists of three non-overlapping highways of different speeds, while in the general network, the K-shortest path choice set has routes which may be sharing common stretches). A third and very important factor is the initial path assignment to the drivers, which are always based on the best routes stored at the end of the network fill-up time. The differences in the performance trends discussed below, as compared to the corridor network, should be evaluated keeping those factors in mind. The conclusive section of this chapter provides further discussion on the effects of these factors.

The difference between the two traffic loading patterns (uniform and peaked) is clear from the vehicle densities in a representative link shown in figures 4.24 and 4.25. The peaked pattern results in higher densities in the network

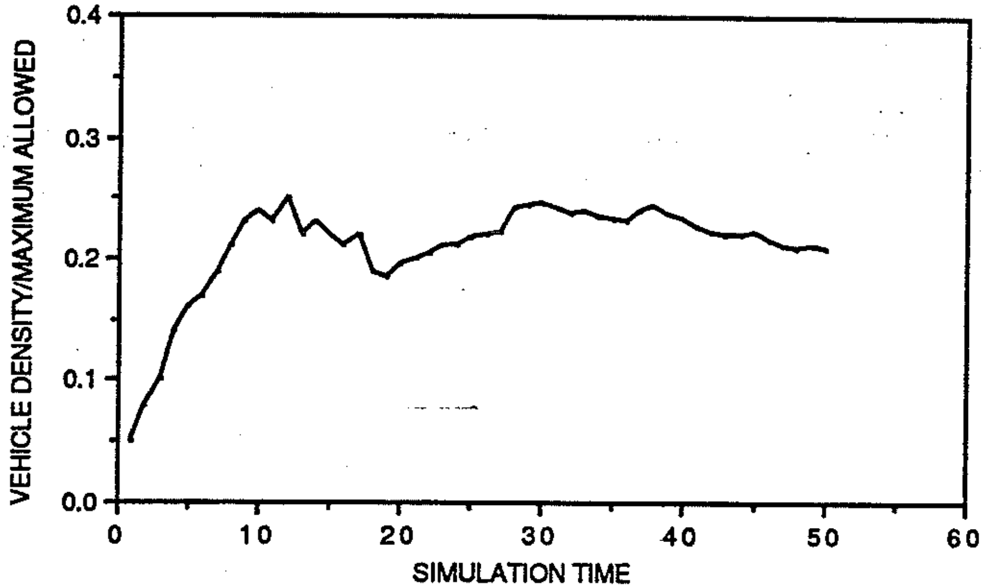

Fig. 4.24 The vehicle density variation in a selected link (North-bound Mo-Pac north of Colorado river) during the simulation of Austin network (Load case : Uniform loading)

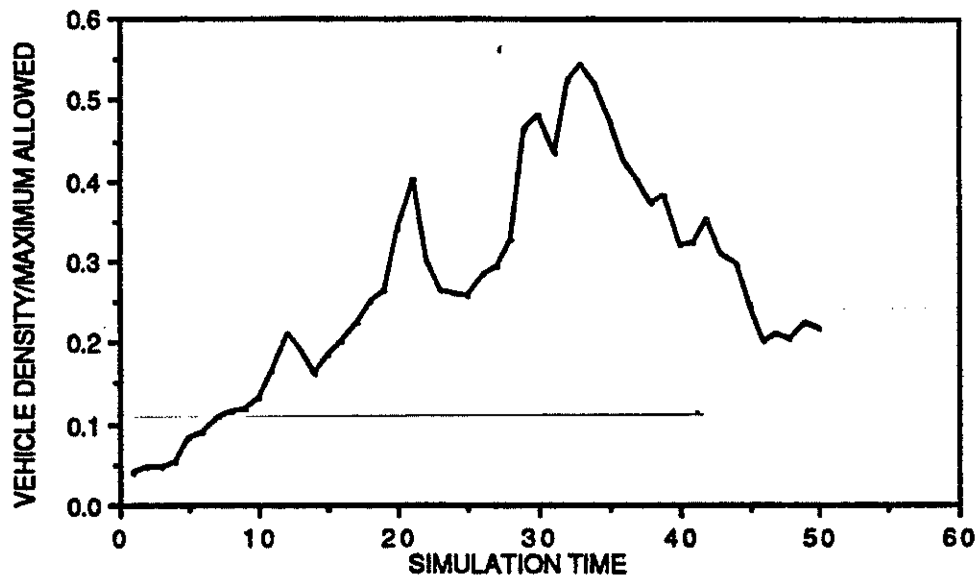

Fig. 4.25 The vehicle density variation in a selected link (north-bound Mo-Pac north of Colorado river) during the simulation of Austin network (Load case : Peaked loading)

and shows higher variations between different information scenarios, pointing to more switching opportunities in the case of the peaked pattern. It is interesting to note that this did not result in better overall trip time benefits for the system or for the group of drivers with information, as compared to the uniform loading case, which is discussed next.

Figure 4.26 depicts the variation of the total system travel time (expressed as usual as a percentage of the base case of no information) for the uniform loading pattern. We see similar patterns of system level trip time advantages, in that most of the possible benefits are achieved with low market penetration of information, with only 25 to 50 percent of the drivers equipped for information. The system does not perform worse under information supply in almost all the cases. However, in contrast to the corridor case, the case of myopic switching does not appear to result in worse performance compared to route-switching with an indifference threshold.

The impact of information on the two groups can be seen in figures 4.27 and 4.28, which present the average trip times of the drivers without information and with information, respectively, as percentages of the no-information case. Again, myopic switching appear to result in better benefits to both these driver groups. These figures show that when the number of drivers equipped for receiving information increases, the average trip times of the drivers with information remain rather constant and the trip time benefits of the drivers not receiving information increases. This also is a different trend than in the corridor case where increasing market penetration of information results in larger benefits to the group without information at the expense of the group with information whose trip time benefits decrease.

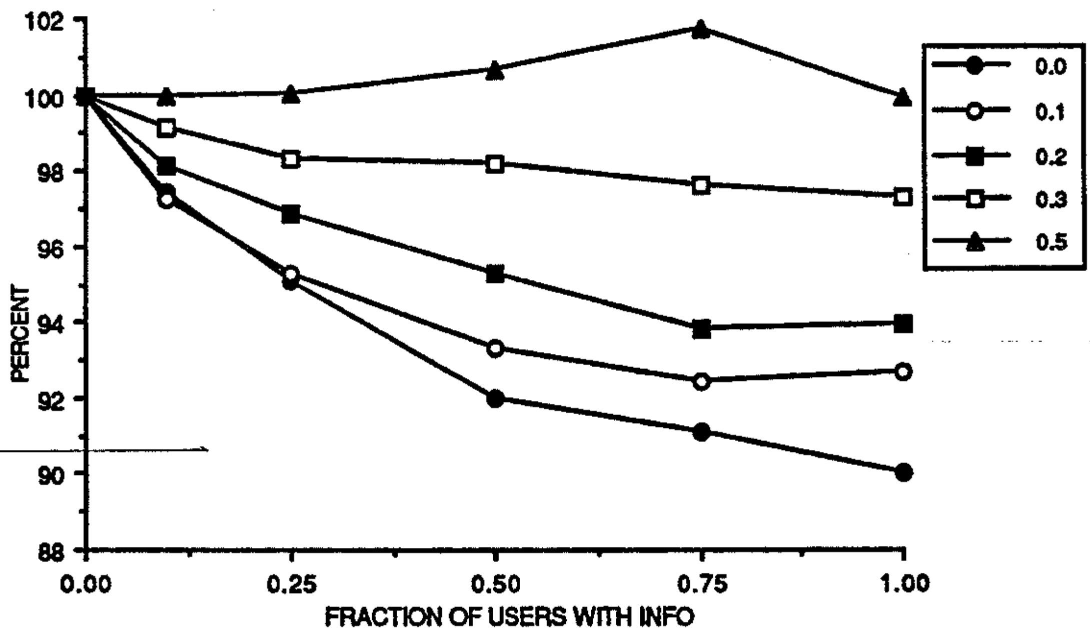

Fig. 4.26. The variation of the total system travel times in the large Austin network as a percentage of the trip time in the base case of no information, under different propensities for driver route switching. Case: Uniform loading. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend.)

