This is a repository copy of *Route Guidance Algorithms Effective for All Levels of Take-Up and Congestion.*.

White Rose Research Online URL for this paper:
http://eprints.whiterose.ac.uk/2243/

Monograph:
Watling, D. (1990) Route Guidance Algorithms Effective for All Levels of Take-Up and Congestion. Working Paper. Institute of Transport Studies, University of Leeds , Leeds, UK.

Working Paper 315

White Rose Research Online

Institute of Transport Studies
University of Leeds

This is an ITS Working Paper produced and published by the University of Leeds. ITS Working Papers are intended to provide information and encourage discussion on a topic in advance of formal publication. They represent only the views of the authors, and do not necessarily reflect the views or approval of the sponsors.

Published paper
Watling, D. (1990) *Route Guidance Algorithms Effective for All Levels of Take-Up and Congestion.* Institute of Transport Studies, University of Leeds. Working Paper 315

ROUTE GUIDANCE ALGORITHMS EFFECTIVE FOR ALL LEVELS OF TAKE-UP AND CONGESTION

David Watling

ABSTRACT

This paper describes work carried out under the EC `DRIVE' programme, the aim being to develop route guidance strategies which direct users to multiple routes between each origin-destination pair, and thereby provide stable and effective guidance even when a large proportion of drivers are guided.

A model is proposed in which guided and unguided drivers have different route choice assumptions, but are still able to interact with one another; the guidance may be based on either user or system objectives. Conditions are deduced under which the resulting route pattern is guaranteed to exist and be stable. To assess the performance of the strategies, simulations are carried out on two real-life networks, for a number of different demand levels, levels of equipped vehicles, levels of error in (or adherence to) the guidance recommendations, and different guidance criteria. The simulations are extended, in order to examine firstly the influence of behaviour of unguided drivers on the benefits obtained, and secondly the performance of the strategies in cases of unforeseen variations in network conditions. Finally, some comparisons are drawn with a route guidance strategy developed in a parallel `DRIVE' project, where only one route is recommended per origin-destination pair.

CONTENTS   Page

FIGURES: **Page**

1. INTRODUCTION

In this report, work carried out by Leeds under DRIVE V1011 ("CAR-GOES") activity B2.4 ("Methods for stabilisation of route recommendations") is described. The objective of this project was to consider algorithms for route guidance which provide efficient routes, independent of the level of take-up (that is, the proportion of the driver population in the network which has guidance equipment) and congestion levels. The work reported builds on the introduction in deliverable 14 (CAR-GOES, 1990a).

Currently, a number of route guidance systems are undergoing field trials - such as the ALI-SCOUT system in Berlin (Von Tomkewitsch, 1987), and AUTOGUIDE in London (Belcher and Catling, 1987). The current operation of such systems is for a single route to be recommended to equipped drivers for an origin/beacon and destination pair, with a new route chosen in the light of prevailing traffic conditions every, say, 5 minutes. Whilst with few vehicles equipped with guidance devices there will be little change to individual link travel times in the network due to guidance, as the level of take-up increases re-routed traffic will eventually congest the recommended route - at the current rate of route updating - and the benefits of guidance are likely to be lost. A natural reaction would appear to be to increase the frequency at which route information is updated as the level of take-up increases. There are a number of problems with such an approach:

(a) Limitations exist on the speed at which the process can be conducted of relaying information on current traffic conditions to the central guidance system, computing recommended routes and sending the route information to the beacons and thence to the drivers. At higher levels of take up, then, it may not be possible to update routes at a sufficiently frequent rate.

(b) At higher levels of take up/congestion, there is likely to be instability in the route recommendations. For example, when two routes with similar characteristics exist between an origin-destination pair, the recommendations are likely to fluctuate between the two routes in successive time intervals. Such instability in the routing pattern - and hence the travel times - would make the task of signal control a very difficult one, even if it were integrated with the route guidance system, and the benefits of route guidance may be lost.

(c) As more vehicles are equipped with guidance devices, it may be expected that errors in the journey time prediction methods would become smaller - since information on current conditions comes only from equipped vehicles, then as take-up increases so does the "sample size". If, however, routes are updated more frequently as take-up increases, there would be a corresponding decrease in sample size, since the total number of vehicles in each update period would decrease as the period becomes smaller. The end product would be that we would lose some of the greater accuracy in the estimation of current conditions which would normally have arisen with higher levels of take-up.

The work in this report is the first step to developing strategies which advise multiple (≥ 1) routes for each origin-destination pair in each update period. In fact here we

only consider the steady state (fixed demand) situation, and so there will be no specific mention of updating of route information; the assessment of the strategies will be based on their benefits under average conditions. The guidance strategies considered are equilibrium-based, in that they consider the effect of re-routed traffic on link travel times. In this way, by anticipating the magnitude of the change in link travel times due to the route advice given, it would be hoped that greater stability would be introduced into the guidance system.

Having given a review of previous work in this area, the strategies to be considered are described. The basic model is one of a multiple user class equilibrium assignment, where the classes are the unguided and guided drivers, the latter subdivided into those following different guidance advice (either system or user optimal) or those receiving poorer information or with less confidence in the recommendations. The advantage of such an approach is that it is able to take into account interactions between equipped and unequipped vehicles - so not only does the route choice of unguided drivers affect that of guided drivers, but also unguided drivers may (in the long term) choose new routes in response to the behaviour of guided drivers.

The concept of multiple user class assignment is introduced, and relevant theoretical work described. Conditions (on the cost functions) are then deduced, under which a unique, stable equilibrium (for both guided and unguided vehicles) may be guaranteed to exist.

The strategies are then investigated in relation to real-life networks, using an adaptation of the simulation/assignment model SATURN. Care is taken to first model the route choice of unguided drivers in a realistic manner, and the scenarios to be studied are then selected, taking into account the recommendations of deliverable 9 (CAR-GOES, 1990b). Two real-life networks, of differing sizes, are studied under scenarios consisting of a number of different demand levels; various levels of take-up of guidance; difference guidance criteria; and a number of levels of information quality supplied to the guidance system (which could be regarded alternatively as levels of adherence of the guided drivers to the recommendations). The performance of a strategy in each case is measured by network-wide quantities, such as total system travel time and average speed, as well as the benefit (or disbenefit) to individual guided and unguided drivers in terms of the change in their average travel time due to guidance.

Two features of this equilibrium-based model are then studied in more detail, as a form of sensitivity analysis. Firstly, the influence of assumptions regarding the behaviour of unguided drivers is studied, where the simulations are carried out again assuming that unequipped drivers stay on the routes they chose before guidance was in operation. Secondly the basic model (with unguided drivers re-routing) is tested in situations of unforeseen variations in network conditions, with the performance and stability of the routing algorithms studied in conditions of random variations in link capacities, but when the routing is based on average capacities.

Finally, in order to gain some comparison between the above multi-route strategies and single route guidance systems applied to the same network, a study is also carried out with the model developed at Leeds for the DRIVE "ASTERIX" project. This

model, as an extension to SATURN, may be used to assess the effect of using guidance to re-route equipped drivers in response to day-to-day variability in demand. A single route only is recommended, based on either minimum actual cost or minimum marginal cost. In particular, from this study, the level of take-up may be determined at which such a single route strategy is no longer effective.

2. REVIEW

Before proposing the route guidance model to be studied here, a review of similar work will be given. Since this paper does not address the problem of responding to incident congestion, the review will not concern itself with the application of guidance systems in such situations, although it is clearly an area where there is great potential for the use of real time information.

The review will concentrate on aspects of particular relevance to this report, notably (where details are available): the guidance strategy implemented (single or multi-route, user or system based objectives); the model used to assess the strategy (dynamic or static, assignment or simulation); the size and other attributes of the test network used, and the current level of congestion; global benefits and the effect on guided and unguided drivers, and the influence of levels of take-up and congestion.

Kobayashi (1979) proposed a model in which unguided drivers chose routes according to (flow-independent) attributes such as road length, number of lanes and number of left or right turns, whereas equipped drivers were guided according to one of three strategies:

(i) single shortest path guidance between each O-D pair
(ii) guidance onto multiple routes using a smoothed "weighted average" of the current and the previous shortest paths
(iii) an heuristic user equilibrium-like guidance, obtained by an incremental assignment technique.

The strategies were implemented with a (dynamic) simulator based on probabilistic queuing theory. They were tested on a sub-network of Tokyo - in order to estimate the benefits of the CACS route guidance system - consisting of some 99 intersections and 286 directional links; a time-sliced origin-destination matrix was estimated for this area. For a 100% take-up of guidance, the results showed that the strategy reducing total travel time by the greatest amount was (iii) and then (ii). Whilst (i) was also able to provide significant overall travel time savings, it was seen that at higher demand levels during the period modelled, there could actually be a (short-term) disbenefit of such guidance. Interpolating the results for this sub-network and two similar ones, Kobayashi estimated a total travel time saving in the whole of Tokyo of 6%. Kobayashi also examined the effect of level of take-up (at 25%, 50%, 75% and 100% of vehicles equipped) on overall travel time for the sub-network considered previously, and found that the benefit of using strategy (i) or (ii) increased as more vehicles were equipped, up to 75% take-up. After then, however, moving to 100% take-up, the benefit may decrease. Unfortunately, strategy (iii) was not investigated in this way. Finally, from studying a random sample of origin-destination pairs, he found that guidance tended to reduce the variance in travel times.

