



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/2242/>

Monograph:

Watling, D. (1990) A Comparison of System Optimal and User Optimal Route Guidance. Working Paper. Institute of Transport Studies, University of Leeds , Leeds, UK.

Working Paper 316

Reuse

See Attached

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



White Rose Research Online

<http://eprints.whiterose.ac.uk/>

ITS

[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.

White Rose Repository URL for this paper:

<http://eprints.whiterose.ac.uk/2242/>

Published paper

Watling, D. (1990) *A Comparison of System Optimal and User Optimal Route Guidance*. Institute of Transport Studies, University of Leeds. Working Paper 316

Working Paper 316

December 1990

**A COMPARISON OF SYSTEM OPTIMAL
AND USER OPTIMAL
ROUTE GUIDANCE**

David Watling

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.

Abstract

The work described in this paper (carried out under the EC 'DRIVE' programme) extends the simulations described in Working Paper 315, with the aim of studying the likely benefits to and reactions of drivers to system optimal (SO) route guidance - in particular, these effects are compared with those obtained under user optimal (UE) guidance. The model used is again one of a multiple user class equilibrium assignment, so that equipped drivers may be directed to more than one route per origin-destination movement.

UE and SO guidance are compared, at different levels of equipped vehicles and demand levels, on the basis of the number of routes they recommend and the similarity of the flows on these routes, as well as link-based properties such as actual flows and queues resulting. These serve to demonstrate the extent to which the routes recommended under UE guidance serve as proxies to those under SO guidance. Secondly, a comparison is made of average (dis)benefits to guided drivers as well as the excess travel time incurred by individual equipped drivers in following SO, as opposed to UE guidance, in order to determine the extent of user sub-optimality of SO routing. Thirdly, input from a parallel DRIVE project, investigating user reactions to guidance information, is used to infer the extent to which drivers are likely to accept the sub-optimality of SO guidance, and the factors which are likely to influence their acceptance. Finally, some preliminary analysis is performed on combined strategies, which aim to strike a balance between the system benefits of SO guidance and the user benefits of UE routing.

TABLE OF CONTENTS

1.Introduction	2
2.Background	2
3.Route guidance model	3
4. Discussion of test results in activity B2.4	4
5.Comparison of route properties under UE and SO routing	5
6.Comparison of link properties under UE and SO routing	8
7.Driver reaction to route guidance	9
8.Acceptance of UE and SO guidance	10
9. Combined user-and system-based strategies	11
10.Conclusion	13
Acknowledgements	14
References	15
Figures	16-28

1. INTRODUCTION

This report describes work carried out at ITS Leeds under DRIVE V1011 (CAR-GOES) activity B2.5, 'Analysis of "Community Criteria" Strategies'. An equilibrium-based model of a route guidance system was developed in activity B2.4 ('Methods for stabilisation of route recommendations'), and simulation on real-life networks used to demonstrate the effectiveness of the multi-route guidance strategies even at high levels of take-up (i.e. with a high proportion of drivers being guided) and of congestion. The simulations are extended here in order to study properties of system optimal (SO) route guidance, and in particular the system and user benefits relative to 'user optimum' (or user equilibrium - UE), routing. The extent to which the routes recommended under UE routing serve as good proxies to those arising from SO guidance is investigated, as well as the magnitude of the disbenefit to the user in being given SO rather than UE advice. Input from DRIVE work area B3 is used to infer the extent to which drivers are likely to accept this sub-optimality. Finally, some initial tests are performed on strategies which aim to strike a balance between the user benefits of UE routing and the greater system benefits of SO routing.

2. BACKGROUND

Looking to the literature, a discussion paper by Boyce (1986) suggests a number of points worthy of investigation in relation to multi-routing guidance strategies. For example, in mentioning the concept of user optimal equilibrium routing, he suggests that 'it should be analyzed to what extent multiple routes are used in the solution of the user optimal model and how the number of routes depends on zone size and level of congestion'. Also, in discussing system optimal routing, he recognises the possibility that some drivers may have to follow a longer route in order that total travel time is reduced; nevertheless, he mentions the possibility of an SO routing strategy giving rise to greater benefits for everyone than UE routing.