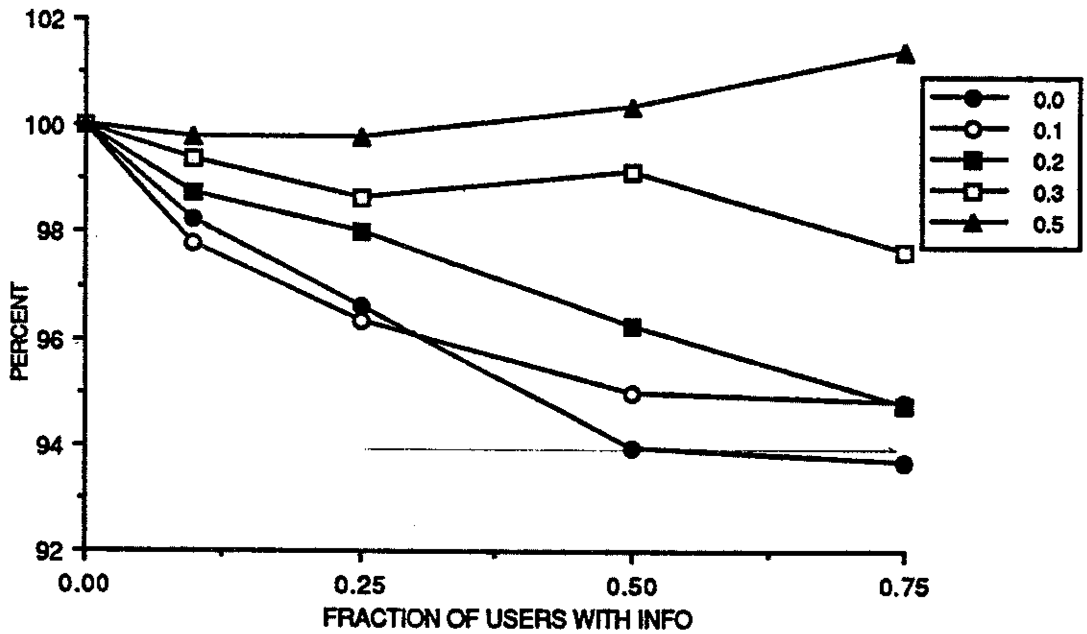

Fig. 4.27. Average trip times for drivers not receiving information in the Austin network as a percentage of the average trip time in the base case of no information, for different driver route-switching propensities. Case: Uniform loading. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

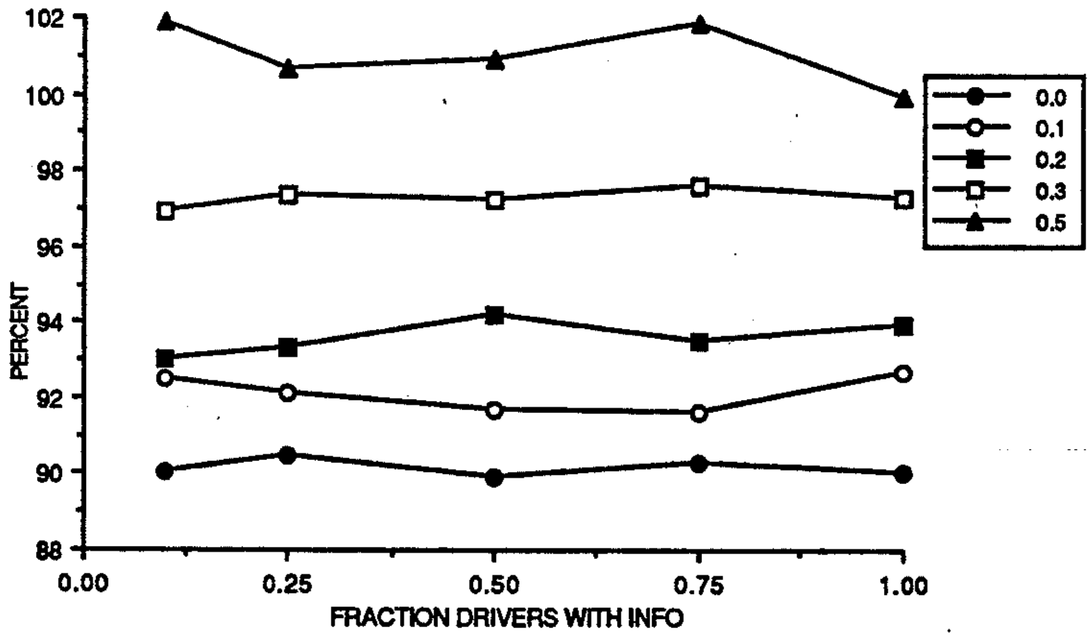

Fig. 4.28. Average trip times of the group receiving information in the Austin network as a percentage of the average trip time in the base case of no information, for different driver roue-switrching propensities. Case: Uniform loading. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

The case of the peaked loading shows some different performance characteristics, which appear to differ from the uniform loading case as well as the corridor cases. Figures 4.29, 4.30 and 4.31 depict the travel time benefits for the system and the two groups of drivers as before. The most notable difference with the uniform loading case is that the benefits derived by the group with information decreases for higher market penetration of information (fig 4.31) while it was staying somewhat constant in the uniform loading case. Of course, this is in line with the trend shown by the corridor cases. The reason is that there are more switching opportunities which causes over-reaction by the drivers and congestion on alternative streets to which they switch. We again see reducing marginal improvement of the system average trip time on increased market penetration, as shown in fig. 4.29. The benefits derived by the group without information is smaller than those in the case of uniform loading (fig. 4.30). This can be explained as follows: the peaked loading case introduces higher level of congestion in the stored paths assigned to the drivers without information, as compared to the uniform loading case. Even though the drivers without information have switched away from these paths, the larger number of drivers in these paths itself reduces the relative benefits from the switching of the drivers with information. As in the case of the uniform loading in the Austin network, but contrary to the corridor simulations (with loading pattern-1), the myopic switching behavior tends to result in better benefits, systemwide and also for the group of drivers with information.

It is also notable that even with higher variation of density levels (implying more route switching opportunities), the dynamic peaked loading results in lesser trip time benefits than the uniform loading in the Austin network. This

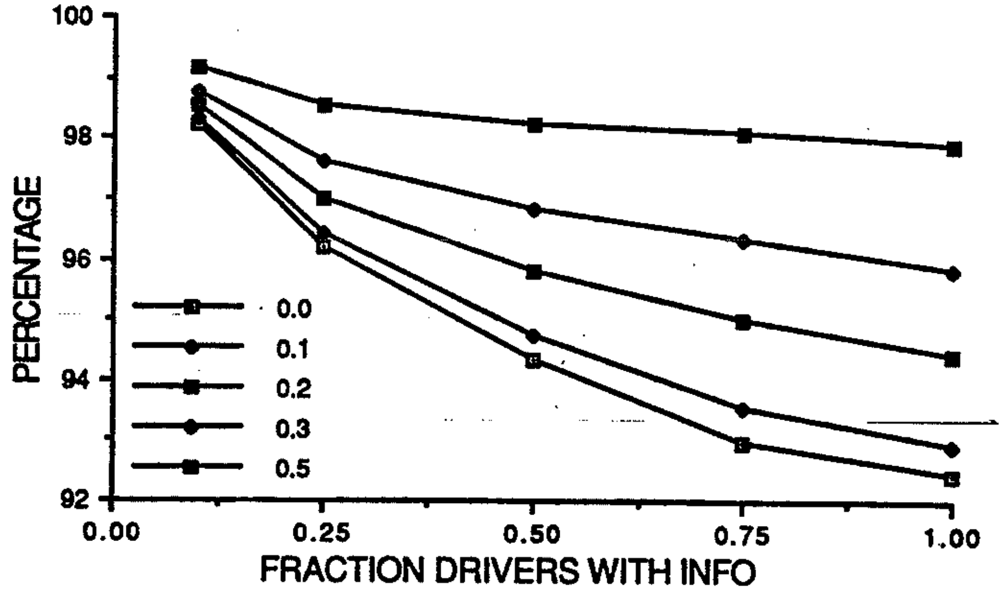

Fig. 4.29. The variation of the total system travel times in the large Austin network as a percentage of the trip time in the base case of no information, under different propensities for driver route switching. Case: Peaked loading. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

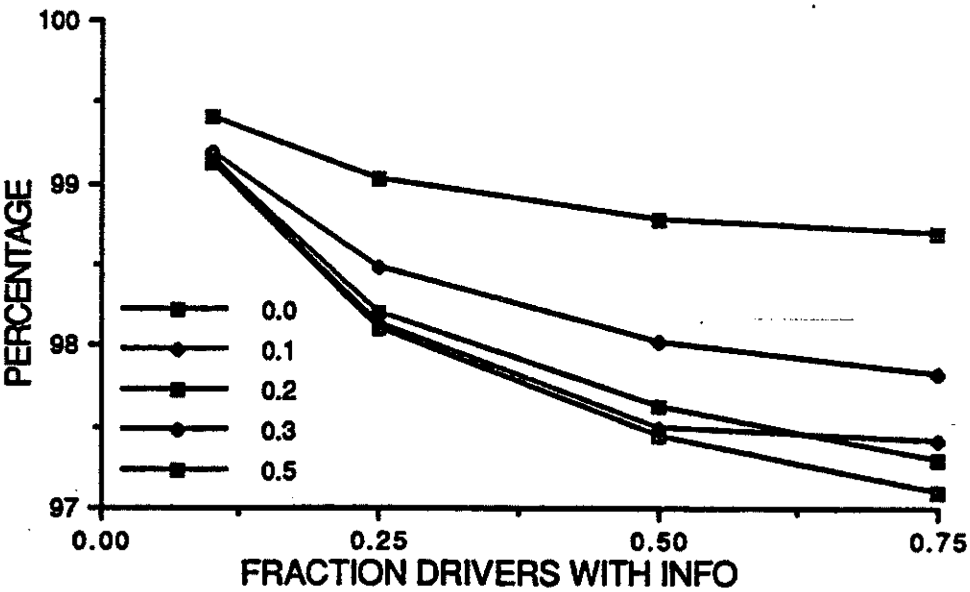

Fig. 4.30. Average trip times for drivers not receiving information in the Austin network as a percentage of the average trip time in the base case of no information, for different driver route-switching propensities. Case: Peaked loading. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