Tsiji et al (1985) later investigated the test area for the CACS system in Tokyo using a quite different model to that of Kobayashi. They considered the case where only a small proportion of drivers were equipped, and assumed that there would be no effect on the behaviour of unequipped drivers of implementing guidance. Furthermore, they assumed that unequipped drivers would gain no benefit from guidance, and so confined their study to the effect on equipped vehicles. Their model had the advantage of being able to incorporate two sources of random variation - not only in the travel times, in order that a reasonable spread of routes may be obtained, but also in the times yielded by the journey time prediction algorithm (that is, the input to the guidance system). On the other hand, a number of other strong assumptions were made: in particular, that travel times on alternative routes were independent, that there were only two alternative routes between each origin-destination pair, and that each was used only by guided drivers or only by unguided drivers. Extrapolating the results they obtained for the test area, they estimated an 11% reduction in total travel time due to guidance in the whole metropolitan area of Tokyo, at a 10% level of take-up.

Al-Deek et al (1989) studied a portion of the SMART corridor in Los Angeles - which included a freeway and three parallel streets. From an initial survey, they concluded that the "preferred route" for drivers was to enter and leave the corridor on the freeway, and inferred that a suitable indicator of the possible benefit of guidance was the difference between the time on the freeway route and that on the minimum time route (the estimate of the minimum time route being updated over time). These assumptions clearly ignore any effect the re-routed traffic may have on travel times in the network, and the response of and (dis)benefit to unguided drivers. In implementing the strategy on the test network - using TRANSYT as a basis for the simulation - they found that the maximum travel time saving for guided drivers was 3 minutes for a 20-25 minute trip.

Smith and Russam (1989) made use of the heuristic dynamic assignment model CONTRAM (Leonard et al, 1979), in order to estimate the potential benefit of the AUTOGUIDE system in London. They used the London Transportation Studies model, consisting of some 7000 intersections and 16000 one-way links, together with 985 zones. Demand was assumed to be fixed over time on a particular day, and so in this way some of the dynamic element to CONTRAM was lost. The origin-destination matrix provided was assumed, however, only to represent an average demand, with day-to-day variability introduced into the model by a randomisation process. Unguided drivers base their route choice on a stochastic user equilibrium-like assignment for the average origin-destination matrix - that is, the route flows for the assignment of the average matrix are scaled up to conform to the actual (realised) matrix - and do not change their routes in response to the new behaviour of guided drivers. Equipped drivers, on the other hand, are assigned to the minimum cost route for each O-D pair according to the actual O-D matrix and current conditions. That is to say, unguided drivers can only choose routes according to average conditions they have previously encountered (and even then they perceive travel costs differently); guided drivers, on the other hand, are routed according to the conditions specific for that particular day/time. Studying levels of take-up of 10%, 20%, 30% and 100%, it was found that there was a system benefit in terms of reduction of total travel time (2.5%-6.0%) and total distance travelled (0.2%-1.3%) for all cases, whilst the average network speed showed a corresponding increase (2.4%-5.3%).

Disaggregating into the separate groups of drivers, it was found that equipped drivers obtained a reasonably constant 6% reduction in travel time, whilst the unequipped drivers tended to benefit slightly more with higher take-up (benefit 2.2%-3.1%) up to the maximum 30% take-up for which their behaviour was studied.

Breheret et al (1989) made use of CONTRAM in a somewhat different way, in order to study the effects of route guidance in both incident and incident-free scenarios. In the incident-free situation, unguided drivers were assumed to follow minimum perceived cost paths, the perceived cost of a link being a random variable with mean the actual cost (given current traffic conditions) - that is, the routing pattern is something akin to a stochastic user equilibrium assignment. It was assumed that unequipped drivers would not change their routes in response to the new conditions brought about by guidance. The strategy of guiding equipped drivers to actual minimum cost routes was tested for a number of levels of take-up, on a test network consisting of some 81 links and 21 nodes. It was seen that both equipped and unequipped vehicles experienced a reduction in travel time, at all levels of take-up. The maximum benefit to unguided vehicles (3%) and the system as a whole (6%) was virtually attained at 20% take-up, after which the changes were small. The greatest journey time savings for equipped drivers (15%) were obtained at 5%-10% take up, after which the benefits decreased considerably. These findings were considered to be a measure of maximum benefits, since in reality the guidance system would not continuously re-compute optimum routes. A more realistic strategy is therefore also studied, in which routes are re-calculated only every 15 minutes. In this case too it was found that the total travel time reduced for all levels of take-up considered, with the maximum saving of around 3% achieved at a take-up of 20%. Finally, a strategy consisting of a single route recommended for the whole peak period was studied and was found to produce a substantial increase in travel time, which became greater as more vehicles were equipped. Perhaps, more interesting than that performance of this strategy is, however, their findings regarding the influence of unguided drivers' behaviour. If, contrary to the assumptions above, unguided drivers do change their routes in response to the new conditions under guidance, then this single route strategy performs much differently, leading to a small (up to 1%) saving in total travel time for levels of take-up of less than 30%.

Koutsopoulos and Lotan (1989) modelled the interaction between equipped and unequipped vehicles in the form of a stochastic user equilibrium assignment of two user classes, where equipped drivers are assumed to have a smaller variance in their perceived cost of travel, to represent the effect of the guidance system improving the quality of the information available to them. They applied the model to the network of Sudbury, Massachussetts, consisting of 204 nodes, 70 zones and 578 links (including 214 centroid connectors). Under user optimal routing (when the perception errors of equipped drivers are completely removed) and with a 100% take-up of guidance, a saving in total travel time of 3%-4% was observed for three different levels of congestion in the demand. As would be expected, as the perception errors of equipped drivers increased in magnitude, the benefit decreased. Studying the effect of level of take-up on total travel time, it was found that under user optimal routing (with or without stochastic errors), time decreased linearly with the percentage of equipped drivers, for all levels of congestion. This finding was found, however, to be sensitive to the model of unequipped drivers; with a decrease of 10% in their perception error variance, total travel time was found to decrease less rapidly as take-

up increased. In terms of the benefits to individual groups - under stochastic user optimum routing, with a low level of error for guided drivers (the only strategy investigated in this way) - equipped drivers were found to experience a decrease in average travel time of around 3%. The benefit decreased with an increase in take-up, but by only a very small amount (range of less than 0.5%). For unequipped drivers, on the other hand, there was a disbenefit at all levels of take-up, although the increase in their average travel time due to guidance was never more than 0.2%. These results on equipped/unequipped drivers appear to relate to the least congested case; other results indicate that the underline{relative} benefit of equipped over unequipped drivers decreases as congestion increases.

Rakha et al (1989) were mainly interested in incidents and the optimisation of signal settings in interaction with route guidance, but in passing studied the case of route guidance alone, in incident-free conditions. They considered a small network consisting of a freeway and parallel arterial, with five connecting roads. The simulation model used was INTEGRATION (Van Aerde and Yagar, 1988), which requires the input of a fixed demand pattern. Unequipped drivers were assumed to follow the minimum free flow travel time route between each origin-destination pair, these being fixed throughout the simulation period. Equipped drivers, on the other hand, were guided to the minimum time route based on current traffic conditions - this information being updated approximately every six seconds, and providing the capability to re-route equipped drivers during their journey. The average network speed before guidance was high - 59kph - although it may be expected to be higher than in a usual urban situation, due to the relatively large amount of freeway travel. For the lowest level of take-up considered (20%), they found a huge reduction in total travel time of some 18%; although this benefit continued to increase with level of take-up, only an additional 3% saving was accrued by 100% take-up. Average network speed and total distance travelled, on the other hand, increased with level of take-up. These findings are clearly dependent to a great extent on the model of unguided drivers, who are assumed not to respond to congestion. Since the length of a link is probably a reasonable proxy for free flow travel time, their study of the impact of guidance is very much like a comparison between a population of distance minimisers (before guidance) and one of time minimisers and distance minimisers (after guidance) - not surprisingly, since the former population would yield minimum total distance travelled in the network, the latter population travels a greater distance but in a smaller time (because minimisation of time, albeit on an individual basis, is now an objective).

A quite different approach to those discussed so far was proposed by Mahmassani and Jayakrishnan (1989), being the first specifically to address the behaviour of equipped drivers in response to guidance. The dynamic model they use comprises a macro-simulator coupled with a set of rules for describing the behaviour of individual guided drivers. It is assumed that the guidance system tells the driver the estimated time to his destination on his current route and on what is predicted to be the quickest route. The driver switches to the recommended route if the percentage journey time saving is greater than some threshold value, this value being modelled as a random variable across the user population. (A device is also included to make switching less likely as the destination is approached). This is intended to represent, for example, drivers not wishing to change when the alternative route is unfamiliar. Two implementational aspects to note are, firstly, that the model has been developed

specifically for a network of the type they use for evaluation purposes (a commuting corridor consisting of three major expressways, connected at regular intervals by cross-over links). Secondly, the computational demands for such detailed modelling are very high (viz. a Cray super computer).

Two quite different "loading patterns" were used to assess the strategy (with the first, more congested, pattern having a "before guidance" travel time of nearly twice as much as that of the second pattern, which tends to spread peak traffic over a greater period). In the case where the switching threshold was identically equal to zero - so that equipped drivers always switch to the quicker routes, when recommended - the maximum saving in total travel time was around 5% at levels of take-up of 10% and 25%; but at 50% and above, the benefits decreased considerably, even to the extent of a large disbenefit. In terms of the effect on individual equipped drivers, the greatest benefit was again found to occur at the lowest level of take-up (10%), with travel time savings decreasing to a 8%-9% <u>increase</u> in average travel time at 75% take-up. Unequipped drivers <u>always</u> benefitted, though never as much as the maximum for equipped drivers with the maximum savings again at low levels of take-up. For switching thresholds other than zero, on the other hand - except when drivers were very unwilling to switch (where there was little change in travel times) - it was seen that for all levels of take-up, equipped drivers, unequipped drivers and the system as a whole (maximum 6% saving) would benefit from guidance. The travel time savings increase with take-up for unequipped drivers and the system as a whole.