It is noted in passing that the existence of at least one (though artificial) network satisfying Boyce's last conjecture above may be established by reference to Braess's paradox (Braess', 1968, and discussed in Sheffi, 1985). In terms of the four link example given by Sheffi, Braess's paradox concerns the construction of an additional (fifth) link in the network; comparing the user equilibrium flow pattern before and after building the link, it is seen that all drivers have a greater travel time in the 'after' situation. In fact, the user equilibrium pattern for the four link network corresponds to a system optimal pattern for the five link case, with the additional link assigned zero flow. Thus, for the five link network, all drivers would have shorter travel times under a SO assignment than under a UE assignment. In one respect this could be regarded as an illustration of an ideal route guidance system: if it were supposed that without guidance in the system, all drivers follow a UE and (with 100% take-up) guidance were used to route according to an SO pattern, then in this example every individual would benefit due to the co-operation between drivers which is possible through such an information system. In another respect, this could be seen as an indication of possible advantages of SO guidance over UE guidance.

In practice, of course, there will be many other issues to consider than in the above example, such as the complex interactions between equipped and unequipped drivers when there is less than 100% take-up. Van Vuren et al (1989) in fact addressed this

problem, by modelling the drivers in a network by a multiple user class equilibrium assignment, in which unguided drivers follow a UE and guided drivers were routed in order to minimise total system travel time (SO routing). In carrying out a theoretical investigation of the positions of the routing patterns for the two groups of drivers, they showed that for any network in practice, under such a model, there could be only at most one route for each origin-destination pair which contained both UE and SO drivers. This means that on any other route on an O-D movement, there may be only UE drivers (experiencing the same travel cost as drivers on the shared route, but with higher marginal costs) or only SO drivers (with the same marginal cost as the shared route, but higher actual cost). It was concluded then that the average cost of the SO drivers must always be greater than or equal to that for UE drivers.

3. ROUTE GUIDANCE MODEL

The model of route guidance to be considered here is proposed in a previous 'DRIVE' deliverable (CAR-GOES, 1991), as well as a working paper (Watling, 1990) and a paper (Van Vuren and Watling, 1991). It is described here briefly, for completeness.

It is assumed that network supply and demand are fixed (i.e. do not vary with time) and that the whole network is available to the guidance system. Average link cost-flow relationships are supplied, with cost here being measured purely in terms of time (and so the words are used interchangeably). The model is based on a four user class equilibrium assignment, with a mixture of deterministic and stochastic costs. Unguided drivers follow a stochastic user equilibrium (SUE) - 'perceived' link costs being independently, normally distributed with means which are the average cost-flow relationships - and (fixed) variances which depend on free flow travel times and a parameter Θ . These form the first user class, with a value for Θ chosen by 'calibrating' the route choice inefficiency of unguided drivers against values reported to have been observed in the literature. The remaining three classes consist of equipped drivers - guided either to a user equilibrium (UE) pattern, a system optimal (SO) pattern or a stochastic user equilibrium (SUE) flow pattern with parameter ψ ($<\Theta$).

The user classes are assigned to the network in interaction with one another - that is, the travel times (and hence the route choice) for one user class are affected by the flows of all user classes, and a combined equilibrium is sought between these classes. There is an assumption here, then, that unequipped drivers will re-route in response to the new behaviour of guided drivers.

It has been shown that a stable, combined equilibrium is guaranteed to exist and be unique for link cost functions c_a of the polynomial form

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

where F_a is the total flow on link a , d_a and b_a are constants and the power k (>0) is a link independent constant.

4. DISCUSSION OF TEST RESULTS IN ACTIVITY B2.4

The above model has been tested extensively on two real-life networks, and the results reported in the references given towards the start of the previous section - the work being

carried out under `DRIVE' activity B2.4. Aspects of these results which are particularly relevant to work under B2.5 will now be discussed.

The networks considered were Weetwood and Barcelona, which - according to the specification of test scenarios for the modelling work in DRIVE V1011 (CAR-GOES, 1990a) - are respectively `medium' and `large'.

The scenarios considered were:

- (a) Weetwood and Barcelona networks
- (b) Demand levels 1, 2 and 3, corresponding to average network speeds before guidance of respectively 35, 25 and 15 km/h (equivalent to 100%, 130% and 160% of the `observed' O-D matrix).
- (c) Levels of equipped vehicles of 0%, 5%, 10%, 20%, 30%, 50%, 70%, 90% and 100%.
- (d) 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).