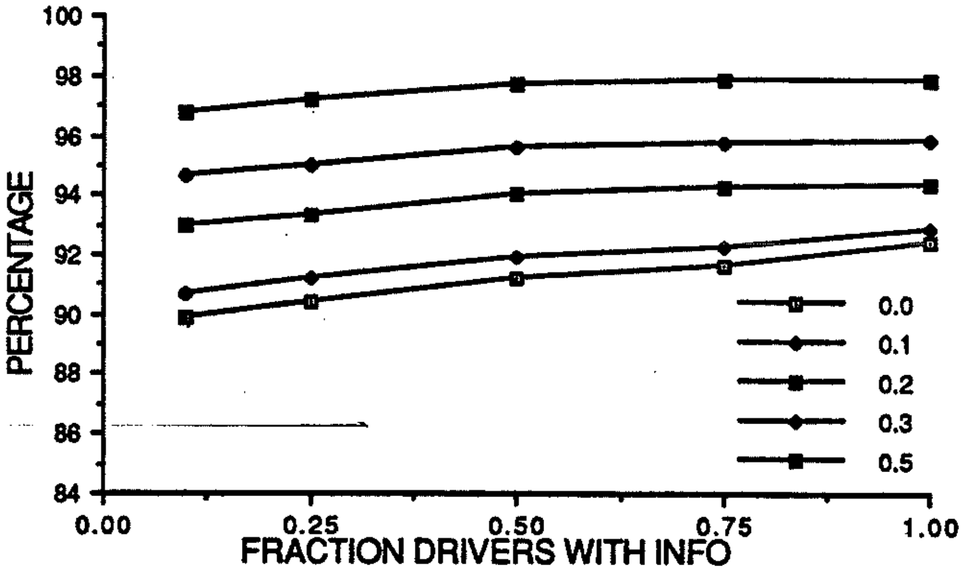

Fig. 4.31. Average trip times of the group receiving information in the Austin network as a percentage of the average trip time in the base case of no information, for different driver roue-switrching propensities. Case: Peaked loading. (Note: One curve for each value of the mean route-switch threshold fraction. The values are as shown in the legend).

is in line with the results from the corridor simulations, where the trip time benefits were less when more drivers were making route switches.

The route-switching statistics of a specific simulation run with the peaked loading case for the Austin network reveal further interesting aspects. These statistics are shown in tables 4.1 and 4.2. The myopic switching behavior (i.e., $\mu_j = 0$) produces many route changes during the trip (table 4.1), while the other cases only rarely produce more than one route change during the trip (table 4.2). The average distances at which the route switches ocuurred show that about 70 to 80 percent of the route switches by the group with information are made within the beginning 0.40 fraction of the trip length and about 50 percent occur within 6 percent of the trip length, in the cases of the Myopic switching. This is strongly indicative of a preponderance of route changing near the starting point. This is important because the initial paths assigned to the drivers are always based on the paths at the end of the fillup period, and may be quite different from the best paths available to a driver at departure. This points to the necessity for better methods of initial path assignment.

The non-mypoic switching cases (table 4.2) show that as the indifference threshold fraction increases, the route switching is attempted farther away from the starting point. The fact that there are no second switches with the non-myopic switching cases implies that the drivers are switching to relatively good routes after which they do not find better routes. Another interesting aspect is that there are more single switches in the non-mypoic switching case within a short distance from the starting-point of the drivers than in the case of non-myopic switching. What this means is that when the drivers switch to alternate routes whenever they find some trip time advantage on them, they cause the trip times on their current paths to be better, thus preventing further switches. This also explains why the

Route switch threshold fraction	Fraction of drivers with info.	Total # of route switches	Route switch number	Number of vehicles making switches	Average fractional trip length of the switches
0.0	0.10	4315	1st	2130	0.0574
			2nd	578	0.2643
			3rd	468	0.3571
			4th	330	0.4586
	0.25	11510	1st	5343	0.0577
			2nd	1440	0.2530
			3rd	1153	0.3398
			4th	844	0.4278
	0.50	26856	1st	10745	0.0588
			2nd	3177	0.2395
			3rd	2569	0.3215
			4th	1936	0.3920
	0.75	44484	1st	16289	0.0594
			2nd	5076	0.2288
			3rd	4173	0.3047
			4th	3287	0.3778
	1.00	62542	1st	21793	0.0593
			2nd	7020	0.2249
			3rd	5864	0.3067
			4th	4718	0.3778

Table 4.1. Route switch statistics for mypoic driver behavior. Case: Austin network, Peaked loading. (Note: Some vehicles made more than 4 switches.)

Route switch threshold fraction	Fraction of drivers with info.	Total # of route switches	Number of single switches	Average fractional trip length of the switches
0.1	0.10	1418	1418	0.1904
	0.25	3583	3583	0.1871
	0.50	7082	7081	0.1948
	0.75	10579	10578	0.1951
	1.00	14035	14034	0.1943
0.2	0.10	1022	1022	0.3365
	0.25	2559	2559	0.3259
	0.50	4999	4999	0.3259
	0.75	7475	7475	0.3272
	1.00	9916	9916	0.3281
0.3	0.10	647	647	0.4167
	0.25	1738	1738	0.4064
	0.50	3390	3390	0.4047
	0.75	5067	5067	0.4126
	1.00	6697	6697	0.4136
0.5	0.10	365	365	0.6664
	0.25	941	941	0.6667
	0.50	1829	1829	0.6621
	0.75	2672	2672	0.6651
	1.00	3540	3540	0.6639

Table 4.2. Route switch statistics for non-myopic driver behavior. Case: Austin network, Peaked loading. (Note: Almost all vehicles made only 1 route switch.)

density patterns in figure 4.24 show marginally smoother density variations in the case of the myopic switching.

For the myopic switching cases, the average fractional distances at whichroute switching occurs is relatively robust under different market penetrations, the reason for which is not very clear. These increases as the indifference threshold fraction increases, which could be due to the increased switching activity nearer to the destination due to heavier congestion at these zones.

4.4. ILLUSTRATIVE RESULTS ON THE MODELLING FRAMEWORK

This section serves three purposes: 1) Give some examples of other insights on network performance that can be derived using the modelling framework, which would highlight its additional capabilities, 2) Verify that the model is 'reasonable' in the way it captures the dynamics of networks under information and 3) Provide a description of its computational capabilities. Subsection 4.4.1 reports on the dynamic changes of the network paths as well as the path choices of a representative driver. Section 4.4.2 provides certain computational results.

4.4.1. Path dynamics and Driver path selection

It is important to understand the dynamic changes that occur in the network with regard to the shortest paths. As the k-shortest paths are stored in this framework, it is useful to gain insights into the dynamics of the k-shortest path sets. This has implications on deciding how many paths should be there in the path sets (i.e., what should be a reasonable value for 'k'), as well as to confirm that such an approach is reasonable.

Table 4.3. provides some illustrative results on how the shoortest paths changed on one specific simulation of the Austin network. The results shown are the average statistics from the path sets for all the O-D pairs compared every 2 miutes of simulation. A few intersting observations can be made from this. On the average about 3 of the paths change their 'rank' in the path-set (i.e., swap positions in the k-path list ordered according to the shortness of the paths or they drop out of the path) during a 2 minute period. Only about 1 path out of the 10 drop out of the path list, though. This means that there is no need to re-enumerate the paths after each simulation time step (0.1 minute in the simulation runs here). It would be sufficient to just recalculate the path trip times every simulation time-step. Thus, the premise that a path trip time aggregation module (which is faster than the k-shortest path module) had to be developed, is verified. The facts that 3 paths change positions in the path list during 2 minutes as well as that the shortest paths changes about 25 percent of the time are also important, as they help in deciding the frequency of shortest-path enumeration. The updating interval used her (1 minute) appears to be adequate based on this.

Fig. 4.32 shows the path followed by a specific driver going towards a downtown node in Austin. We see that the initial path that the driver started on, was quickly changed to an alternative path which is not changed any more. This behavior of a route-switch in the early part of the trip is confirmed also by the route-switching statistics provided in the last section. It also implies that, as the initial path assignment is based on a stored set of paths which do not depend on the prevailing conditions at departure, the driver is able to find an alternative route