The second loading pattern studied satisfied stochastic, dynamic user equilibrium conditions, in the sense that "no user can improve his/her random utility by unilaterally switching either departure time or route", where the components of utility are travel time and the difference between desired and actual arrival time. For all switching thresholds considered (except the one with a high mean, where there was little change), a benefit was observed for both equipped and unequipped vehicles at all levels of take-up. In the degenerate case of a threshold equal to zero, the maximum benefit for the system (4% reduction in total travel time) and for unequipped drivers (2%) was at a take-up of 25%-50%, whereas the benefits to equipped drivers steadily decreased with increasing take-up (from a maximum 12% saving at a 10% take-up). For the other, randomly dispersed thresholds, although the maximum benefits were similar to those in the zero threshold case, the behaviour was somewhat different at higher levels of take-up. The benefit to unguided drivers and the system as a whole increased with take-up, although the <u>rate</u> of this change decreased; meanwhile, for equipped drivers, the travel time savings decreased, but again at a decreasing rate.

Finally, Van Vuren et al (1989) studied a model of route guidance which is closely related to that considered here (and will be discussed again, later in this report). Their steady state model was one of a two user class assignment in which unguided drivers were assumed to follow a user equilibrium and guided drivers were routed according to a system optimal flow pattern with interactions between guided and unguided drivers thereby modelled. Conditions were derived under which the combined flow pattern is guaranteed to be unique and stable. Furthermore, it was shown that in such a system, the equipped and unequipped drivers could share at most one route for each origin-destination pair; guided drivers on other routes take a

travel time which is higher than that of the unguided drivers. Numerical results were confined to a demonstration on an artificial two-link network.

3. BACKGROUND: MULTIPLE USER CLASS ASSIGNMENT

Having reviewed previous work in simulating route guidance systems, the model to be studied in this report may be developed. It will be based on the concept of multiple user class equilibrium assignment, and so a brief introduction to this area will be given, along with relevant theoretical results.

As a steady state model of driver route choice under long term average conditions, the user equilibrium (UE) assignment proposed by Wardrop (1952) - in which each driver is assumed to be aiming to minimise, non-cooperatively, his own personal travel cost - is accepted by many to be a reasonable approximation to "average" driver behaviour. Two key assumptions are that drivers define "cost" in the same way and that they <u>know</u> the cost of travelling along each route between their origin and destination. Studies show (for example, Wootton and Ness, 1989 - from where the figures to follow are taken) that although time minimisation is the predominant route choice criterion (more than 70% of drivers), distance minimisation (10%) and having no known alternative route (10%) are also reasons given for choosing a particular path. In such cases, when it is possible to divide the demand, a priori, into "user classes", where the cost definition is the same within a user class (but may differ between classes), a multiple user class equilibrium assignment may be defined. <u>Within</u> each user class, a user equilibrium is obtained, but there is a dependence between the classes, in the sense that the flow of one user class affect the costs of another user class (Van Vliet et al, 1986). As will become clear later, the fact that different cost functions may be specified for each user class makes the concept particularly relevant to the modelling of guided and unguided vehicles in a route guidance context. The framework also allows network restrictions to be applied to certain user classes, so that it is straightforward to define a limited (strategic or tactical) network available to equipped drivers - although the study of such strategies is not within the scope of this report.

The second assumption in Wardrop's model which was highlighted above - that drivers know the relevant travel costs - may also be relaxed to some extent. (It is noted that Wootton and Ness found that "only 50% of drivers seeking either their quickest or shortest routes succeed in finding the route they desire"). The reasoning is that if drivers do not know the costs exactly, they will not necessarily choose the user optimum routes when they aim to minimise their own cost of travel. Such an effect may be reproduced by a stochastic user equilibrium (SUE) assignment - developed by Daganzo and Sheffi (1977) - in which the link costs are treated as random variables, but with the same behavioural assumption as Wardrop (ie drivers <u>aim</u> to minimise personal travel cost). Because of the mutual dependence between cost and flow, the link flows are then also random variables, the assignment being the expectation of these random variables. The randomised costs are usually referred to as <u>perceived costs</u>, with the random perturbations known as perception errors. The name "perceived cost" stems from the idea that drivers may perceive cost differently - in a (macro) SUE assignment, this micro behaviour is modelled by the random variation in the sample mean link cost (since if the individual perceived driver costs are random variables, then in repeated sampling the <u>mean</u> of these costs

will also be a random variable). It will be proposed here that a SUE assignment is a reasonable model of the route choice behaviour of drivers who do not have access to a guidance device.

The concept of stochastic equilibrium assignment and multiple user classes were brought together in a paper by Daganzo (1983); the main results of this paper are outlined here. Daganzo considered a family of link cost functions of the form

$$c_{ai} = d_{ai} + \text{\ss}_i \, t_a \, (\underline{F})$$

where for each link a and user class i,

c_{ai} is the cost to user class i of using link a

$F_a = \Sigma_j \, \alpha_j \, f_{aj}$

$\underline{F} = (F_1, F_2, ...)$

f_{ai} is the flow of user class i on link a

t_a is a continuously differentiable function

and

d_{ai}, α_i (>0) and $\text{\ss}_i$ (>0) are finite constants.

The idea of representing the flow of all user classes on a link in terms of a single measure F_a is consistent, for example, with the concept of expressing a combination of car and lorry flow in terms of "passenger car units" or "passenger car equivalents".

The perceived cost of travel C_{ai} for link a and user class i is assumed to be given by

$$C_{ai} = c_{ai} + \Phi_{ai}$$

where Φ_{ai} is a random variable (the "perception error") with an expectation of zero.

Daganzo imposed two sets of conditions - firstly on the cost functions and secondly on the perception errors- under which a stable equilibrium flow pattern exists and is unique:

(1)The <u>inverse</u> of $\underline{t}$ $(\underline{F})$ is a monotonically increasing, continuously differentiable function in the domain where it takes finite values, is defined for all $\underline{t}$ $(=(t_1, t_2, ...))$, and is uniformly bounded. Furthermore, $\underline{t}$ $(\underline{F})$ is defined for all feasible flow patterns.

(2)For each user class i, the components of $\underline{\Phi}_i = (\Phi_{1i}, \Phi_{2i}, ...)$ are mutually independent, independent of the costs c_{ai}, and have densities which are finite, have at most a finite number of discontinuities and have finite second moments.

Now, condition 2 <u>excludes</u> the possibility of some of the perceived costs being

deterministic; that is, for example, if there are no perception errors for a certain class of drivers. In the case of models with some deterministic aspects, two approaches are suggested. One option is to approximate the deterministic costs by stochastic costs with sufficiently narrow densities. Alternatively, Daganzo offers modified properties of the assignment, which may be shown to follow when the second condition is relaxed to:

(2')Condition 2 holds, but some (or all) of the Φ_{ai} are allowed to be zero.

Under this revised condition, a stable equilibrium is guaranteed to exist and will be unique with respect to link flows (though not necessarily with respect to user class link flows).

Daganzo also suggests two solution algorithms for the multiple user class SUE problem - one in terms of link costs and the other based on flows. The latter was chosen for the purposes of this study, being a more standard formulation (an extension of the "method of successive averages", suggested by Powell and Sheffi (1982) for the single user class SUE problem) and being already implemented within SATURN. In order to prove the convergence of his flow-space algorithm, Daganzo made the additional assumption that the Jacobian of $\underline{t}$ ($\underline{F}$) is symmetric. The scheme is implemented as follows:

(1)Set $f_{ai}^{(0)} = 0$, $\forall$ a, i, where $f_{ai}^{(r)}$ refers to the estimate of the equilibrium user class flows at iteration r. Set r = 0.

(2)Calculate $\underline{F}^{(r)}$ from

$$F_a^{(r)} = \Sigma_j \; \alpha_j \; f_{aj}^{(r)}$$

and hence the costs $c_{ai}^{(r)}$ corresponding to $\underline{F}^{(r)}$ ($\forall$ a, i).

(3)   For each user class i:

(a)Sample a set of link error terms Φ_{ai} ($>a$) from the specified probability distribution, by a pseudo-randomisation process, and set

$$C_{ai}^{(r)} = c_{ai}^{(r)} + \Phi_{ai} \; \forall \; a$$

(b)Perform an all-or-nothing assignment for this user class using the randomised costs $C_{ai}^{(r)}$ - yielding a set of user class link flows $g_{ai}^{(r)}$ (for all a)

(c)Set

$$f_{ai}^{(r+1)} = (1 - 1/n) \; f_{ai}^{(r)} + 1/n \; g_{ai}^{(r)} \; (\forall \; a)$$

(d)   Set

$$F_a^{(r)} = F_a^{(r)} - \alpha_i \; f_{ai}^{(r)} + \alpha_i \; f_{ai}^{(r+1)} \; (\forall \; a)$$

and recompute the costs corresponding to $\underline{F}^{(r)}$ - store again in $c_{ai}^{(r)}$ ($\forall$ a)

(4) Set r = r+1 and return to step 2 until the pre-determined number of iterations are complete.

Two points are worth noting regarding the details of the above algorithm. On a particular iteration, the costs are updated after loading each user class; an alternative would have been to recalculate the costs only after loading all user classes. In the former approach, the aim is in some way to anticipate the effect on link costs for the next iteration, of the newly loaded flows of user classes already considered on the current iteration and thereby to improve the rate of convergence to the equilibrium link flow pattern. As stated earlier, however, it may not be possible to guarantee the uniqueness of the user class link flows, and in this sense the latter approach may produce a more even distribution of flow between the user classes on a link basis. This is because in the former approach, after loading user class 1 flows on a particular iteration, for example, a link may become congested and then very few user class 2 flows will be loaded onto it.

The second point to note is that the randomisation process in step 3(a) of the above algorithm follows that standard SATURN procedure of "sampling" a new set of link error terms before building the minimum cost tree from each origin, and so in fact there are a number of randomisations performed for each application of this step.

4.   MUC ROUTE GUIDANCE MODEL

Having introduced the concept of multiple user class equilibrium assignment, with possibly a mixture of stochastic and deterministic costs, the route guidance model will now be introduced.