Of particular interest here is the difference between a strategy in which all equipped drivers are guided according to a UE routing pattern, and one in which guidance is used to route them according to an SO pattern. In terms of the effect of guidance on total system travel time, it was seen that for all the levels of congestion considered, the UE and SO routing patterns grew more dissimilar as the level of take-up (% of drivers equipped) increased. As the congestion increased, on the other hand, the two strategies performed more similarly (for all the levels of take-up) - for example, at the highest demand level (average network speed 15 km/h) considered for Weetwood, the difference between the travel times for the two strategies was always less than 1% of the total.

At a level of aggregation lower than the system-wide comparisons above, it is possible to compare the effect of the strategies on the average travel time, separately for equipped and unequipped vehicles. As such a comparison is particularly relevant for this paper, the graphs showing the results are repeated here for ease of reference (figures 1 to 6). In particular, for individual equipped drivers, it may be seen that for any given level of congestion, the percentage saving in average travel time under UE routing is approximately constant (of the order of 5%) for all levels of take-up considered (from 0.1% - not shown in the graphs - right up to 100% of vehicles equipped). For SO routing, the benefits to equipped drivers are of a similar order of magnitude to those under UE routing for levels of take-up of 50% or more. The greatest difference between the two strategies is the effect on equipped drivers at levels of take-up of 30% or less, this class of drivers often experiencing an average disbenefit; with the extreme being in figure 5, an increase in average travel time of over 5% for Barcelona at demand level 2 (25 km/h average speed) and a level of take-up of 5%.

These test runs - together with the work discussed in section 2 - raise a number of issues to be addressed in this paper. Firstly, how are the reduced average individual benefits of equipped drivers under SO (as opposed to UE) routing distributed amongst the drivers?

In particular, what is the maximum disbenefit relative to UE routing?) How do the routing patterns tend to differ between the two strategies? Does one tend to recommend more routes than the other, and how many of the routes carry a substantially different flow, for the two cases? On a link level, how do the guided flows and travel times compare between the two strategies - are they substantially different, and on how many links? Does one strategy tend to avoid over-capacity links more than the other? How are the answers to these questions related to the level of take-up and the level of congestion?

The analysis of the results described above will be extended in order to answer these questions, which will provide some insight into the first two points in the specification of the project (proximity of UE and SO routing). The final point will be addressed using input from initial evidence obtained by ITS Leeds under DRIVE CAR-GOES activity B3.3 into the acceptance rate of user sub-optimal information.

Finally, the performance of 'combined' strategies will be investigated, which aim to strike a balance between the greater system benefits of SO routing and the greater user benefits (particularly for levels of take-up of less than 30%) under UE routing.

5. COMPARISON OF ROUTE PROPERTIES UNDER UE AND SO ROUTING

In this section, properties of the recommended routes under UE and SO routing will be discussed. Route characteristics are clearly of great importance in the context considered here, but it should be noted that some care has to be taken in interpreting the results obtained. Properties of such an equilibrium assignment guarantee uniqueness only of the link flows, not of the route flows, and it is possible therefore that different route flow patterns will be obtained with different solution algorithms. Furthermore, the technique used to obtain these results - an extension of the 'method of successive averages' - has the property that at iteration n , for a particular user class, the link flows are the average of the flows on that link arising from the n all-or-nothing route flows determined by that stage - that is, a proportion $1/n$ of each route flow contributes to the link flows. For the Weetwood network, then, with 200 iterations performed, it is the case that any 'silly' route chosen during the early stages of the algorithm (when the current estimate of the equilibrium costs is still poor) will contribute a flow of 0.5% to the final estimate of the equilibrium flows for every time it is chosen. For this reason, it was decided to investigate properties only on 'dominant' routes - that is, those routes contributing at least 5% to the final user class flows. In this way, it was hoped also that the problem of non-uniqueness of the route flows would be avoided to some extent, since only the more 'obvious' routes were being considered, which could be expected to be chosen in any solution algorithm. As a further restriction, only dominant O-D movements were considered - that is, those with a total flow of more than ten.