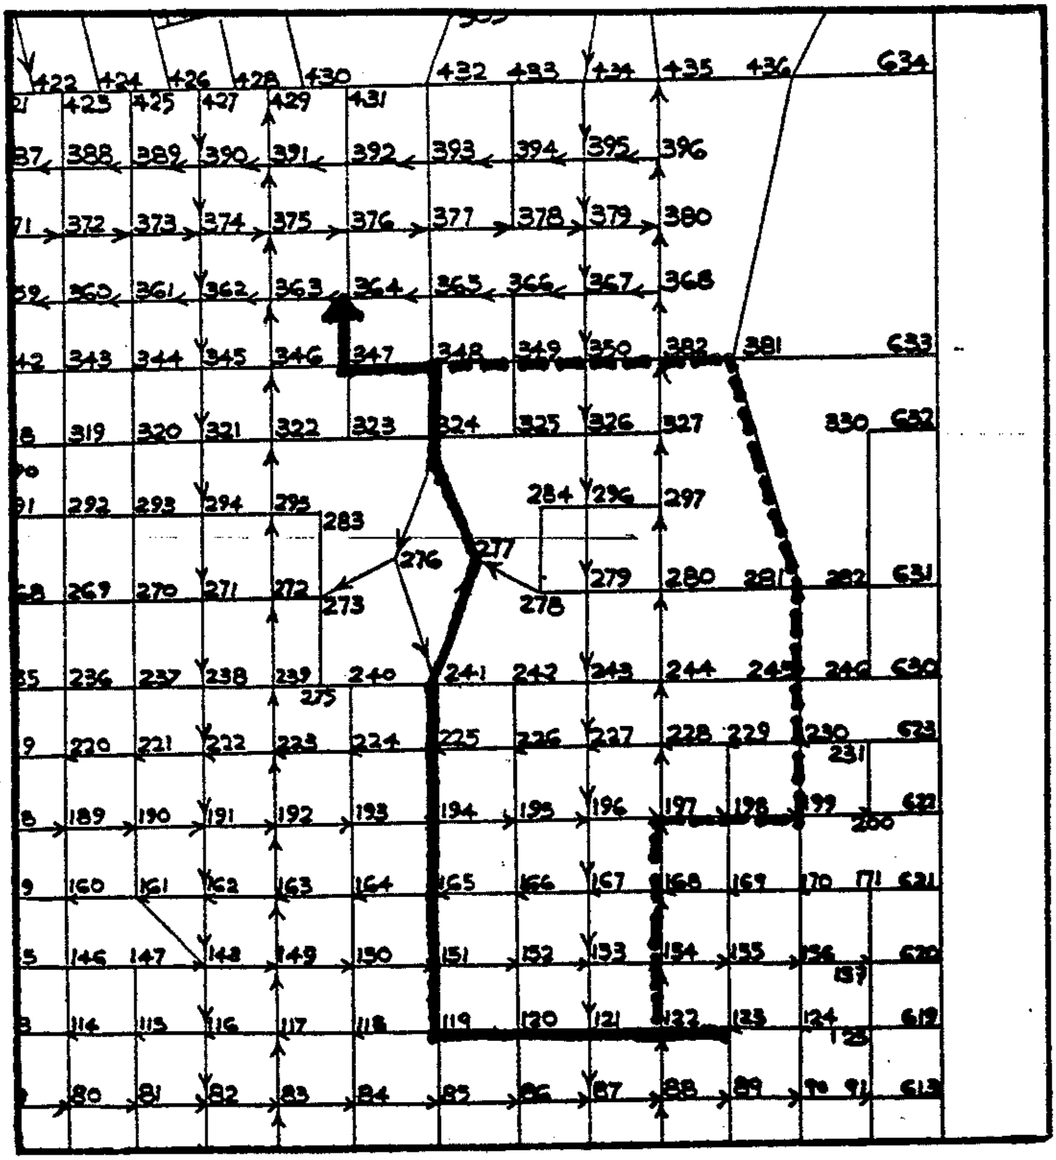

Fig. 4.32 Initial path and the path after a switch for a representative driver in the Austin network.

Time	Average # of paths out of 10 that changed	Average # of paths out of 10 that dropped out	Fraction of cases where the shortest path changed.
2.0	2.16	0.51	0.18
4.0	2.79	0.62	0.21
6.0	3.06	1.03	0.24
8.0	3.66	0.94	0.28
10.0	3.87	1.32	0.27
12.0	3.71	1.13	0.29
14.0	3.49	1.05	0.27
16.0	3.89	1.28	0.29
18.0	2.94	1.01	0.25
20.0	3.60	1.19	0.28
22.0	3.55	1.06	0.28
24.0	3.61	1.32	0.27
26.0	3.16	1.14	0.25
28.0	3.95	1.42	0.26
30.0	3.04	0.98	0.27
32.0	3.68	1.03	0.29
34.0	2.79	0.99	0.23
36.0	2.90	1.02	0.25
38.0	2.96	1.09	0.25
40.0	3.10	1.04	0.26

Table 4.3. Path statistics from successive K-shortest path sets.

very quickly. Thus the information supply here is effectively similar to that of pre-trip information supply. This underscores the concerns expressed earlier about the initial path assignment.

4.6. Computational Results

Three different aspects of the computational intensity of the general modelling framework are studied using the Cray YMP-8/64 computer: 1) the computational intensity of the various component modules of the framework, 2) the computational effort when the model is applied to networks of different sizes with different demand levels and 3) the comparison between the sequential and vectorized path-updating routines. No results on the corridor simulation model are provided as this model was developed for a specific network structure and was found to have very small computational times (about 30 seconds for a 2-hour simulation of the three-highway corridor) compared to the general network, due to its fast path-processing capabilities.

Table 4.4 provides the representative results on the computational intensity of the program routines as found in a 50 minute simulation of the Austin network. We see that about 75 percent of the execution time is spent on path processing, as compared to about 20 percent in vehicle simulation. It is seen that the path trip time aggregation routine executes about 7 times faster than the k-shortest path enumeration routine. This means that calling the enumeration routine only once in 10 simulation time steps reduces the computational effort tremendously. While there is no easy way to judge the efficiency of the algorithms and data structures used here, it can be seen that most of the premises under which the path-

	Total time (seconds)	Time per call
Main program	32.8	32.8
Vehicle Moving	98.4	0.2
Decision Modelling	2.1	---
Initial Path Selection	0.4	---
K-shortest path finding	112.3	4.49
Path trip time aggregation	312.0	0.70

Total Execution Time = 558.0 seconds

Table 4.4. Computational Performance of the Component modules.

Case	Nodes	Arcs	Vehicles	Computation Time (sec)
5 by 5 grid	25	80	1250	6.58
	25	80	2500	7.73
10 by 10 grid	100	360	5000	49.2
	100	360	10000	52.8
Austin	630	1770	22000	558.0
	630	1770	28000	576.0

Table 4.5 Computational Intensity of different network cases.

processing routines are developed carefully are proved reasonable, especially the computational intensity in modelling this component of network dynamics.

Table 4.5 provides results on applying the framework to various network sizes. Three different sizes are considered. A 5 nodes by 5 nodes network, a 10 nodes by 10 nodes network and the Austin network. Both the grid networks have two-way streets which are each 500 ft long between the nodes. For the grid networks, the total demand was divided equally among all the nodes and the vehicle generation is uniform over a 50 minute period. We see that the computational intensity does not show a highly non-linear increase with the number of nodes (or number of arcs, which itself is roughly linear with number of ~~nodes in city~~ networks). In fact the times increase ten-fold as the number of nodes increases four-fold, from 5 by 5 to 10 by 10. There is an eight-fold increase as the number of nodes increase another 6-fold to the Austin network case. This is a very significant result, pointing to the fact that even larger networks may be simulated without an unacceptable increase in computational effort. The computation times vary only very little as the demand levels vary. This has to be expected, as the main loops in the simulation are written in such a way that all the links (in fact, link-segments) are looped over and the vehicle loop within this has relatively constant upper limit (due to maximum number of vehicles in each link-segment) regardless of overall demand levels.

4.5. CONCLUSIONS

Several conclusions can be derived from the simulation studies reported in this chapter. Most of these are qualitative in nature, but nevertheless may be considered sufficiently generalizable, due to the variety of information scenarios simulated for networks and loading conditions of extremely dissimilar nature.

The most significant conclusion is that traffic networks could obtain significant benefits from implementing driver route information systems. There were only very few cases where the system performed worse under information and even in those cases the system trip time increases were small (less than 2%). Up to 10 percent improvement was obtained under information, which is probably more than what can be obtained in congested urban traffic systems with any of the existing traffic engineering solutions (except, perhaps demand peak-spreading, which is usually a socio-political solution beyond the traffic engineer's control).

In most of the cases, it is found that relatively low levels of market penetration (about 25 to 50 percent of drivers with information equipment) is sufficient to achieve almost all the possible benefits. It is also seen that the driver population with information may receive decreasing benefits when more drivers are equipped. These, coupled with the fact that the drivers without information receives increasing benefits, open up an array of issues regarding equitable and cost-effective designs of information supply schemes.

Certain results from the simulations of the Austin network may appear contradictory to the results from the corridor simulations. Most significantly, it is seen that the perfectly myopic behavior of switching to a route that provides a trip time advantage, however small, results in worse system performance in the corridor case. On the contrary, for the Austin network, myopic switching results in the best performance. One reason for these seemingly contradictory results could be the way the initial paths are assigned to the drivers in the general network simulations. These paths are randomly selected from the best paths stored at the end of the fill-up times, which may be considerably different from the best paths later on. This means that an initial switch would almost always provide them a better path with considerable benefits. This and the fact that route-switching is rarer in the general network due to the very similar paths in the k-shortest paths imply that those who do not make a switch would travel on a worse route. Such effects due to the initial path assignments can be avoided by incorporating a calibrated model of initial path selection. At the present time no such models are available and it is difficult estimate the effect of the initial paths in the results reported herein.

Several aspects of network dynamics need to be studied further, which cannot at present be attempted due to insufficient data. These include the pre-trip route selection behavior of the drivers, precise modelling of their travel near the destination etc. Dynamic models need to be developed for the drivers' learning behavior and day-to-day adjustment of route and departuref times as well as their en-route decisions without information. Similarly more accurate origin-destination matrices need to be developed.

Chapter 5

NETWORK EQUILIBRIUM UNDER INFORMATION

One important concern in evaluating urban traffic networks under information supply is the dynamics of the effectiveness of information, namely the evolution of the driving patterns in the network under information. The utility that each driver derives from driving on a certain route under information may be different from the utility from the route he/she was driving on when route information was not available. Thus information supply may induce perturbations in the traffic system, causing it to go through an evolutionary period when the drivers adjust their departure time and route decisions till they possibly reach a state when they do not make adjustments and the traffic system is under 'equilibrium'. Thus, the advantages that information supply provides is more meaningful when analyzed from the perspective of the equilibration characteristics of the system. This chapter discusses this aspect, focusing mainly on how a system under equilibrium without information reequilibrates when information is provided and how sensitive the equilibration is to the utility function assumed.

The multiple-facility-single-destination modelling program is used for this study, due to the ease with which discrete alternatives for driver decisions can be defined in this model and the very significant computational advantage it offers over the general network model. Section 5.1 provides a brief background review that complements the discussion in chapter 1 (section 1.3) on this topic. Section 5.2 explains the methodology (iterative simulations). Section 5.3 discusses the results.