The demand for travel (as represented by the mean origin-destination matrix) is assumed to be fixed, as are the network supply conditions. It is also assumed that the whole network is available to the guidance system. Average cost-flow relationships are supplied for each link. Throughout this report, "cost" will be measured purely in terms of travel time, and so the words cost and time will be used interchangeably below.

The model consists of four user classes, the demand level for each being a fixed (known) proportion of the origin-destination matrix. Three of the user classes correspond to vehicles equipped with a guidance device, and for two of these three equipped classes it is assumed further that the guidance system is provided with perfect information and that the guided drivers adhere totally to the route recommendations. The first class consists of unguided drivers, each of which aims to minimize his own personal cost of travel, but in general fails to do so because of imperfect knowledge of the traffic conditions. This class is modelled by a stochastic user equilibrium (SUE), the "perceived cost" for each link following some specified distribution (discussed later). The second class is a subset of the equipped vehicles where each driver is guided so as to minimize his own personal travel cost. The perfect information assumed to be available to the guidance system is used to eliminate the perception errors, i.e. they follow a Wardrop user equilibrium. The third class consists of a second subset of the equipped drivers, which are guided so as to minimize the total system cost ("system optimal" - SO), again using the perfect information available. The fourth class comprises the remaining equipped drivers. The aim of the guidance system for this class is again to recommend routes according to a UE pattern; however, in order to represent the effect of errors in the journey time prediction methods or of drivers not adhering completely to the recommendations, they are modelled by a SUE, but with a distribution for the stochastic variations which is different to that of the unguided drivers.

The four user classes interact with one another, in the sense that the flows of one user class affect the costs, and hence the route choice, of the other user classes. In this way, the assumption is that under such steady state conditions, the unguided drivers will tend to change their routes in response to the new route choice of the guided drivers (the influence of this assumption will be investigated at a later stage in the report).

Now, Van Vuren et al. (1989) concluded that for the case of a guidance system with user equilibrium unguided drivers and system optimal guided drivers, the only link cost functions c_a of the family which was established by Van Vliet et al. (1986) to ensure existence and uniqueness of a multiple user class equilibrium, were of the polynomial form:

$$c_a = d_a + b_a F_a^k \qquad \qquad ...(1)$$

where F_a is the total flow on link a, d_a is a constant representing fixed effects such as free flow travel time, b_a is a constant and the power k (>0) is a link independent constant. In the more general four user class model considered here, we cannot use the same result of Van Vliet et al., since the properties were established only for the deterministic cost case, whereas here we have a mixture of stochastic and

deterministic costs. We may, however, use results established by Daganzo (1983) for a similar (though more general) family of cost functions to those of Van Vliet el al., but for the case where some of the classes may have stochastic costs. Then, in a similar way to Van Vuren et al., by applying the work of Daganzo, it follows that the equilibrium for our more general four user class model is guaranteed to exist and be unique (with respect to link flows and user class/link costs) for cost functions of the form (1). In this case, c_a is the actual link travel time for all drivers; in the assignment, however, each class will be associated with a different cost: the unguided drivers making random perception errors with (1) as the mean; the guided SUE drivers experiencing different random errors due to imperfect recommendations, etc.; the guided SO drivers using marginal costs corresponding to the actual costs (1); and the guided UE drivers using the actual costs (1). The fact that Daganzo's results may be applied to guarantee the above conditions on the equilibrium may be verified as follows.

It is well known that a system optimal assignment in the one user class case with link costs c_a may be obtained by a user equilibrium assignment with marginal link costs c'_a given by

$$c'_a = c_a + F_a \; \frac{dc_a}{dF_a}$$

To obtain, then, the required routing pattern with actual link costs (1), the user class costs c_{ai} for link a and user class i must be

$$c_{a1} = c_{a2} = c_{a4} = d_a + [b_a F_a^k]$$

$$c_{a3} = d_a + (k+1) [b_a F_a^k]$$

where user class 1 consists of the unguided drivers, and the remaining classes are the guided drivers, following (respectively) UE, SO and SUE routing; perceived costs are therefore stochastic for user classes 1 and 4, and deterministic for user classes 2 and 3. It may be seen that the user class cost functions above are indeed of the form required to apply Daganzo's work.

Furthermore, Daganzo's conditions require that the variance of the perceived journey time distribution is flow independent. This condition has been noted variously by authors investigating the single user class stochastic user equilibrium case: Sheffi and Powell (1982), Daganzo (1982) and Sheffi (1985). In the latter reference, Sheffi suggests - for a probit-based route choice model - the use of a standard deviation of link a perceived cost of θc_{oa}, where c_{oa} is the free-flow travel time and θ (>0) is a constant. In the guidance model proposed here, we have also chosen to use Normally distributed perception errors for the unguided drivers, but with a standard deviation of θc_a^{UE}, where c_a^{UE} is the travel cost for link a corresponding to a (deterministic) user equilibrium flow pattern for all drivers. This is preferred because it is more closely related to the idea that larger perception errors are made with larger travel times and greater congestion, rather than using the free flow travel time which may be more related to the physical characteristics of the link (for example, an uncongested freeway would have a relatively large free-flow travel time and would thus counter-intuitively tend to result in large perception errors). The guided SUE drivers are

modelled in the same way, except that their link travel time standard deviation is Ψc_a^{UE}, where $0<\Psi<\theta$ (ie guidance tends to reduce the size of the errors made by equipped drivers). The errors are distributed independently between user classes and between links. It is noted that this model is somewhat unrealistic in one respect, since the journey time prediction methods will tend to be more precise with larger levels of take-up - data on actual travel times relayed to the guidance system from the beacons will relate only to equipped vehicles, and so an increase in level of take-up will essentially lead to an increase in sample size. It would be expected, then, that the variance of the random errors would be a decreasing function of the level of take-up. Since no suitable relationship of this kind was available, however, it was necessary to retain the assumption of a constant error variance relative to the proportion of vehicles equipped.

5.   TEST RESULTS

The guidance strategy was implemented using an adaptation of the simulation/ assignment model SATURN (Van Vliet, 1982) and the solution algorithm of Daganzo, as described above.   Scenarios to simulate were chosen according to the recommendations of deliverable 9 (CAR-GOES, 1990b).   The two real-life networks considered were those of Weetwood (an area of Leeds) consisting of 70 zones, 104 intersections and 440 links, with approximately 20,000 origin-destination trips in the base (demand level 1) matrix; and of Barcelona comprising some 110 zones, 820 intersections and 2547 links, with around 60,000 origin-destination trips in the base matrix.   The cost functions used were of the form (1), where k was given the value 5 for both networks.

For each network, the guidance model was implemented under

(a)three different demand levels, corresponding to an average network speed (before guidance) of approximately 15, 25 and 35 kph;
(b)nine different levels of equipped vehicles: 0%, 5%, 10%, 20%, 30%, 50%, 70%, 90% and 100%; and
(c)three different routing criteria - with equipped drivers either all guided as a UE, all guided as a SO or all guided as a SUE (with two different levels of error in this latter case).

The desired network speeds were obtained by running the model a number of times with different origin-destination matrices, obtained by multiplying the base matrix by a scalar factor.   Coincidentally, it was found that for both networks, the demand levels defined in (a) corresponded to 100%, 130% and 160% of the observed origin-destination flows.

It is noted that additional scenarios, of a level of take-up of 0.1% and 1%, were also studied, although it was not convenient to display these results (in any case, the strategies developed here are specifically carried at higher participation levels); the results, may, however, be referred to in passing in the text.

Finally, in order to decide upon a suitable value for the parameter θ, which determines the link travel time variances for the unguided drivers, an idea due to Breheret et al (1990) is used.   For a number of values of θ, the average inefficiency

I(θ) is calculated, given by

I(θ) = 100 (p(θ) - 1) %

where

p(θ) = $\dfrac{\text{Total system travel time under SUE}(\theta)}{\text{Total system travel time under UE}}$

and where SUE(θ) means an SUE assignment for the whole O-D matrix, with parameter θ. That is, assuming that drivers are aiming to follow a UE, I(θ) is a measure of the average excess travel time incurred by their perception errors.

For the purposes of this report, for a given network, a value for θ is then chosen which gives rise to an inefficiency of the order of 5%-6% for each of the demand levels considered. The reasoning behind this is that various studies have shown that the percentage wastage caused by drivers not fulfilling their objective of choosing the minimum time or minimum distance route is of this order - for example, Lunn (1978) estimated the average excess on all journeys in Great Britain, excluding commuter trips, to be at least 5% of total costs; Wootton et al (1981) arrived at figures of 4%-6.5% inefficiency; and Jeffrey (1987), from analysing times and distances corresponding to a sample of journeys made in the U.K., concluded that the average inefficiency of drivers was around 6%. The use of inefficiency to build a suitable route choice model for unguided drivers is appealing, in that it allows the calibration of the model against observed data (though in a very coarse manner).

The values of $I(\theta)$ for a number of values of θ are given in figures 1 and 2. For Weetwood, there is a clear pattern of an increased $I(\theta)$ with increased θ or greater demand. The value $\theta=0.3$ is chosen for the purposes of further investigation, giving an average inefficiency of 6% over the three demand levels. For the Barcelona network, the pattern is somewhat different, with much less difference between demand levels and with the possibility of $I(\theta)$ decreasing with greater demand (demand level 1 showing greatest inefficiencies). It is still the case, however, that $I(\theta)$ is an increasing function of θ. $\theta=0.4$ is chosen for future study.