Now, results discussed in the previous section indicated that the most interesting area for a comparison of the benefits to guided drivers under UE and SO routing was at levels of take-up of 30% or less. Since a detailed comparison of all the scenarios described in section 4 would produce too many results to assimilate easily, it was decided to restrict attention to a subset of these scenarios, according to the observation made at the start of this paragraph:

- (a) Weetwood network
- (b) Demand levels 1 and 2, corresponding to 100% (35 km/h) and 130% (25 km/h) of the observed O-D matrix

- (c) Levels of equipped vehicles of 10% and 30%
- (d) Two different routing criteria - with all equipped vehicles guided either as a UE pattern or as a SO pattern

Considering firstly demand level 1 at 10% take-up, figure 7 displays the frequency distribution of the number of recommended routes per origin-destination (O-D) pair, with respect to the number of O-D pairs, under each of the routing strategies. There is a slight shift in frequency from one recommended route to two recommended routes in supplying SO rather than UE guidance. As summary measures of the graphs, the mean number of routes per O-D pair is 1.27 under UE routing and 1.33 under SO routing. A similar pattern in a shift of the number of recommended routes - to a greater frequency of O-D movements, with more than one recommended route under SO routing - is evident in figures 8 (30% take-up, demand level 1), 9 (10% take-up, demand level 2) and 10 (30% take-up, demand level 2). The average number of recommended routes are respectively 1.29 and 1.54 for UE and SO routing in figure 8; 1.35 and 1.39 in figure 9; and 1.39 and 1.78 in figure 10. Although not shown in the figures, a 50% take-up at demand level 1 was also studied, where the average number of route per O-D movement was 1.32 (UE) and 1.86 (SO). This is consistent with the pattern that under UE routing, the average number of routes changes very little with level of take-up, but increases quite sharply under SO routing. There also appears to be a greater average number of routes with increased congestion although there may be some influence here of the way in which only dominant O-D movements are considered.

The results also have implications for the implementation of such multi-route strategies - with a number of O-D movements requiring as many as seven recommended routes with a substantial percentage of the O-D flow. Although the strategies are in general multi-route ones, it may be seen that of the order of 70% of O-D pairs require only a single recommended route; the great majority of the remainder require two routes. This may suggest that, should the implementation of a true equilibrium strategy be considered impracticable, then a multi-route strategy based on at most two recommended routes per O-D pair may prove to be effective - although this clearly needs investigation.

A comparison of the route flow patterns under the two strategies was carried out, with each origin-destination movement categorised according to whether the same routes were used in the same proportions under UE and SO guidance, or the same routes were used but with different proportions of the O-D flow, or the routing patterns were different in that at least one route was not common to the two strategies. Figure 11 contains the results of these comparisons, for both demand levels and for both levels of take-up. As before, only origin-destination movements with a total flow greater than ten are considered - hence the different total number of O-D pairs for the different demand levels. All routes are considered, however - not just the dominant ones. The results show little difference between different levels of take-up. In all cases, around 130 O-D movements have exactly the same route recommendations under the two types of guidance, and around 30 O-D pairs use the same routes but in different proportions. The remaining 60-70% of O-D movements are made using different routing patterns for equipped drivers under the two strategies.

Figure 12 compares the route flow patterns once again, but the aim here is to determine whether the patterns are substantially different under the two strategies. Firstly, only dominant (>5% of flow) routes are considered. Secondly, it is tested whether UE and SO guidance recommend the same routes in similar proportions. For a given O-D pair, two routing patterns are said to be similar if they use the same dominant routes and if the

percentage of O-D flow on any given one of these routes differs by less than five between the two routing patterns. The categories are now: same dominant routes in similar proportions, same dominant routes in dissimilar proportions and different dominant routes. Again, there is little difference in the results between levels of take-up. For approximately a half of the origin-destination movements, the routing patterns under UE and SO guidance are approximately the same. Virtually all of the remaining O-D movements use different dominant routes under the two strategies; hardly any (of the order of 10 O-D pairs) use the same set of routes but in substantially different proportions. This is consistent with previous findings for the same network, (Watling, 1990) which indicated that under SO routing, equipped drivers tended on average, to travel a substantially greater distance than without guidance, whereas UE routing gave rise to a saving in distance travelled. This was the first indication that the actual routes used under the two strategies had to be substantially different, at least for some O-D movements.

The final, and perhaps most interesting, comparison of route attributes concerns the travel times for guided drivers under the two strategies. Again considering only dominant routes and O-D flows, the results in the bar chart in figure 13 (demand level 1, 10% take-up) give - for each route recommended under SO routing and each O-D pair - the percentage excess travel time of a vehicle following an SO route relative to the travel time on any UE route for that O-D movement. (Recall that in theory, for a given O-D relation, all UE routes should by definition have the same travel time; in practice, of course, this will not be exactly the case, and so the travel time used was in fact the average of the estimates of the UE route travel times at termination of the solution algorithm).