5.1. BACKGROUND REVIEW

An introduction of some of the research efforts on modelling in-vehicle information systems based on network assignment was given in chapter 1 (section 1.3). A detailed discussion of the underlying assumptions of network assignment is presented here.

The traditional network assignment theory as it developed in the late 60's based on earlier efforts by Wardrop (1952) and Beckmann (1956), was concerned with finding link flows in traffic networks under static conditions with known traffic demands among the different origins and destinations. Finding the link flows is approached as an assignment problem where flows are assigned to the network links under optimality or equilibrium conditions. Optimality means that the total trip time cost based on flow-dependent link costs is minimized subject to flow conservation constraints, which results in the best total cost in the system. Under such assignment, the each drivers may not be achieving the optimal route travel costs from his/her individual point of view. As a result, they may change their routes (unless their route choices are controlled externally) and the network may not be under equilibrium. When the drivers 'cannot unilaterally change routes and achieve a better *route trip time cost*', as defined by Wardrop (1952), the network is under 'user equilibrium' (UE). Finding the network flows under user equilibrium is also an assignment problem similar to the optimal assignment problem, but with a different equivalent objective function (Beckmann, 1956) that can be proved to result in user equilibrium conditions. While system optimal assignment is important from the point of view of controlling the traffic in a network, user equilibrium is important from the point of view of predicting how the system would (or can) perform on its own in a steady-state.

User equilibrium results in unequal flows on different routes (between an O-D pair) whose trip times are the same at equilibrium, and there is no apparent intuitive reason why different numbers of drivers should drive on routes with the same trip time. This resulted in a further refinement of the equilibrium conditions as the case when 'no driver can unilaterally change routes to gain better *perceived travel time*, which is defined to be a random variable'. The stochastic user equilibrium (SUE) assignment problem based on this definition was first formulated by Fisk (1980) and Daganzo and Sheffi (1977). The perceived trip time on a particular route is considered a random variable T_r (over the driver population) such that $T_r = t_r + \epsilon_r$, where t_r is the deterministic component (actual trip time cost) and ϵ_r is a random part. If ϵ_r is a random variable that is independently and identically Gumbel[62] distributed over the drivers and the routes, the demand using the different route alternatives will be in a logit form, which is,

$$d_r = D \frac{e^{-t_r}}{\sum_{all\ r} e^{-t_r}} \qquad 5.1$$

where, d_r is the demand on a particular route between an O-D pair for which the total demand is D. Based on the distribution assumed for the random term, other

[62] The cumulative Gumbel distribution is given by
$F(\omega) = Pr(\epsilon \leq \omega) = e^{-e^{-\omega+E}}$,

where E is the Euler's constant (E = 0.5708....)

forms of demand division are possible, such as the probit form that results from assuming a Normal distribution. The average trip time on each route is a function of the traffic flow on that route, and it has been shown that the actual route flows can be found using mathematical programming problem, using an equivalent objective function. For the logit case, Fisk developed the first solution method (1980). For the probit case, the problem was formulated by Sheffi and Powell (1982). Solution techniques can be found in Sheffi (1985). This usually results in unequal flows on different routes between the same O-D pair. The drivers *perceive* the routes to have the same trip times, but in reality the routes have different average trip times. This explains the different flows on these routes, thus removing the non-intuitiveness of the original user equilibrium, which was mentioned in the last paragraph. The original user equilibrium is in fact a special case of the stochastic user equilibrium. When the variance of the perceived travel time is zero, SUE results in the original UE.

One main drawback of the static assignment analysis is that the O-D demands and the resulting link flows are considered not to vary in time, which is hardly a correct assumption in most urban contexts. If the time-varying demand pattern is known, then the problem is to find the link flows during different time periods that result in the optimal system costs or equilibrium conditions. This is called the dynamic network assignment problem. The system optimal dynamic assignment problem was first formulated by Merchant and Nemhauser (1978). Carey (1987) has explored the convexity and solvability of the problem under certain linearity assumptions.

The Dynamic User Equilibrium (DUE) problem is more complicated than the static problem, as alternative definitions are possible for equilibrium conditions. For instance, dynamic equilibrium can be based on just the departure time decisions of the drivers in idealized cases where the route choice is not of concern (Cosslet, 1977; Abkowitz,1981; De Palma et al, 1983). A more realistic case is when the routes and the departure times are both selected by the drivers. Mahmassani and Herman (1984) analyze the dynamic equilibrium in a two-route network using link traffic flow relationships. One definition of such a dynamic user equilibrium would be that 'no driver can unilaterally change his/her route *at departure* to obtain better trip times'. This is the definition used by some of the researchers. For instance, Janson (1990) considers the case of equilibrium route choice based on trip times under time-dependent O-D demand and Mahmassani and Herman (1984) consider the route *and* departure time choice at equilibrium based on utilities. An alternative definition is provided by Wie et al (1989) and Ran and Shimazaki (1990) which is that 'no driver can unilaterally change his route *at any instant during the trip* to obtain better trip times' (departure time choice was not considered by them under the assumption of fixed demand variation). Such a definition implies that the problem is in terms of variables that are continuous time-based functions requiring optimal control formulations (which are discrete-time formulations of the same problem. As should be clear by now, there is no consensus yet on what kind of equilibrium really exists in networks (if such equilibria do exist) and hardly any empirical data has been collected on this aspect.

The intent in this chapter is to study a Stochastic Dynamic User Equilibrium (SDUE), meaning that 'no driver can unilaterally change his departure time or route to obtain better *perceived utility*', which is also the definition used by Ben-Akiva et al (1986). It can be seen that the stochastic equilibrium here is slightly different from the original definition of Sheffi and Daganzo (1977), in that it is based on the perceived utility instead of just trip time. The reasons for basing this on utility are as follows. If it is based on just the travel times, the drivers can leave early or late when there is no congestion and we obtain unrealistic solutions. Thus, we need to base it on a utility function involving the schedule delay (earliness or lateness on arrival). In addition, under information, the perceived trip time is the same as the actual trip time, or the variance of the perceived trip times over the drivers (for a particular route at a particular time) is very small. On the other hand, the utility perceived by the drivers is still a random variable, and thus a stochastic equilibrium based on utility is more meaningful under information. The equilibrium is assumed to be based on the perceived utilities *at departure*, which means that the route decisions en-route are not based on the utilities. Still, when the departure decisions are at equilibrium, the en-route decisions by the drivers (based on trip times alone) are deterministic and thus equilibrium conditions do prevail during the trip. Next, the methodology adopted for equilibrium analysis is discussed, followed by some more comments on the underlying assumptions.

5.2. METHODOLOGY

The driver departures are assumed to be based on a selection from among a few discrete route and departure time combinations using the utility that can be derived from the selection. It is assumed that there are N_T departure time alternatives and N_R route alternatives, thus forming $N_T \cdot N_R$ combinations for the driver to choose from. The utility U_{tri} perceived by driver i from a trip starting at departure alternative $\{t,r\}$ where t is one of the time alternatives and r is one of the route alternatives, is composed of two terms.

$$U_{tri} = V_{tri} + e_{tri} \tag{5.2}$$

where, V_{tri} is the deterministic (systematic) part of the utility for driver i on selecting the departure alternative $\{t,r\}$. In general, this could be different for different drivers due the systematic attributes such as location, age, income etc., but only identical drivers are assumed in this study and hence this part is the same (V_{tr}) over all drivers selecting the alternative $\{t,r\}$. e_{tri} is the random part that is different for different drivers, departure times and routes.

The division of the total traffic demand into various fractions choosing each combination at equilibrium, is based on their V_{tr} value, and the distribution of the random term. The random term is assumed to be independently and identically Gumbel distributed (see footnote 1 of section 5.2) here, resulting in a logit demand division,

$$d_{tr} = D\frac{e^{-V_{tr}}}{\sum_{all\ t,r} e^{-V_{tr}}} \tag{5.3}$$

where, D is the total demand and d_{tr} is the demand choosing departure alternative $\{t,r\}$. The implicit assumption here is that the departure choice is independent of the route choice[63].

The next three sections describe (1) the network studied, (2) the function used for systematic utility and (3) the iterative technique used to find the equilibrium, in that order.

5.2.1 The case-study: Network and Experimental set-up

The corridor network with three parallel highways is modelled in this study using the efficient corridor program described in section 3.8. The selection of such a case-study is based on two reasons: (1) The fast path-processing capabilities of the parallel-facility-single-destination simulation program helps in carrying out the iterative simulations within the available computational resources and (2) the network structure with independent routes of similar lengths is preferable for the assumption that the random terms of the perceived utility are independently and identically distributed (IID).

The network and the assumed demand are the same as those described in Chapter 4 (sections 4.2.4.1 and 4.3.1). The three highways in the network are each of 9 mile length and have free-flow speeds of 55, 45 and 35 mph. All three lead to the same destination, with the 6 miles farthest from the destination generating driving demand at 1600 veh/mile over the peak period of 80 minutes studied. All drivers are assumed to have identical work start-times (time = 80).

The peak-period is divided into 16 discrete departure time alternatives, each of 5 minute length. The drivers choose from any of the three routes, and thus

---

[63] Thus, we assume a 'joint logit' structure. Of course, it is possible to relax this assumption by using a 'nested logit' structure with certain assumptions on the dependence between the two dimensions of choice, namely the route choice and the departure time choice.

for each driver there are 16 $\times$ 3 = 48 departure alternatives, each with a certain systematic utility (as explained in the next section), which is the average utility derived by all the drivers that select that alternative. The demand in each of the one-mile sectors of the corridor farthest from the destination are assigned to 48 alternatives. Certain sensitivity studies were also carried out on the utility function (see next section) by adjusting the coefficient of early-side schedule delay. Two other values, 0.00084 and 0.00063 were tried in addition to the calibrated value of 0.00042.