There are two studies (the findings of which were described in Section 2) with which some comparison may be drawn on this point of modelling unguided drivers. Breheret et al (1990), in using a Uniform error structure for perceived costs (with flow dependent range), found the relationship between inefficiency and spread parameter to be highly network dependent and demand dependent - because of this, and because they give no indication as to the size of the networks or the absolute levels of congestion, it is difficult to draw any further parallels with this work. Koutsopoulos and Lotan (1989), on the other hand, used a very similar `before guidance' model to that considered here, the most notable disparity being their use of a flow dependent perception error variance of $\varphi c_a(F_a)$. For their study on a network of a similar size to Weetwood, they used a value of $\varphi=0.5$, which gave rise to an inefficiency of around 4% (relative to the UE, $\varphi=0$, case) for the three demand levels considered.

Note: It is evident that quite large values of the spread parameters are required to give 'realistic' inefficiencies. In one respect this is unappealing, since - as it makes sense to truncate the perceived travel time distributions at zero - the randomisation may be biased (although for the values chosen below for the Weetwood and Barcelona networks, the bias will tend to be very small).

The model was applied to the networks described, with unguided drivers modelled by a SUE with $\theta=0.3$ for Weetwood and $\theta=0.4$ for Barcelona. From initial studies of the Weetwood network, it appeared that in order to obtain a reasonable degree of convergence a large number of iterations would be required - the stopping criterion chosen was the completion of 200 iterations. Although this may have been the case too for Barcelona, the size of the network meant that such a large number of iterations would be computationally prohibitive, and so only 30 iterations were carried out in this case. The results are given in Figures 3 to 14. In the figures on system benefit (expressed in terms of total travel time, Figures 3-5 and 9-11), each of

16

the four strategies is represented - for example, the trend labelled SUE(0.2) refers to the strategy in which all equipped vehicles are guided according to a SUE with $\Psi=0.2$. In the figures on individual benefits (figures 6-8 and 12-14), results relating to the UE (broken line) and SO (continuous line) routing strategies only are given, with each separated into guided (star symbol) and unguided (square symbol) vehicles.

The first point to note is that in a small number of situations, there is evidence of strange behaviour - firstly, in figure 5, for Weetwood at demand level 3 with 100% of vehicles equipped, the total travel time arising from SO routing is greater than that arising from UE routing, which clearly should not be the case. Investigating the algorithm further for a number of different cases, it was found that the flows under SO routing tend to converge slower than under UE routing (convergence being measured by the indicator suggested by Sheffi & Powell (1982) for SUE assignment, and applied here to the total link flows). The results given here were obtained by applying the algorithm for a pre-determined number of iterations; it appears to be, then, that the results are due to the algorithm achieving a greater degree of convergence under UE, than SO, routing. Secondly, for Barcelona at higher demand levels (figures 10 and 11) a similar problem is evident - for example, in figure 11, an assignment with 70% guided according to a SO routing gives a smaller total travel time than a 100% SO routing. This is most likely due to the smaller number of iterations permitted for Barcelona, where in some cases the flow pattern may not have stabilised by the time the algorithm is terminated. The Barcelona results therefore have to be viewed with a little more scepticism than the Weetwood ones.

In terms of the system benefit (as measured by total travel time), it can be seen that for all of the routing strategies considered and for both networks, guidance offers an improvement over the base (no guidance) situation, for all levels of equipped vehicles and all demand levels considered. In all cases, the travel time saving becomes greater as the level of take-up increases (with one or two exceptions - see comments above on convergence), following an almost linear trend in the Weetwood case. Below a 50% take-up, the percentage saving in total travel time tends to be a higher with higher demand levels, although in all cases the differences between demand levels are not great. Concentrating specifically on the UE and SO routing strategies, figures 3-5 show that for Weetwood the percentage saving in total travel time due to guidance increases with demand for most levels of take-up - for example, between 10% and 90% take-up, UE routing gives savings of 0.5%-4.0%, 0.9%-5.6% and 1.0%-7.1% respectively for demand levels 1, 2 and 3; SO routing, on the other hand, gives respective savings of 0.7%-6.4%, 0.8%-6.6% and 0.8%-7.0%. The differences between demand levels are, however, not great, and the same is true for the Barcelona network. In this latter case, though, there is a slight decrease in the benefits attainable at higher levels of take-up as demand increases (figures 9-11) - the savings for 10%-90% take-up are 0.9%-6.0%, 1.2%-5.5% and 1.2%-5.3% under UE routing for demand levels 1, 2 and 3 respectively, and under SO routing the respective savings are 1.5%-8.5%, 1.5%-7.7% and 1.8%-5.8%. On the whole, for both networks, the pattern is as one may expect, with an increase in total travel time savings as Ψ decreases for SUE routing (down to $\Psi=0$ for UE routing), all of these giving rise to larger total travel times than SO routing.

In terms of the benefit to individuals (the percentage decrease in average travel time with the guidance system in operation), it may be seen that for a UE routing,

equipped drivers are always better off with guidance. A striking feature is that this benefit is approximately constant for all levels of take-up - notably, the benefit is achieved at a very low percentage of equipped vehicles (in fact at 1% take-up, not shown in the graphs), and does not start to decrease at some participation level. From figures 6-8, it may be seen that under UE routing, guided drivers save around 4%, 5% and 6% respectively for Weetwood demand levels 1, 2 and 3, whereas (figures 12-14) the savings are around 8%, 7% and 6% for Barcelona demand levels 1, 2 and 3. Under such a routing scenario, the change in travel time for unequipped drivers is always small relative to the benefit to guided drivers (always smaller than 3%, and usually less than 1%, with an actual disbenefit for the Barcelona network at demand level 1), although there appears to be slightly greater benefit to them as congestion increases from demand level 1 to demand level 3. In most of the situations, the benefit to individual unguided drivers also tends to increase with level of take-up.

With SO routing, there is a disbenefit to individual guided drivers on average, for `lower' levels of take-up (`low' being levels of equipped vehicles less than of the order of 10%-30%); for higher levels of take-up, on the other hand, guided drivers experience a saving in travel time which increases with level of take-up. Above 50% take-up the savings for equipped drivers under SO routing are 3%-7% for Weetwood and 5%-9% for Barcelona, depending on level of take-up and demand level. The journey time saving for unguided drivers tends to be somewhat larger here than with UE routing, particularly at lower levels of take-up, with guided and unguided drivers benefiting similar amounts at higher levels of take-up.

Comparing the UE and SO routing strategies for both networks, it may be seen that UE guidance will always benefit the equipped drivers most, with only limited benefits to the unequipped drivers, but giving rise to considerable system benefits. Not surprisingly, SO routing primarily benefits the unguided drivers - at the expense of the guided drivers - at lower levels of equipped vehicles. However, equipped drivers start benefiting too when their numbers increase. At the highest levels of equipped vehicles (over 50%-70%), under such a routing strategy the guided drivers may even benefit more than the unguided drivers, despite being guided to minimum marginal cost routes. The system benefits of SO routing are higher than with UE routing, but may not warrant the disbenefits to equipped drivers at lower participation levels.

The above discussion has been concerned with the effect of route guidance on travel time - it seems reasonable to use this as the main factor for assessing the strategies, as it has been assumed that drivers measure cost purely in terms of time. However, other useful indicators have also been studied - namely, distance travelled and average speed - and results corresponding to these are given in Figures 15 to 20 (distance) and 21 to 23 (speed), for the Weetwood network only.

Firstly, considering figures 15 to 17, under UE routing it may be seen that the total distance travelled is always reduced with guidance in operation, decreasing as level of take-up increases (with a maximum 4%, 5% and 6% saving for demand levels 1, 2 and 3 respectively). Similarly, with the SUE routing strategies there is always a benefit, which increases with take-up. For the SO routing scenarios, on the other hand, the total distance travelled under guidance increases with take-up (up to the point at which around 30% of drivers are equipped), and there is not a benefit in the respect until about 50%-70% take-up.

Examining the effect on individuals (figures 18 to 20), under UE routing a similar pattern is evident with the change in average distance travelled as there was with the change in average travel time - namely, that equipped drivers benefit a constant amount (around 5%-6% saving in distance travelled) at all levels of take-up.

Unequipped drivers, on the other hand, lose out under such a strategy with an increase in distance travelled of up to 1%-2% at higher levels of take-up. The pattern for SO routing is quite different, with equipped drivers experiencing a sharp, substantial increase in distance travelled at the lowest levels of take-up (even higher at 1% take-up - not shown on the graphs) - the disbenefit at this stage is between 9% and 13%, dependent on the demand level. This is as one may expect an SO routing to work, with equipped drivers guided to longer, but less congested, routes. As more vehicles become equipped, the disbenefit becomes smaller in magnitude, but does not again reach the average distance travelled at 0% take-up until the stage at which about 50%-70% of drivers have guidance devices. Unequipped drivers, on the other hand, travel a shorter distance on average under SO routing than without guidance in operation, regardless of the proportion of drivers equipped - although the savings are small (around 1%).

Finally, in figures 21-23, the average speed in the network is shown as a function of the proportion of equipped vehicles. Only the SO and UE routing patterns are shown, as the SUE routing patterns are very similar to the UE ones, and only serve to obscure the other results. The graphs indicate that under UE routing, the speed does not vary much with the level of take-up, but is nevertheless always higher with guidance in operation than it is without. The maximum increase in speed is 2%, with the benefits generally being greater at higher demand levels. With the SO routing strategy, on the other hand, a steady increase in average speed is evident as the level of take-up increases, which flattens out as the point is approached at which all drivers are equipped (the maximum increase in speed is 5%-6%). The reason for the increased speed, from considering results already discussed, is the greater distance which tends to be travelled under SO routing, but in a shorter time, relative to the base (no guidance) situation.