For the scenario considered here, the excess travel time on an SO route ranged from -3.0% to 60.9% of the time under UE routing. That is, some drivers experience a disbenefit of as much as 60% by being routed according to system, as opposed to user, objectives. For ease of display, disbenefits larger than 25% are grouped together in the plot. Studying the general shape of the graph (ignoring the final group in order to obtain a true histogram), it may be seen that the distribution is very skew. For almost half of the equipped drivers, the travel times under SO routing are within 1% of those under UE routing, and these are unlikely to detect any difference between the two strategies on an individual level. Of the order of 7% of equipped drivers are in fact better off by between 1% and 3% under the SO strategy. The remaining guided drivers (over 40% of the total number equipped) lose out in following SO rather than UE guidance, with about 5% of them taking more than 15% longer in this case.

Figure 14 illustrates the corresponding results for demand level 1 and 30% take-up. In this case, the excess travel times range from -11.0% to 67.9%. The general impression is that the distribution is significantly less skew than for the 10% take-up case. Again, around 7% of the equipped drivers make a significant (>1%) gain following SO rather than UE guidance. The benefits are somewhat greater than at 10% take-up, with the great majority saving between 1% and 5% in travel time relative to UE routing. Around 40% of drivers are unlikely to perceive any difference in being routed according to the two strategies (with less than 1% difference in travel time). Over half of the equipped drivers lose out significantly under SO routing, with around 5% of them experiencing an increase in travel time in excess of 15% of that under UE guidance. Although the figures are not given here, the pattern for 50% take-up was consistent with these findings, in that the distribution became more symmetric with an increase in the level of take-up; at the same

time, the maximum possible benefit of following SO, as opposed to UE, advice also increased (with the range of excess travel times at 50% take-up being -15.8% to 69.1%) and the number of people losing out from such guidance decreased.

For demand level 2, figure 15 (10% take-up) and 16 (30% take-up) demonstrate a similar sort of pattern of a decrease in skewness and increase in the magnitude of the benefit for those who are actually at an advantage following SO rather than UE advice. With 10% of drivers equipped, the range of excess journey times for SO routing is -3.6% to 76.8% of the UE route times. Again for approximately half of the equipped group, the SO route time will be within 1% of the UE time. Only around 3% of guided drivers benefit significantly more (>1%) under the SO strategy; on the other hand, 3% of this group take more than 25% longer under this guidance criterion. In total, around half of the drivers take significantly longer in the SO than in the UE case, At the higher, 30%, level of take-up, the range of excess journey times for SO routing widens, with a minimum of -22.0% and maximum of 78.3%. Around 40% of equipped drivers experience no more than a 1% change in travel time in following SO as opposed to UE routing; about 20% of drivers benefit by more than 1% in preferring the SO strategy; and for approximately 3% of equipped drivers, their travel time is more than 25% greater under SO guidance than under UE guidance.

To conclude, the above results indicate that to obtain the greater system benefits of SO routing, a small sub-group (of the order of 4%) of the equipped drivers will have to follow routes which are more than 25% longer (in terms of travel time) - and, at the extreme, more than 70% longer - than the UE routes. On the other hand, for about half of the equipped drivers there is little difference in following either of the guidance strategies, and for some equipped drivers (up to 20% in some cases) there is a significantly greater benefit under the SO strategy, particularly at higher levels of take-up.

6. COMPARISON OF LINK PROPERTIES UNDER UE AND SO ROUTING

Link-based attributes of the strategies will now be considered: in particular, the flows of guided drivers; the link travel times; and the links which carry a flow which is over capacity.

Figure 17 - corresponding to demand level 1 at 10% take-up - shows that around 50% of the guided link flows under UE routing are within $\pm 5\%$ of those under SO routing. The ratio of the remaining flows are spread reasonably evenly over a wide range - with, for example, around 10% of the guided link flows under SO routing being as much as 100% larger than those under UE routing. A very similar pattern was observed for the other three scenarios, with there being no obvious dependence on congestion or level of take-up; for this reason the results are omitted.

Studying next the link travel times under the four scenarios, it was seen that around 90-95% of links had times under SO routing which were within 5% of those under UE routing. The range of the ratio of the time under SO guidance to that under UE guidance varied considerably with the level of take-up: for demand level 1 at 10% take-up, this range was 0.80 - 1.20 (96% of links within 0.95 - 1.05), and at 30% take-up it widened to 0.55 - 1.45 (93% of links within 0.95 - 1.05); for demand level 2, with 10% of drivers equipped the range was 0.80 - 1.15 (95% within 0.95 - 1.05) and with 30% equipped it was 0.60 - 1.45 (91% within 0.95 - 1.05).