5.2.2 Utility Function

The systematic utility function used for this study is a modified version of one calibrated by Hendrickson et al (1984) based on data from commuters in the Pittsburgh area. The original calibrated utility function included certain terms which were meaningful in the Pittsburgh study, even though Hendrickson et al did not find them statistically significant. After deleting those terms which are not pertinent in this study, we get the following utility function that was used for the iterative simulations.

$$V_{tr} = -0.021\,TT_r - 0.021\,CT_{tr} - 0.00042\,(SDE_{tr})^2 - 0.148\,SDL_{tr} + 0.014\,(SDL_{tr})^2 \qquad 5.4$$

where, for trip time alternative t and route r,

V_{tr} = Systematic Utility

TT_r = Free-flow travel time on route r.

CT_{tr} = Additional trip time (over free-flow time) due to congestion.

SDE_{tr} = Early-side schedule delay = max((Work start time - Arrival time),0)

SDL_{tr} = Late-side schedule delay = max((Arrival time - Work start time),0)

This utility function is simple and includes variables that are directly obtainable from the simulation, making it a suitable one for the current purpose.

5.2.3 Iterative equilibration technique

The intent is to find a Stochastic DUE departure pattern. The total traffic demand from each one-mile sector gets divided into demands selecting the different departure time and route alternatives with a logit split according to the utilities (based on the utility function specified in the last section). The utilities are the result of the traffic conditions caused by the demand split too, which imply that we need to find a solution where the utilities assumed for a certain demand split matches the utilities that result from such a traffic pattern. Of course, a logit split does result in SDUE as explained in the last section. An iterative technique is used to achieve such a departure pattern (i.e., demand splits). The procedure starts with a set of assumed utilities for each departure alternative and assigns the traffic accordingly. From the simulation, a new set of utilities are obtained, which is used to assign the traffic for the next iteration. Equilibrium is reached when the assigned traffic pattern in two iterations are the same or are sufficiently close.

A smoothing procedure is required to achieve a faster convergence as well as to ensure the convergence to one solution, as there is no quarantee that there is a unique solution. The procedure used here is called 'the method of successive averages' (MSA). This method is usually suggested for stochastic user equilibrium assignment for an iterative search along a descent direction for the minimum of the equivalent objective function (see Sheffi, 1985. pp. 326-340). Even though the problem here is different in that it is not an equivalent optimization, but is rather a search process, the nature of the problem renders it particularly suited for this kind of a smoothing procedure. The technique is as follows.

Given: Free-flow trip times on the network routes (these can be found by static addition of free-flow trip times on the links, which in turn depend only on link-lengths and free-flow speeds on the links)

Step 0. Iteration index, $n = 0$. Assign the volumes selecting each departure alternative at iteration 0 according to a logit split (see Eqn. 5.3) based on free-flow utilities. $\bar{x}_n = \{x_{rt} \forall r,t\}_n$ are the set of volumes assigned during this iteration.

Step 1. Perform the simulation using $\bar{x}_n$ and find the actual utilities of the departure alternatives. Calculate the volume splits, $\bar{y}_n = \{y_{rt} \forall r,t\}_n$ according to these utilities

Step 2. The volumes assigned for the next iteration are,

$$\bar{x}_{n+1} = \bar{x}_n + \frac{1}{n}\bar{y}_n$$

Step 3. If $|y_{rt}-x_{rt}| \leq N, \ \forall\, x_{rt} \in \bar{x}_n,\ y_{rt} \in \bar{y}_n$, STOP ($N$ = say, 2).

Else, $n = n+1$ and go to step 1.

Comments are in order here about the methodology. First, there is no guarantee of a unique equilibrium state. In fact, it seems quite probable that multiple equilibria exist. The main reason for this is that the utility function is not convex with respect to the volumes. It may be noted that static stochastic equilibration problems formulated in terms of link flows do have unique equilibria, due to the facts that the formulation is in terms of link flows and that the trip times are convex with respect to the link flows in the static case (see Sheffi, 1985. pp. 318-322). However, the present problem is in a dynamic setting with utilities that are not convex with respect to the route traffic demands and are undefined with respect to the link traffic flows. In addition to this, the enroute driver route-switch decisions affect the 'smoothness' of the 'utility space' resulting from the demands, meaning that the utilities from two almost-identical departure patterns could be quite different, thus causing added non-convexity to the problem.

Due to the possible existence of multiple equilibria, the method of successive averages is only a way of 'forcing' the system to converge to one of the equilibrium states. If instead of using a smoothing procedure, a behavioral rule is used for the day-to-day decisions and adjustments of route and departure times by the drivers, we could get traffic volumes from iteration to iteration that show the actual states that a system may go through under information, while equilibrating. This would be similar to the case studied by Mahmassani and Chang (1986) under boundedly-rational user equilibrium conditions and the stochastic demand adjustment mechanism (though it is based on MSA and some relatively simple learning rules) studied by Vythoulkas (1990), both for day-to-day adjustments of driver decisions in networks without information. In future, when data becomes available on such driver behavior (mainly of learning) under information, it will be possible to study the systems evolution towards convergence. It may also be noted that networks need not have unique dynamic equilibrium states in reality either. In fact, some researchers believe this is the case and have proposed dynamic user equilibrium *under external control* (Ran and Shimazaki, 1990). In view of these aspects, it can be seen that the methodology here is to find a possible equilibrium state and to derive insights from that.

It may also be noted that the methodology remains unchanged in the case of reequilibration from one equilibrium state, except that the iterations start with the utilities at first equilibration instead of the utilities based on free-flow speeds. This is used to study the effect of information on a network that is in equilibrium without any information supply. The first equilibrium can be found using the above methodology, and the next equilibration is carried out starting from the equilibrium departure pattern from the first equilibrium, but with a fraction of the drivers making route decisions enroute using information. The implication here is that the equilibrium reached is a function of the initial starting conditions, only in which case these reequilibration studies are sensible. Different equilibrium states do indeed result on reequilibration as the studies show. This raises interesting issues on the stability of networks with equilibrium under information, which are not studied in detail here. A discussion of the results follow in the next section.

5.3. RESULTS

The main results from the iterative equilibration runs are shown in tables 5.1 through 5.4. Table 5.1 shows the results of equilibration of the network with no information supply, with three assumed utility functions. Utility function U1 is the function calibrated by Hendrickson et al (1984) and described in equation 5.4. Utility functions U2 and U3 are two modified utility functions with the absolute value of the coefficient of SDE (Early-side Schedule delay) increased to 0.00063 and 0.00084 respectively from the 0.00042 value of U1.

Iterations	Original Utility Function (U1)	Modified Utility Function (U2)	Modified Utility Function (U3)
0	-4267.3	-4710.6	-4871.0
1	-6136.6	-6928.2	-7349.7
2	-6045.7	-7045.0	-7852.4
3	-6058.1	-7019.3	-7853.8
4	-6072.4	-7009.4	-7881.1
5	-6076.9	-7016.4	-8095.2
6	-6077.3	-7020.8	-8133.8
7	-6077.1	-7025.6	-8091.8
8	-6077.3	-7030.1	-8148.1
9		-7031.7	-8083.1
10		-7031.3	-8042.6
11		-7032.0	-8152.7
12		-7031.2	-8131.9
13		-7032.3	-8156.7
14		-7031.3	-8172.3
15		-7032.3	-8136.6
16		-7031.3	-8152.9
17		-7032.3	-8144.7
18		-7031.3	-8144.8
Avg Trip time	10.01 min	10.67 min	12.08 min
Avg. Schedule delay (early)	22.51 min	19.11 min	16.81 min
Avg. Schedule delay (late)	0.49 min	0.72 min	1.12 min

Table 5.1 Variation of system utility during equilibration and the average values of schedule delay and trip time at equilibrium.

	IB = 0.0, INF = 0.5	IB = 0.0 INF = 1.0	IB = 0.2 INF = 0.5	IB = 0.2 INF = 1.0
Total system utility at reequilibration under information	-5919.6	-6052.7	-5881.7	-5883.3
Equilibrium Avg Trip time (Overall)	9.66	9.88	9.58	8.91
„ (group with info)	9.50	9.88	9.40	8.91
„ (group without info)	9.82	--	9.78	--
Equilibrium Avg Earlyside Schedule Delay (Overall)	22.35	22.29	22.24	22.25
„ (group with info and early)	26.06	25.94	25.52	25.77
„ (group without info and early)	25.71	--	25.91	--
Equilibrium Avg lateside Schedule Delay (Overall)	0.47	0.52	0.48	0.48
„ (group with info and late)	3.38	3.68	3.46	3.50
„ (group without info and late)	3.55	--	3.59	--

Notes: All times are in minutes. IB = Average Indifference Threshold Fraction. INF = Fraction of Drivers with Information. Total Equilibrium System Utility at the start = -6077.3

Table 5.2. Reequilibration under information from a no-information equilibrium state. Utility function : U1 (Original calibrated function. Coefficient of SDE term = 0.00042)