6.   SENSITIVITY TO VARIABILITY IN NETWORK SUPPLY

A number of features of the strategy proposed and the model assumed will now be investigated, beginning in this section with a study of the "stability" of the strategies in conditions where there are unforeseen variations in link capacities. The guidance system and the unguided drivers are assumed to base their routing decisions on the average capacity of each link - denoted by h_a for link a, say. In order to reproduce the effect of day-to-day variations in capacity, for example, the actual capacity for link a is modelled as a Normally distributed random variable with mean h_a and standard deviation δh_a (where $\delta > 0$) (NB The distribution is truncated at zero, and so may be biased to a small degree). Since SATURN, on which the guidance model is based, calculates the parameters in the flow-delay curve by fitting it to one point at zero flow and free flow travel time and to one at capacity and travel time at capacity (where free flow travel time and travel time at capacity are assumed here to be fixed, user-specified quantities), then link travel times will be different using the actual capacity as opposed to the average capacity.

In order to gain an insight into the influence of such variations, the Weetwood network was investigated at demand levels 1 and 2, for a take-up of 0%, 5%, 10%, 20%, 30%, and 50%, and under both UE and SO routing strategies for the guided drivers. In each case, the link capacities in the network were completely randomised two hundred times, the total travel time being calculated each time. The mean and

standard deviation of the total travel times was then computed (care was taken to draw the same two hundred randomised capacities for each scenario). The results obtained were as follows:

Demand level	% guided	Guidance criterion	δ	Mean total time	s.d. total time
1	0	-	0.05	4403.4	53.7
1	5	UE	0.05	4383.5	53.7
1	10	UE	0.05	4379.6	53.9
1	20	UE	0.05	4355.3	53.2
1	30	UE	0.05	4335.0	53.6
1	50	UE	0.05	4293.4	54.6
1	5	SO	0.05	4384.2	52.4
1	10	SO	0.05	4368.8	50.1
1	20	SO	0.05	4341.5	48.1
1	30	SO	0.05	4316.6	46.5
1	50	SO	0.05	4243.2	46.1
1	0	-	0.10	4531.6	134.2
1	5	UE	0.10	4509.6	133.9
1	10	UE	0.10	4506.4	134.2
1	20	UE	0.10	4480.1	132.3
1	30	UE	0.10	4458.5	133.4
1	50	UE	0.10	4415.0	135.2
1	5	SO	0.10	4508.5	130.5
1	10	SO	0.10	4489.8	125.9
1	20	SO	0.10	4455.8	119.9
1	30	SO	0.10	4427.1	116.3
1	50	SO	0.10	4346.6	115.4
1	0	-	0.15	4841.9	352.9
1	5	UE	0.15	4814.0	348.5
1	10	UE	0.15	4813.3	353.4
1	20	UE	0.15	4781.3	342.9
1	30	UE	0.15	4756.4	343.3
1	50	UE	0.15	4707.1	343.1
1	5	SO	0.15	4808.9	340.4
1	10	SO	0.15	4783.0	332.7
1	20	SO	0.15	4732.4	313.8
1	30	SO	0.15	4695.0	306.9
1	50	SO	0.15	4597.4	303.0
2	0	UE	0.05	8376.0	210.9

2	5	UE	0.05	8312.2	210.5
2	10	UE	0.05	8299.4	211.8
2	20	UE	0.05	8245.1	211.7
2	30	UE	0.05	8200.3	210.7
2	50	UE	0.05	8083.6	211.5

Demand level	% guided	Guidance criterion	δ	Mean total time	s.d. total time
2	5	SO	0.05	8364.7	211.5
2	10	SO	0.05	8306.3	208.9
2	20	SO	0.05	8232.6	206.9
2	30	SO	0.05	8187.2	205.4
2	50	SO	0.05	8065.1	203.7
2	0	-	0.10	8877.1	536.6
2	5	UE	0.10	8804.8	533.2
2	10	UE	0.10	8790.9	536.6
2	20	UE	0.10	8731.9	536.6
2	30	UE	0.10	8683.3	531.5
2	50	UE	0.10	8555.4	534.1
2	5	SO	0.10	8862.1	535.6
2	10	SO	0.10	8794.1	531.8
2	20	SO	0.10	8708.2	523.5
2	30	SO	0.10	8655.4	520.0
2	50	SO	0.10	8518.0	520.5
2	0	-	0.15	10093.0	1421.8
2	5	UE	0.15	9999.4	1403.6
2	10	UE	0.15	9983.7	1414.4
2	20	UE	0.15	9911.7	1406.5
2	30	UE	0.15	9853.6	1391.2
2	50	UE	0.15	9696.1	1382.6
2	5	SO	0.15	10069.4	1415.3
2	10	SO	0.15	9978.7	1408.3
2	20	SO	0.15	9861.6	1375.2
2	30	SO	0.15	9791.0	1364.7
2	50	SO	0.15	9617.0	1363.4

Looking at the results as a whole, it can be seen that in all cases, guidance decreases mean total travel time, irrespective of the level of variability in capacities; the benefit in this respect increases with level of take-up. As in the standard (fixed capacity) case, one or two counter-intuitive results may occur at lower levels of take-up for demand level 2, where SO routing may give rise to a slightly longer mean total travel time than UE routing; in all other cases, SO routing gives rise to a lower mean travel time. The advantage of SO routing appears to be greater when there is greater variability in link capacities and a high level of take-up.

However, perhaps the most interesting feature of the results is the variance in total travel time. For UE routing, on the whole only very small changes (which may be

positive or negative) to the total travel time standard deviation arise as a result of guidance; the exception appearing to be at the highest level of variability in capacities ($\delta = 0.15$) with a high (30% or more) level of take-up, where a significant decrease occurs in variability of total travel time. Under SO routing, however, there is always a significant reduction in the travel time standard deviation with guidance in operation, with this saving increasing with level of take-up. The greatest reduction is over 15%, achieved at the highest level of take-up considered, for demand level 1 and $\delta = 0.15$, with a similarly sized benefit for the same scenario when $\delta = 0.05$. At the higher demand level, the benefits tend to be smaller, but still significant (3%-4% reduction in standard deviation).

To conclude, then, this small study has to some extent demonstrated the stability of the two routing strategies, with significant (mean) total travel time savings over the `no guidance' situation, even when there is variability of which the strategy cannot take account. The SO routing strategy appears to be particularly appealing in this respect, being more reliable in terms of network performance, with the variance in total travel time decreasing significantly as more drivers are equipped with guidance devices.

7.   INFLUENCE OF UNGUIDED DRIVER ASSUMPTIONS

It has been assumed, in the equilibrium-based strategy proposed, that unequipped drivers will respond to the new route choice of guided drivers, and hence still follow a stochastic user equilibrium, in interaction with the guided drivers. Given that an SUE assignment is acceptable in the first place as a reasonable route choice model, it seems natural to assume that - in the long run - unguided drivers will still seek minimum cost routes when guidance is in operation. Such an assumption does, however, neglect any loyalty unguided drivers may have to the routes they used before guidance was implemented (compare with Mahmassani and Jayakrishnan, 1989). Moreover, if the results obtained here under steady state conditions are to be extrapolated to suggest the potential benefits of a <u>dynamic</u> route guidance system, then it would seem appropriate to investigate scenarios in which unequipped drivers are not able to react fully to the behaviour of guided drivers.

This will be achieved by investigating the other extreme to that already considered, in which unguided drivers stay on fixed routes and do not react in any way to the re-routing of guided drivers. This is achieved by performing a stochastic user equilibrium for the whole demand matrix; for a particular guidance scenario, the link flows arising from the 100% SUE assignment are multiplied by the proportion of drivers who are not equipped with guidance, and the result loaded onto the network as fixed flows, before performing the assignment for the guided drivers.

The test runs considered previously have all been repeated for the Weetwood network under this alternative model of unguided drivers, and the results are given in figures 24 to 29. Of particular interest here is a comparison with the results obtained when unguided drivers were assumed to re-route (figures 3 to 8); since the results are bound to coincide for 0% and 100% take-up, the interest will be in the relative benefits at intermediate points. Considering firstly the system travel time at demand level 1 - figures 3 and 24 - the most striking feature is the greater difference between the effects of SO and UE routing at lower levels (up to around 50%) of take-up for the

fixed route assumption over the re-routing one. In the fixed route case, the SO routing strategy is now a pure optimisation of total travel time with respect to guided drivers, with unequipped drivers not able to reduce the potential benefits. There is little effect on the (system) performance of the UE and SUE routing strategies of the unguided driver assumptions at this demand level. In more congested conditions (figures 4, 5, 25 and 26), similar comments apply, although the relative advantage of SO routing over UE routing decreases as demand increases.

In terms of the effects on individual drivers, at demand level 1 (figures 6 and 27) the changes in average travel time both for guided and for unguided drivers are similar under the two sets of assumptions. The unguided drivers under SO routing save slightly more in the fixed route case (over the re-routing case) at lower levels of take-up (<20%). The greatest disparity is, however, in the travel times of guided drivers under SO routing, particularly up to around 50% take-up. Whilst under such a strategy there is a similar disbenefit to equipped drivers under the two sets of assumptions at 5% and 10% take-up, the fixed route case gives rise to much greater travel time savings for guided drivers at 20%, 30% and 50% take-up.

At demand level 2 (comparing figures 7 and 28), the figures for UE routing are again similar in the two cases. For SO routing, on the other hand, guided drivers are on average slightly better off in the fixed case, as are unguided drivers at low percentages of equipped vehicles. Finally, in the most congested situation considered (figures 8 and 29), there is generally a slightly smaller average benefit for guided drivers with the UE routing strategy under the fixed route case compared with the re-routing case. The most distinctive feature is the greater similarity between UE and SO routing for guided drivers under the fixed route assumption.

To summarise, then, for a UE (or SUE) routing strategy the benefits of route guidance do not appear to be significantly affected by the response of unequipped drivers (assuming, as has been done throughout this paper, that it is possible to accurately model and predict this response). Under SO routing at lower levels of take-up, greater system benefits may be gained if unequipped drivers stay on fixed routes, as opposed to re-routing to new minimum perceived cost routes; it tends to be the equipped drivers who are gaining out of this.