Considering, finally, the links which carry a flow which is over capacity. The figures below give the number of links which are over capacity under the two strategies, and the total number of vehicles queued on these links at the end of the simulation period (that is, the difference between the demand flow and the capacity):

Scenario queue	Guidance	No. of links	Total
DL 1, 10% take-up 71.7	UE	6	
	SO	4	
51.8	UE	5	
	SO	2	
DL 1, 30% take-up 61.6	UE	5	
	SO	2	
26.1	UE	25	1429.4
	SO	26	1427.7
DL 2, 10% take-up	UE	24	1386.2
	SO	24	1340.6

It can be seen that under both guidance strategies, the queues become smaller as the level of take-up increases. SO routing, however, leads to less over-capacity flow than UE

routing in all four scenarios, with the greatest difference (in percentage terms) at the lower demand level, where there is greater potential for the SO strategy to find uncongested (but greater distance) routes, which avoid congested links.

7. DRIVER REACTION TO ROUTE GUIDANCE

The aim of this section is to tie together the results of the simulation discussed so far in this paper, with the findings of a parallel DRIVE V1011 activity, B3.3 "Drivers' reactions to accuracy of guidance information" (CAR-GOES, 1990b; also see Bonsall 1991a and 1991b). This study was concerned with driver reaction to the accuracy of guidance information - in particular, the results obtained using the interactive simulation 'game' IGOR (Bonsall and Parry, 1990) lead to some interesting conclusions.

Briefly, the approach for using IGOR to collect data on driver responses to route guidance was as follows. A hypothetical network was displayed on a computer screen, and the user allowed to make a number of journeys without guidance, in order to gain some familiarity with the network and the sort of travel times which may be encountered (the travel times vary randomly, in order to simulate variations in traffic conditions, and so the user will experience different travel times on different journeys on the same route). The user then made a number of journeys with guidance advice given at each junction (which may or may not suggest the minimum time route given current conditions), accepting or rejecting the advice given. In the study, over 300 participants from the UK and France were involved, with over 10,000 route choice decisions made.

Many factors were investigated for influence on the acceptance of advice, but the most interesting in the context of the study reported in this paper were: the quality of (current) advice; the quality of previous advice; and the presence of visible congestion at the point of the route choice decision.

Now, on the first point, each time a route was advised, a measure called the "duffness" was calculated, being the ratio of the time to the destination on the recommended route to the time on the minimum time route, expressed as a percentage (that is, in the context of figures 13 - 16, duffness = % excess journey time + 100%). The results showed that in the case where the advice given was to follow the minimum time route, the recommended route was followed on around 82% of occasions. A huge range of sub-optimal advice was considered with routes recommended whose journey time was more than four times that of the minimum time route. The results were displayed as a graph of acceptance of advice against duffness with an apparently sharp decrease in frequency of acceptance with an initial increase in duffness from the perfect information, 100% duffness case [the raw data were in fact grouped, with the obtained by plotting the midpoints of each class. Looking to the data, a great many users were given perfect advice, but very few recommendations were given for routes with only a small - up to 5% - excess time, making the data somewhat less reliable in this range.

It is not clear, therefore, to what extent routes which are only slightly slower than the minimum time one are accepted]. For example, around 70% of advice was accepted which recommended a route which was (exactly) 10% slower than the minimum time one; with this acceptance frequency falling to 55% with a 20% excess journey time, to around 35% with a 50% excess time, and to an acceptance of about 25% at 70% excess time.

The second point of interest is the influence of the quality of previous guidance information on acceptance. For studying such effects, it is useful to use an 'inverted' form of a graph given by Bonsall, which appears in the project report (CAR-GOES, 1990). Here, acceptance of advice is given as a function of the quality of current advice, for various given levels of previous advice. It is seen that if previous advice had been good (excess

journey times of no more than 5%), then whilst acceptance of current advice for a 10% excess time route was still around 70%, the acceptance profile for poorer current advice flattened out. Under such a quality of previous advice, current advice which gave rise to a 30% excess time was accepted on only slightly less than 70% of occasions, and current advice for a 70% excess time route would have an acceptance of as high as 65%. It is noted, on the other hand, that there was apparently little effect of being previously given good advice on drivers currently being given optimal advice - their acceptance was still around 80%.