	IB = 0.0, INF = 0.5	IB = 0.0 INF = 1.0	IB = 0.2 INF = 0.5	IB = 0.2 INF = 1.0
Total system utility at reequilibration under information	-7592.8	-7618.0	-6799.6	-6807.7
Equilibrium Avg Trip time (Overall)	11.03	11.40	10.25	10.15
„ (group with info)	11.04	11.40	9.97	10.15
„ (group without info)	11.02	--	10.53	--
Equilibrium Avg Earlyside Schedule Delay (Overall)	19.57	19.39	18.67	18.69
„ (group with info and early)	24.36	23.63	22.86	22.52
„ (group without info and early)	23.58	--	22.21	--
Equilibrium Avg lateside Schedule Delay (Overall)	0.96	0.95	0.71	0.68
„ (group with info and late)	5.41	5.32	4.15	3.96
„ (group without info and late)	4.99	--	4.17	--

Notes: All times are in minutes. IB = Average Indifference Threshold Fraction. INF = Fraction of Drivers with Information. Total Equilibrium System Utility at the start = -7031.3

Table 5.3. Reequilibration under information from a no-information equilibrium state. Utility function : U2 (Modified from the calibrated function. Coefficient of SDE term = 0.00063)

	IB = 0.0, INF = 0.5	IB = 0.0 INF = 1.0	IB = 0.2 INF = 0.5	IB = 0.2 INF = 1.0
Total system utility at reequilibration under information	-8754.0	-9212.0	-7728.2	-7955.3
Equilibrium Avg Trip time (Overall)	12.40	13.20	11.26	11.40
„ (group with info)	12.64	13.20	11.03	11.40
„ (group without info)	12.15	--	11.49	--
Equilibrium Avg Earlyside Schedule Delay (Overall)	17.48	17.65	16.90	17.02
„ (group with info and early)	22.69	22.81	21.12	21.10
„ (group without info and early)	21.23	--	20.81	--
Equilibrium Avg lateside Schedule Delay (Overall)	1.23	1.46	0.93	1.05
„ (group with info and late)	6.15	6.46	4.91	5.44
„ (group without info and late)	5.93	--	4.69	--

Notes: All times are in minutes. IB = Average Indifference Threshold Fraction. INF = Fraction of Drivers with Information. Total Equilibrium System Utility at the start = -8147.4

Table 5.4. Reequilibration under information from a no-information equilibrium state. Utility function : U3 (Modified from the calibrated function. Coefficient of SDE term = 0.00084)

Table 5.1 shows the sensitivity of the equilibration process to the coefficient of the early-side schedule delay. This is an important variable in the utility function in that it captures the trade-off between the trip time and the schedule delay. A lower absolute value for this coefficient means that a departure alternative that results in lower trip time but higher schedule delay is preferable to the driver. This is clear from the results, as the average trip times at equilibrium indeed increase when the coefficient is low. The fact that about 20 % increase in trip time would be acceptable to the drivers when the absolute value of this coefficient increases to 0.00084 from 0.00042, clearly shows the importance of evaluating trip time improvements with information strategies with proper accounting for the schedule delay increase that might accompany such improvements. The results also point to the necessity to undertake the empirical research for calibrating accurate utility functions.

Tables 5.2 to 5.4 show the results of the network reequilibrating under information, starting from equilibrium states (which are reached after the iterations shown in table 5.1). In the case of all the three utility functions and both the values of average indifference threshold fraction (IB = 0, for the myopic switching and IB = 0.2, for the more conservative boundedly rational switching), we see that the utility at reequilibration under information is better when a smaller fraction of people are equipped for information. In line with some of the conclusions from the corridor simulations reported in chapter 4, myopic switching (IB = 0) can be seen to cause reequilibration to a worse total utility state, as compared to the equilibrium without information.

The reason for the worsening of the total system utility function on reequilibration is clearly the increase in early-side schedule delay and the decrease in the trip time. Note that the system utility is is lower even though th etrip times are lower, due to the higher schedule delay. In most cases the increase in schedule delay is rather small, but the resulting worsening of the utility is considerable. This is a direct result of the form of the utility function, where the disutility increases quadratically with respect to the early side schedule delay. This again, points to the need for reliable utility functions.

Another interesting observation is that the schedule delays are in general higher for the drivers receiving information. This could be because these drivers achieve lower trip times thanks to the information but end up in an equilibrium with lower trip times (and worse utilities) as the mechanism here that causes departure time adjustments that are in the descent direction of the disutility (name ly the logit split that maximizes the random utility) may be acheiving a local minimum. This may not be the case in reality, as the drivers could change departure times according to different adjustment mechanisms (i.e., learning behavior) and reach different equilibrium states. Of course, if they make boundedly-rational departure choice adjustments (see Section 2.3.1) as opposed to optimizing decisions, they may accept the lower utilities (see Mahmassani and Jayakrishnan, 1988, for simulation results on similar trade-offs in the day-to-day adjustments of drivers in the context of road-closure perturbations), and reach a 'worse' equilibrium state.

Due to the larger coefficients for the late-side schedule delay in the utility function, the average schedule-delay on the late-side is quite small in magnitude and as such does not seem to provide much meaningful insights with regard to the trade-offs with trip time and early-side schedule delay. This aspect can be studied by changing its coefficient in a sensible manner, which however is not attempted here. As the information supply does not increase the average trip times in general (according to the results of chapter 4), the late-side schedule delay should anyway not be expected to increase under information either.

5.4. CONCLUSIONS

This chapter examined the effects of information on the equilibration characteristics of a specific network. It is interesting to find that based on a utility function of the drivers that includes schedule delay also along with the trip time, there are indeed cases where information causes the system to reach equilibrium states with worse utility (lower trip times but higher earliness). This of course depends on the utility function that is assumed. Further research is needed to develop and calibrate realistic utility functions so that such analyses can be performed to find the impacts of in-vehicle information systems. This aspect has not been given enough emphasis in the current plans for implementation of such systems to gain maximum trip time advantages. Based on the results from this chapter it appears essential that this issue is addressed carefully in future.

Chapter 6

CONCLUSION

This chapter provides concluding remarks on the research effort reported in this thesis. The chapter starts with a summary and discussion of the overall conclusions in section 6.1; Section 6.2 discusses the author's views on the various aspects of the work that constitutes contributions to the state of the art. The last section discusses the possible areas in which further modifications can be incorporated to the simulation framework as well as the future research directions related to this topic.

6.1. OVERALL CONCLUSIONS

Due to the nature of the simulation-based analysis, the conclusions derived here are mostly qualitative. The results are based on a relatively few cases too. Nonetheless, some of the insights developed may be generalizable and are valuable. Only a few networks were studied, but they were carefully selected to explore the applicability of the framework to evaluate realistically large traffic networks under information and also to study the fundamental aspects of traffic dynamics under information. As various information scenarios are studied on two different types of networks, one with a specific corridor structures and the other with a general multiple-destination structure, the conclusions may be applicable to a variety of network scenarios. At the same time, as may be expected in this type of research where only a few representative cases are studied, some of the results which are apparently conflicting may need to be correctly interpreted.

An important conclusion from the simulation studies is that information supply in traffic networks results in performance improvements that are not necessarily directly proportional to the 'level' of information supply. For instance, the overall network travel times appear to reduce sharply when only small fractions of vehicles are equipped for information supply, but as the percentage of the equipped vehicles increases, we get decreasing marginal reduction in the system-level total trip times. It was found that only less than 50% (sometimes less than 30 %) of the drivers need to be equipped to receive information, for the system to achieve more than about 80 % of the advantages. This pattern is seen in all of the simulations carried out, regardless of the traffic loading patterns or the assumed levels of driver propensity to switch routes using the information, even though the general network cases did not show as much of advantages with low fraction of vehicles as the corridor cases (mostly due to the nature of the route assignments of uninformed drivers, as discussed in Chapter 4).

The results do confirm that by and large the information supply does provide benefits in terms of system-level trip times. This is intuitively appealing and was generally assumed to be the case when the demonstration and/or implementation efforts on route guidance and information systems were initiated around the world, but nevertheless was rarely confirmed with extensive realistic analyses. The improvements in total system travel times are generally less than 10 percent in comparison with the case of no drivers being equipped for information. While this improvement may not appear to be very significant, it should be noted that it is more than what can be achieved with any of the existing conventional traffic engineering approaches in heavily congested networks.

The next conclusion is about the effect of information on the two major groups of drivers in a network: those who are equipped for information and those who are not. The results show that, in general, the drivers without information are not adversely affected by information supply to the other group of drivers. On the contrary, they received benefits in all the cases simulated and in some cases even obtained more benefits than those receiving information.

The observations regarding the impacts on the two groups of drivers (with and without the equipment to receive route information) for various levels of information supply are also meaningful. It is seen that in most cases, the average trip time of the drivers without information increases with increasing fractions of drivers with information. On the other hand, the average trip time benefits derived by equipped drivers decrease as the fraction of equipped drivers increase.

Perhaps the most interesting and probably somewhat controversial conclusion from these simulations is that there indeed could be cases where supplying information can be detrimental to the system performance. This was found in the simulations of the 3-highway corridor with 75 to 100 percent of the drivers equipped and the drivers making myopic decisions in route switching (always switching to the best path displayed). Of course, such worsening of performance may not always occur, as observed from the limited cases simulated for the general Austin network. The reason could be that corridor networks with fewer facilities get congested easily due to over-reaction of the drivers while the general networks offer more opportunities to disperse the traffic. While further simulations are needed to confirm the definite reasons behind such a worsening of performance, the fact that it did occur in some cases point to the necessity for careful analysis when information systems are put in place.