8.   AN ALTERNATIVE SINGLE ROUTE STRATEGY

The strategies so far considered have been distinctive in the way that they direct traffic between each origin-destination pair to multiple routes (ie there is at least one recommended route). Such strategies have been seen to provide efficient guidance, even with a high proportion of vehicles equipped. In order to obtain some comparison with strategies in which only a single route is recommended between each O-D pair, a number of test runs were also carried out using programs developed at Leeds in the DRIVE "ASTERIX" project.

The procedure - known as SATRAP - is again an extension of the SATURN package, and is described in deliverable 2 of that project (ASTERIX, 1989). The aim of the strategy is to re-route guided drivers in response to day-to-day variability in demand, and so the approach is quite different to that considered thus far (it is closer to the approach of Smith and Russam (1989) described in Section 2). A particular set of

options of this procedure were chosen, yielding the following model of the "before guidance" situation:

(a) The "average day" route choice is obtained by a UE assignment of the average O-D matrix (that is, the O-D matrix supplied).

(b) Today's O-D matrix is obtained by a pseudo-random sampling: today's O-D flow Q_{ij} between origin i and destination j is assumed to be distributed Normally with mean q_{ij} and coefficient of variation $ß$, where q_{ij} is the corresponding element in the average day O-D matrix.

(c) Today's assignment is obtained by using the routes chosen in the average day case, with traffic assigned to routes in the same _proportions_ as the average day case. (That is, the average day route flows are scaled by a factor which is the ratio of the average day O-D flow to today's O-D flow).

It is assumed that network supply conditions are fixed, that unequipped drivers do not re-route in response to guidance, and that the guidance system is provided with perfect information, the recommendations of which are adhered to by the equipped drivers.

The model was applied to the Barcelona network (with thirty iterations for step (a)) using very similar flow-delay curves to those used in the equilibrium-based strategy - that is, the curves are identical for flows up to capacity, but in the SATRAP procedure the standard SATURN assumption is used of a linear build-up of delays over capacity. This was not considered to be a serious problem, as the aim was only to gain some insight into the performance of single route strategies. A small number of scenarios were considered; these were some (but not all) combinations of:

(i) Two levels of variability in the O-D matrix - $ß = 0.2$ and $ß = 0.7$

(ii) Guidance to the minimum actual cost route or to the minimum marginal cost route between each O-D pair (the latter strategy intended to push the flows closer to a system optimal assignment).

(iii) A number of levels of take-up - the number and range of these depending on the particular situation.

(iv) Two demand levels; level 1 corresponding to an average network speed for the average day assignment/O-D matrix of approximately 35km/h (100% of the "observed" matrix) and level 2 to a speed of 25km/h (150% of the observed matrix).

Two points are noted regarding these test runs. Firstly, the "before guidance" model was not calibrated in order to provide a reasonable inefficiency, as for the equilibrium-based strategies. This could have been achieved by basing the average day route choice on a SUE, and obtaining a reasonable level of inefficiency by the choice of the link travel time variances and the parameter $ß$ controlling the variability in the O-D matrix. This means that it will not be possible to compare the single route and the equilibrium-based strategies in absolute terms. Secondly, in drawing the

elements of today's trip matrix at step (b) of the model, the distribution is in fact truncated at zero. This means that in using a coefficient of variation as high as 0.7, the randomisation is almost certainly biased. This has the effect that the mean number of trips is increased in the randomisation process, and so - relative to ß = 0.2 - there is an increase in congestion as well as an increase in variability.

Considering firstly ß = 0.2 and demand level 1, figure 30 shows the total system travel time as a function of the proportion of vehicles equipped, both for re-routing to the minimum actual cost route and for re-routing to the minimum marginal cost route. Whilst, as mentioned above, little can be inferred about absolute benefits, the relative change as take-up is increased shows a clear pattern. Up to around 10% take-up there is an overall benefit of actual cost re-routing, although the greatest saving occurs at around 6% take-up. By the stage at which 20% of drivers are equipped, the recommended route for each O-D pair is so congested that the system is much worse off than with no guidance at all. Minimum marginal cost routing gives rise to significantly greater system benefits than minimum actual cost routing, with the greatest saving over the before guidance situation at a take-up around 9%. Again, however, as more vehicles become equipped, such a strategy will eventually be detrimental to system performance (at somewhere between 15%-20% take-up).

Figure 31 shows (for the same scenarios) that for guidance based on actual costs, the greatest benefit for equipped drivers occurs when few people have guidance devices (in fact 0.1% take-up - not shown in the figure - gives rise to a slightly larger benefit than 1% take-up too). The saving decreases as more vehicles become equipped, with guided drivers losing out (relative to the "no guidance" situation) when 5% or more vehicles are equipped. At very low levels of take-up (1% or less), unequipped drivers benefit through such guidance, but to a lesser extent than equipped ones. Contrary to the pattern for guided drivers, those without guidance experience a greater reduction in average travel time as the level of take-up increases, although there is some evidence of the benefit stabilising (at around 0.5%). Turning attention to minimum marginal cost routing, the plot shows the pattern one may have expected - with equipped drivers losing out to a large extent (with the magnitude of the disbenefit increasing with level of take-up), at the expense of unequipped drivers who always benefit.

Figure 32 shows the system performance of minimum actual cost re-routing under more congested conditions (again with ß = 0.2). As was observed for demand level 1, the greatest benefit is with around 5% of vehicles equipped. In absolute terms, the saving in vehicle hours is also similar at the two demand levels - although as a percentage of the total travel time without guidance, the more congested conditions lead to a much smaller benefit. The detrimental effect of such a strategy at demand level 2 is first obvious at between 7% and 10% take-up - somewhat lower than with demand level 1.

Finally, figures 33 and 34 show (respectively) the system and individual benefits under greater variability in the O-D matrix (ß = 0.7, demand level 1). The greatest system benefit of both actual and marginal cost re-routing is virtually achieved at 5% take-up, with the maximum saving attainable of similar magnitude for these routing criteria. After this stage, as more vehicles become equipped the system travel time increases, with a disbenefit (relative to the "no guidance" situation) becoming

apparent at around 10% - 15% take-up. The total saving in vehicle hours is slightly more here than with the less variable O-D matrix (ß = 0.2), although the benefit relative to the travel time before guidance is virtually the same in the two cases.

Figure 34 suggests a rather different picture for the individual effects of such guidance to that obtained with ß = 0.2. In the case of actual cost re-routing, all the benefit is to equipped drivers, decreasing with level of take-up. Unequipped drivers at the same time experience a slight increase in average travel time due to guidance, the disbenefit increasing with take-up. The pattern for marginal cost routing is not too dissimilar - with large benefits to equipped drivers when only a small percentage have a guidance device, reducing as the level of take-up increases, with the effect on guided drivers again relatively small (two orders of magnitude smaller).

9.   CONCLUSION

A model of a route guidance system has been proposed in terms of a multiple user class equilibrium assignment, with vehicles divided into equipped and unequipped classes, the former being subdivided further dependent on the routing criterion used and the quality of the information supplied. Guidance is used to route vehicles either to a "user optimum" or to a system optimum flow pattern, assuming that without guidance drivers aim to follow a user equilibrium but fail to do so because of perception errors in their evaluation of travel times.   Furthermore, unequipped drivers are assumed to respond to the new route choice of guided drivers, and seek a new user equilibrium routing.

For cost functions of a particular polynomial form, it is shown that such routing strategies, in combination with the route choice of unequipped drivers, are guaranteed to lead to a unique and stable equilibrium flow pattern.

The main advantage of such an equilibrium-based strategy is that it spreads the traffic between multiple routes on each origin-destination movement, and so would be expected to lead to effective guidance even when a high proportion of drivers are equipped with a guidance device. The test runs on two real life networks were used to investigate such a property, as well as the performance of the strategy in a number of different scenarios. It was seen that both user optimal (whether or not subject to predictive errors or to drivers not adhering completely to the recommendations) and system optimal guidance reduced the total system travel time in all situations, the benefit being an increasing function of the level of take-up. The maximum savings were of the order of 5% - 8% reduction in total travel time at 100% take-up; it should be remembered, however, that absolute measures of these may be somewhat misleading as they are highly dependent on the model of unguided drivers - the study of relative benefits is much more appropriate. The level of congestion appeared to have less effect on the benefits, although there was some indication of slightly greater percentage savings in more congested situations below 50% take-up.

Under user equilibrium (UE) routing, individual guided drivers experienced a significant reduction in average travel time, this being approximately constant (of the order of 5%, varying with demand level) for all levels of take-up. Unguided drivers also tended to benefit from such guidance, but always to a much lesser degree than guided drivers (usually less than 1% reduction in average travel time); their savings tended to increase as the percentage of equipped drivers or the level of congestion increased.

System optimal (SO) routing was found to lead to a slightly greater reduction in total travel time than UE routing, particularly in the least congested scenarios. The effects on individual groups of drivers are, however, quite different in the two cases. SO routing was seen to primarily benefit unequipped drivers, significantly improving their position in comparison with UE routing. At lower levels of take-up (10% or less) this saving tends to be at the expense of equipped drivers, who may experience an increase (of as much as 5% in the extreme) in average travel time due to guidance. For higher levels of take-up, equipped drivers will benefit too from SO routing, with the saving in average travel time growing in similarity (lending to the order of 5%) as

the level of take-up increases.

In terms of other performance indicators, it was seen that UE routing decreased the total distance travelled in the system (up to a maximum of around 5% decrease), all the benefit being to equipped drivers, with unequipped drivers tending to be worse off than without guidance in operation. SO routing, on the other hand, led to an increase in total distance travelled of up to 1% at lower (<30%) levels of take-up due to guided drivers being directed to significantly longer (by up to 13%) but less congested routes. There was little effect on unequipped drivers. The greater distance travelled in a shorter time meant that the average speed increased under SO routing by up to 5%. UE routing, too, tends to increase average speed, but to a much lesser extent.