Furthermore, the results indicated that if previous advice had been somewhat poorer (5%-15% excess journey time), the acceptance of (current) advice which directed users to a route with a 10% excess time would fall to about 60%, and for a 30% excess time the acceptance would be less than 50%. With even poorer previous advice, the acceptance of current advice was seen to decrease further.

Finally, the effect of corroborating information on acceptance was studied using a stated preference experiment, conducted after the 'journeys' made on the simulator. As a 'worst case' it was found that only 10% of advice would be followed if guidance directed the driver to the left, but a road sign pointed right, the 'crow-fly' direction was right, drivers ahead were turning right, there was a major road to the right and a minor road to the left, and traffic appeared congested to the left but free-flowing to the right. The experiment investigated the influence of each of these factors in turn, and found that the one which corroborated the guidance advice the most was the fact that the following road on the recommended route appeared to be free-flowing.

8. ACCEPTANCE OF UE AND SO GUIDANCE

Bringing together the results of the last two sections, it is possible to draw a number of conclusions regarding the acceptance of both UE and SO guidance.

Firstly, if UE guidance in the network model is equated to guidance to the minimum time route in the IGOR study, then it may be concluded that UE routing advice will be accepted on around 82% of occasions. For SO guidance, on the other hand, reactions will tend to be rather different, depending on individual routes and how much longer (time) they are than the minimum time one. For example, for routes with a 10% excess time (relative to the time under UE guidance), acceptance falls to 70%. Around 95% of guided drivers under SO routing were seen to have an excess travel time of less than 15%, and for these the frequency of acceptance will be at worst around 65%. For the remaining minority of guided drivers, in particular those losing out the most - with an excess time of 50%-70% - the advice is followed on around 25%-35% of occasions.

As a whole, then, the results appear to indicate that the acceptance of SO routing advice will compare reasonably well with that of UE routing for the great majority (95%) of drivers who do not lose out by more than 15% in following SO guidance. However, at the other extreme, for drivers losing out the most under SO guidance, the frequency of acceptance may be only a third that of UE guidance.

Secondly, the behavioural work indicated that if previous advice had been good, then advice which was sub-optimal to a large extent (up to 70% greater travel time over that of the optimal route) would be accepted at a frequency only slightly less than that of UE

guidance (65% acceptance as opposed to 80% acceptance overall of UE guidance). This suggests that acceptance would be good, were a plan to be implemented in which UE routing were usually recommended, except when - say - a severe incident occurred and a switch could be made to SO routing in order to regain stability in the system. The result showed, on the other hand, that frequent use of SO advice could result in a poor acceptance rate - for example, for drivers previously advised routes with a 5%-15% excess time, current advice to use a route with 30% excess journey time would be accepted on less than 50% of occasions, and for poorer qualities of previous/current advice the acceptance would be significantly less.

Thirdly, and finally, other features of SO guidance may make it more appealing to the user (relative to UE guidance). It was seen, at the network modelling stage, how SO routing tended to avoid queues to a greater extent than UE routing. The behavioural work, on the other hand, showed that drivers were more likely to accept guidance if the next road on the recommended route was free-flowing, this being the most influential piece of corroborating evidence of those studied. It would be expected, then, that SO routing would be accepted more often than other user sub-optimal strategies.

9. COMBINED USER-AND SYSTEM-BASED STRATEGIES

A key issue in implementing any route guidance system is the guidance strategy to be used. If the guidance equipment is to be sold to drivers, then there is a clear conflict of interest between the equipped drivers who want to improve their own situation, and the traffic authorities, whose objectives include reducing overall traffic congestion. The comparisons made in this paper appear to indicate that system optimal guidance is likely to lead to large disbenefits to some equipped drivers, relative to following user optimal guidance; however, benefits to the system are, of course, larger under system optimal guidance. A natural solution to this problem would seem to be to seek a strategy which combines the objectives of the user and the system in some way; for example, is it possible to improve on the system performance of user optimal guidance without affecting the user benefits to too great an extent? Although the investigation of such strategies did not form part of the specification for the project reported here, this was seen as a natural progression in the light of the results obtained so far. Initial evidence of the performance of two 'combined' strategies is given below.