The observations discussed in the above two paragraphs raise important questions such as why only a certain fraction of drivers should pay for information equipment while their effort is only helping the fraction of drivers who are not paying anything. Such equity issues, coupled with the fact that most of the systemwide benefits can be achieved with low market penetration raise interesting challenges in designing cost-effective information supply strategies which are socially acceptable and equitable. This is an important aspect which should be addressed in connection with the current plans to implement such information systems around the world.

The network equilibration studies reported in chapter 5 add a new dimension to the conclusions on the system performance under information, namely, the evolution of the system towards a (stochastic dynamic) user equilibrium. When the utility function of the drivers includes a penalty for early or late arrivals at the work place, the performance of the system cannot be correctly summarized or meaningfully evaluated with just the trip time statistics because the schedule delays (early or late) associated with a particular departure time are also of significance. Thus performance comparison on the basis of total utilities is probably more reasonable.

We see that while the system equilibrates under information (or more correctly, reequilibrates from an equilibrium under no information) resulting in better trip times for higher switching propensity of the drivers, the total system utilities are lower. This is due to the higher schedule delays associated with such lower trip times. Again, we see that conservative route-switching with an indifference threshold (towards minimal advantages) provides the best system performance in terms of utility, at equilibrium. Of course, further simulations with more appropriate utility functions (which may be calibrated specifically for systems under information) are necessary for arriving at more general conclusions.

6.2. RESEARCH CONTRIBUTIONS

The research undertaken here is among the first, to the author's knowledge, that incorporates the three essential elements that influence the dynamics of urban traffic networks under information, namely (1) the traffic flow dynamics, (2) the driver route-selection behavior and (3) the dynamic variation of path characteristics under information supply. The most significant contribution of the research is the development of a comprehensive modelling framework that integrates the components that model these three dynamic elements. In the past, the complexity of the interplay among these elements and the computational requirements in tying them together into a modelling system that can be applied to large-scale networks have prevented the development of such tools.

The utility of the simulation framework is two-fold. First, the framework is capable of modelling relatively large networks under information supply, and can be used to design or evaluate alternative information strategies before their implementation on existing city networks. This would help in limiting the need for large-scale real-world experimentation which may be prohibitively expensive. Second, the simulation framework could conceivably be a part of a real-time route guidance/information system, where the information supply strategy is determined by real-time simulation (or, quasi real-time simulation with stored scenarios). The framework, with a few enhancements, is applicable for the former purpose while its use for the latter purpose requires development of detailed interfacing capabilities depending on the particular real-time system in which it is used. In either case, the framework can be expected to be reasonably helpful in performance prediction.

Another contribution is the highlighting of a significant determinant of the performance of information-driven networks, namely the driver behavior that is sometimes equivalent to non-compliance with the information provided. The interesting results under various levels of route-switching propensity show the need for considering this aspect while evaluating information systems. It is notable that the actual levels of switching propensity will depend on the particular network context, but the results from simulations using this framework can be used to decide what kind of route-switching propensity is 'best' for a particular network case. Such insights may even be used for driver education regarding switching behavior that contributes towards user or system optimality. One of the strong conclusions from this research is that such driver behavior of not switching to an alternative route which the information system displays as advantageous, must be included in any evaluative modelling framework: one aspect that has hardly been considered in the past.

The development of a framework that is adaptable to future improvements in the state of the art in this area is also a contribution of the research. The framework has a modular structure so that better models for driver behavior, initial path selection, traffic flow dynamics etc can be incorporated to it with minimal effort. This is in fact one of the reasons why this framework is being strongly considered as a major component of a significant research effort that the Federal Highway Administration has initiated at the University of Texas dealing with optimal route guidance in urban traffic networks.

This framework is among the first in the area of large-scale traffic network modelling that utilizes a K-shortest path approach which, as explained in Chapter 2, has a number of merits than the conventional shortest path modelling. This is an important component of the framework that lends it the flexibility for adaptation to the future advances in driver behavior and route guidance strategies. The K-shortest path routine is efficiently written and uses efficient data structures such as binary heaps, so that the simulations can be done with reasonable computational effort, even with a relatively large number of paths between each origin and destination in the network. This would prove very useful in future studies of information strategies such as, for instance, routing to suboptimal paths and driver selection from choice sets with multiple alternative routes.

The next contribution is that the research developed a modelling framework that efficiently utilizes the recent computational advances, specifically the supercomputing capabilities. Path-processing is the most computationally demanding component of the framework. The heap-based K-shortest path, while it is not very efficient in vectorization, is developed in a manner that is suitable for multitasking. Efficient routines have been developed for the aggregation of link trip times to route trip times. Separate routines, sequential and vectorizable, have been developed for this routine so that the framework can be efficiently implemented on either computational environment.

Several new substantive insights have been attained from the simulation studies providing an intuitive understanding of the dynamics of information-driven networks. This will be valuable in the future development of simulation experiments as well as field studies to develop calibrated models of component phenomena for incorporation into the framework.

6.3. FUTURE RESEARCH

Modelling of traffic networks under information is a rapidly emerging area. Considering the demonstration and implementation efforts envisaged now, and the associated fundamental research efforts initiated in the various related areas, the state of the art is expected to improve dramatically in the near future. A discussion that places the research reported here in the context of such developments is warranted along with pointers to the kind of research that is important.

Development and calibration of behavioral models of driver response to information is of primary significance, because it is among the most important determinants of the performance of information-driven networks. For such models to be incorporated into this simulation framework in a useful manner, they may have to be disaggreagate (individual-level) in nature. Models have to be developed and calibrated in two broad areas: 1) the pre-trip decisions and 2) the en-route decisions.

It has been found in recent research (Mahmassani and Chen, 1991) that pre-trip decisions may play a very significant part in determining the system performance. In this research itself, the lack of a realistic pre-trip decision component may have caused shortcomings, especially in the selection of reasonable initial departure patterns and the simulation of realistic base-cases with no enroute information supply. For instance in the Austin network simulations, the reason that myopic enroute switching behavior produced the best system results may be hypothesized to be the lack of a better assignment of the initial paths of the drivers (which may have caused it to be always better to have an initial change of routes for some significant benefits).

Development of such models of pre-trip route selection is even more important in the context of the dynamic day-to-day evolution of the traffic systems under information as they provide a means of carrying out iterative simulations to model realistic medium or long-term equilibrium (if they indeed exist) in traffic networks. A few aspects of importance in such models would be, 1) the nature and availability of pre-trip information, 2) the driver knowledge about the average conditions on the routes in the network, 3) the relative usefulness of the information, in terms of switching opportunities, on each route and 4) the reliability of the information supply.

In the case of enroute decisions, alternative forms of route switching and/or selection models need to be developed. The possible aspects of driver behavior which have to be studied are, 1) enroute habit formation (such as liking their current route) by the drivers, 2) imaptient behavior when their personal goals regarding acceptable values of various route attributes are not met (especially when route-switch opportunities are rare) and 3) perception of the reliability of the enroute information supply.

There are two possible approaches for collecting the data required to develop such behavioral models: 1) in the field, by means of recording the driver decisions and the associated circumstances, or 2) simulation-based laboratory experiments, with real people making driver decisions in simulated traffic systems. The field data collection method could be much more expensive than the laboratory approach, as the information supply system has to be in place and the data requirements for statistically calibrating any kind of reasonable model would require equipping a large number of drivers with on-board display systems. Furthermore, it may be impossible to study alternative guidance strategies (such as displaying more paths than the best, for example) without extensive hardware restructuring. On the other hand, in a laboratory setting, which would be considerably less expensive and more flexible in studying alternative scenarios, the drivers could be behave in a 'relaxed' manner without the pressures of real-life driving, resulting in relatively less reliable models.

On the traffic simulation side, one possible improvement may be the development of codes for parallel computing environments. As Connection Machines with thousands of processors are now available, the traffic movement in all the network links may be carried out in a simultaneous manner by assigning one processor for simulating each link. Alternatively, one processor may be assigned to each vehicle in the network. The first method would require fewer processors but would involve more inter-processor communication than the other. Extensive rewriting of the code and possible change of the programming language from FORTRAN to C may be necessary to achieve such parallelization.

One current deficiency in the framework is that the vehicle delays at intersections and interchanges are not modelled in detail. Additional modules may have to be added to model pre-timed and/or actuated signals as well as simple stop control at the network intersections. Research is already underway at the University of Texas and at the University of California, Irvine for incorporating such intersection modelling capabilities into this framework.

From the information supply perspective, the framework at present does not model the display of route information based on predicted travel times. For the existing and planned network information systems around the world, this component is only under development. When such travel time prediction models become available, they will have to be incorporated to the framework, for it to be able to model an entirely new array of information supply strategies.

One possible approach at this point is to perform the vehicle movement in each simulation time step based on simulations over a specified period in the immediate future which provides link trip time profiles into the future. This would require a modification in the route trip time aggregation routine to add up trip times dynamically (selecting appropriate value from the link trip time profiles). Another possibility is to approach this problem from a probabilistic point of view with variables capturing route switching probabilities. The profile of the turning probabilities at each node in the past can be extrapolated into the future to calculate the link concentrations and the resulting trip times in the future. Both of these approaches could prove to be rather involved and many additional details of dynamic prediction of travel times will have to be carefully addressed.

As the developments related to traffic network information/guidance systems are still in the arly stages, many areas of future reserach can be identified. There are many avenues of exciting, innovative and challenging research in modelling and analyzing networks under information. The brief account above is not meant to be exhaustive, but it is hoped that it is useful. In the same vein, the entire research effort reported in this thesis is not expected to be the final one on this topic, but rather an earnest beginning at developing a useful tool to study a relatively complex system of profound significance to the quality urban life.