Various features of the proposed model were then investigated further. Firstly, the effects of unforeseen variability in link capacities on network performance were studied, with guidance based, as before, on average (rather than actual) conditions. Results indicated that both strategies were reasonably stable even when unexplained variability was introduced. The SO routing strategy was particularly appealing, as it was seen to lead to a significant reduction (up to 15%) in the standard deviation of total travel time relative to the situation before guidance was in operation.

Secondly, the influence of the assumption was investigated that unequipped drivers re-route in response to guidance. Modifying the model so that instead - at the other extreme - they kept on fixed routes, it was found that very similar results were obtained for UE (and SUE) routing as those which had arisen from the original model. The most significant different was with SO routing at lower levels of take-up, with greater system benefits gained when unequipped drivers kept to fixed routes (it tending to be the equipped drivers who were better off than with the former model).

Finally, the performance of a single route (per O-D pair) strategy, which responds to day-to-day variability in demand and which was developed at Leeds in a different "DRIVE" project, was studied for one of the test networks considered previously. The main findings were that the greatest system benefit of such guidance (whether to minimum actual cost or minimum marginal cost routes) was at around 5% take-up, for the limited number of scenarios considered. For levels of take-up higher than 10%-15%, such guidance is likely considerably to increase total travel time relative to the "no guidance" situation.

To conclude, then, the project has been successful in achieving its aims of developing (multi-route) strategies which are indeed effective at all levels of take-up and congestion, and whose existence and stability is guaranteed by theoretical results.

The limitations of the study have, however, to be recognised, in that the modelling has been based only on "average" (steady state) conditions. Furthermore, the review given in this report of past route guidance modelling work served to illustrate how model dependent results may be. Looking to the future, there is a clear need for a route choice model which is able to take account of the dynamics and uncertainty of urban networks - including "within day" and "between day" variability in network supply and demand, as well as the learning process of drivers as they repeat particular movements. Many route guidance issues still remain, including the

reaction of drivers to system optimal advice (which may sometimes be poor from a users perspective) and the development of strategies which give rise to greater stability in network performance (eg ones which are relatively robust to unforeseen "incidents").

Acknowledgements

References

Al-Deek, H. & Martello, M. & May, A.D. & Sanders, W. (1989) "Potential benefits of in-vehicle information systems in a real life freeway corridor under recurring and incident-induced congestion". Proceedings of the 1st Vehicle Navigation and Information Systems Conference, Toronto, 288-291.

ASTERIX (1989) DRIVE V1054 "System and Scenario Simulation for Testing RTI Systems" Deliverable 3.

Belcher, P.L. & Catling, I. (1987) "Electronic route guidance by AUTOGUIDE: the London demonstration". Traffic Engineering and Control 28, 586-592.

Breheret, L., Hounsell, N.B. & McDonald, M. (1990) "The simulation of route guidance and traffic incidents". Presented at 22nd Annual Universities Transport Studies Group Conference, Hatfield (UK), (unpublished).

CAR-GOES (1990a) DRIVE V1011 "Integration of Dynamic Route Guidance and Traffic Control Systems" Deliverable 14.

CAR-GOES (1990b) DRIVE V1011 "Integration of Dynamic Route Guidance and Traffic Control Systems" Deliverable 9.

Daganzo, C.F. (1982) "Unconstrained extremal formulation of some transportation equilibrium problems". Transportation Science 16, 332-360.

Daganzo, C.F. (1983) "Stochastic network equilibrium with multiple vehicle types and asymmetric, indefinite link cost Jacobians". Transportation Science 17, 282-300.

Daganzo, C.F. & Sheffi, Y. (1977) "On Stochastic Models of Traffic Assignment", Transportation Science 11 (3), pp253-274

Jeffery D. (1987) "Route guidance and in-vehicle information systems". In: Information Technology Applications in Transport, Utrecht, V.N.U. Science Press, 319-351.

Kobayashi F. (1979) "Feasibility study of route guidance system". Transportation Research Record 737, 107-112.

Koutsopoulos, H.N. & Lotan, T. (1989) "Effectiveness of motorist information systems in reducing traffic congestion". Proceedings of the 1st Vehicle Navigation and Information Systems Conference, Toronto, 275-281.

Leonard, D.R., Tough, J.B. & Baguley, P.C. (1978) "CONTRAM - a traffic assignment model for predicting flows and queues during peak periods". Transport and Road Research Laboratory, Crowthorne (UK), TRRL Report LR 841.

Mahmassani, H.S. & Jayakrishnan, R. (1989) "System performance and user response under real-time information in a congested traffic corridor". Presented at

Seminar on Urban Traffic Networks, Capri, Italy, June 1989.

Powell, W.B. & Sheffi, Y. (1982) "The Convergence of Equilibrium Algorithms with Predetermined Step Sizes", Transportation Science 16 (1), pp45-55.

Rakha, R., Van Aerde, M., Case, E.R. & Ugge, A. (1989) "Evaluating the benefits and interactions of route guidance and traffic control strategies using simulation". Proceedings of the 1st Vehicle Navigation and Information Systems Conference, Toronto, 296-303.

Sheffi, Y. (1985) "Urban transportation networks". Prentice-Hall, Englewood Cliffs, New Jersey.

Sheffi, Y. & Powell, W.B. (1982) "An algorithm for the equilibrium assignment problem with random link times". Networks 12, 191-207.

Smith J.C. & Russam K. (1989) "Some possible effects of Autoguide on traffic in London". Proceedings of the 1st Vehicle Navigation and Information Systems Conference, Toronto, 282-287.

Tsuji, H., Takahashi, R., Kawashima, H. & Yamamoto, Y. (1985) "A stochastic approach for estimating the effectiveness of a route guidance system and its related parameters". Transportation Science 19, 333-351.

Van Aerde, M.W. & Yagar, S. (1988)

Van Vliet, D. (1982) "SATURN - a modern assignment model". Traffic Engineering and Control 23, 578-581.

Van Vliet, D., Bergman, T. & Scheltes, W.H. (1986) "Equilibrium traffic assignment with multiple user classes". Proceedings of PTRC Summer Annual Meeting, Seminar M, 111-122.

Van Vuren T., Van Vliet D. & Smith M.J. (1989) "Combined equilibrium in a network with partial route guidance". In: Yagar S. and Rowe S.E. (eds.) "Traffic Control Methods", Engineering Foundation, New York, pp 375-387. Presented at Engineering Foundation Conference on Traffic Control Methods, Santa Barbara.

Van Vuren T. & Watling D.P. (1991) "A multiple user class assignment model for route guidance". Paper to be presented at Transportation Research Board Annual Meeting, Washington, January 1991 (to be published through Transportation Research Record).

Von Tomkewitsch, R. (1987) "ALI-SCOUT: a universal guidance and information system for road traffic". Presented at 6th IEE International Conference on Automotive Electronics, London.

Wardrop, J.G. (1952) "Some Theoretical Aspects of Road Traffic Research", Proceedings, Institution of Civil Engineers II (1), pp.325-378.

Wootton, J. & Ness, M. (1989) "The Experience of Developing and Providing Driver Route Information Systems", Proceedings of 1st Vehicle Navigation and Information Systems Conference, Toronto, pp71-75.

Figure 1: Inefficiency for Weetwood

Figure 2: Inefficiency for Barcelona

Figure 3: System Travel Time, Weetwood Demand Level 1

Figure 4: System Travel Time, Weetwood Demand Level 2

Figure 5: System Travel Time, Weetwood Demand Level 3

Figure 6: Individual Travel Times, Weetwood Demand Level 1

Figure 7: Individual Travel Times, Weetwood Demand Level 2

Figure 8: Individual Travel Times, Weetwood Demand Level 3

Figure 9: System Travel Time, Barcelona Demand Level 1

Figure 10: System Travel Time, Barcelona Demand Level 2

Figure 11: System Travel Time, Barcelona Demand Level 3

Figure 12: Individual Travel Times, Barcelona Demand Level 1

Figure 13: Individual Travel Times, Barcelona Demand Level 2

Figure 14: Individual Travel Times, Barcelona Demand Level 3

Figure 15: Total Distance Travelled, Weetwood Demand Level 1

Figure 16: Total Distance Travelled, Weetwood Demand Level 2

Figure 17: Total Distance Travelled, Weetwood Demand Level 3

Figure 18: Individual Distances Travelled, Weetwood Demand Level 1

Figure 19: Individual Distances Travelled, Weetwood Demand Level 2

Figure 20: Individual Distances Travelled, Weetwood Demand Level 3

Figure 21: Average Speed, Barcelona Demand Level 1

Figure 22: Average Speed, Barcelona Demand Level 2

Figure 23: Average Speed, Barcelona Demand Level 3

Figure 24: System Travel Time, Weetwood Demand Level 1
(fixed route unguided)

Figure 25: System Travel Time, Weetwood Demand Level 2
(fixed route unguided)

Figure 26: System Travel Time, Weetwood Demand Level 3
(fixed route unguided)

Figure 27: Individual Travel Times, Weetwood Demand Level 1
(fixed route unguided)

Figure 28: Individual Travel Times, Weetwood Demand Level 2
(fixed route unguided)

Figure 29: Individual Travel Times, Weetwood Demand Level 3
(fixed route unguided)

Figure 30: System Benefit for Single Route Strategy,
Barcelona Demand Level 1 (ß=0.2)

Figure 31: Individual Benefits for Single Route Strategy,
Barcelona Demand Level 1 (ß=0.2)

Figure 32: System Benefit for Single Route Strategy,
Barcelona Demand Level 2 (ß=0.2)

Figure 33: System Benefit for Single Route Strategy,
Barcelona Demand Level 1 (ß=0.7)

Figure 34: Individual Benefits for Single Route Strategy,
Barcelona Demand Level 1 (ß=0.7)