The first strategy considered is a 'mixed' one. This consists simply of a given proportion α ($0 \leq \alpha \leq 1$) of the equipped drivers being SO routed and a proportion $1 - \alpha$ being UE routed. This strategy falls quite naturally into the general four user class model described in section 3, and therefore existence and stability properties of the combined equilibrium apply equally well to this strategy. It is easily seen that with this mixed strategy, if $\alpha = 0$ a full UE routing strategy is obtained, and if $\alpha = 1$, then the strategy is SO routing.

The second strategy is a 'compromise' one, and this requires a modification to the four user class model considered previously. Now, it is recalled that within the multiple user class (MUC) framework proposed for route guidance (Watling, 1990), a UE routing strategy is implemented by performing a MUC equilibrium assignment with a cost for equipped drivers on link a of

$$c_a(F_a) \quad [\text{UE routing}]$$

where F_a is the total flow on link a . To implement an SO routing strategy, on the other hand, the same equilibrium assignment is performed, with a cost for equipped drivers on link a of

$$c_a' = c_a + F_a \frac{dc_a}{dF_a} \quad [\text{SO routing}]$$

- that is, the marginal cost of travel on link a corresponding to an actual cost of c_a .

The compromise strategy to be considered here will have a cost of

$$c_a' = c_a + \alpha F_a \frac{dc_a}{dF_a} \quad [\text{UE/SO compromise}]$$

for some specified parameter α ($0 \leq \alpha \leq 1$).

Although this is a different definition for α to that of the mixed strategy, it is convenient for later discussions to give the two parameters the same letter. In both cases, it is the case that $\alpha = 0$ corresponds to a UE routing strategy, and $\alpha = 1$ to SO routing.

The compromise strategy is implemented by replacing the costs c_a' of the user class for SO routing with the costs c_a'' above. Existence and stability of the combined equilibrium for this modified model are still guaranteed for the polynomial cost function (1) used, since the costs for the 'compromise' class are then

$$c_a'' = d_a + (\alpha k + 1) b_a F_a^k$$

and so fall into the family of cost functions used to infer theoretical properties of the standard four user class model.

The two strategies described above were applied to the scenarios specified in section 5, for values of α of 0, 0.2, 0.4, 0.6, 0.8 and 1. The effect on total travel time for each of the scenarios is given in figures 18-21, with the corresponding average effects on individual equipped drivers in figures 22-25.

The performance of the mixed strategy is in general disappointing. In terms of total travel time, the strategy often improves on UE routing, but there is no general pattern as to a value for α which will achieve an improvement in all situations; and even for a given scenario, the system performance may be quite unstable with respect to α . More importantly, the strategy does not appear to have achieved its goal in terms of its effect on individuals, with equipped drivers following the SO advice under the mixed strategy losing out more on average than under a full SO routing strategy (and those following the UE advice winning more on average than under a full UE routing).

Considering the compromise strategy, the most interesting results appear to be (for all scenarios) those obtained with $\alpha = 0.2$, where system benefits are close to, or even slightly better than, those of SO routing ($\alpha = 1$), and the average benefit to equipped drivers is reduced relative to UE routing by less than 11%. That is, under this strategy equipped drivers still save between 3% and 5% on average travel time. Increasing α further, however, reduces the benefit to equipped drivers, without offering any significant improvement in system travel time. The fact that the compromise strategy with $\alpha = 0.2$

may sometimes give rise to a smaller total travel time than a pure SO routing may be due to an apparently slow rate of convergence of the algorithm under SO routing (identified in a previous paper - Watling, 1990 - and discussed in more detail by Van Vuren and Watling, 1990). It would seem, then, that the compromise strategy with $\alpha = 0.2$ also has the advantage of being more efficient than SO guidance, in that a better estimate of the equilibrium appears to arise for a given amount of computation time (that is, for a given number of iterations).

10. CONCLUSION

This paper has been concerned mainly with the comparison of the effects on equipped drivers of two route guidance strategies - user optimal and system optimal. The route guidance model is formulated as a multiple user class equilibrium assignment, in order to allow multiple routing between an origin-destination pair (to ensure stability) and to represent interactions between equipped and unequipped drivers.

Applying the model to many scenarios (level of take-up, congestion level, etc.) on two real-life networks, it was seen that the average effect of guidance on an equipped driver was quite different under the two guidance criteria for levels of take-up of 30% or less. In general, under UE guidance, the average travel time saving for equipped drivers was approximately constant (of the order of 5%) for all levels of take-up; under SO routing, on the other hand, equipped drivers could even experience an average disbenefit at lower levels of take-up, although as more vehicles become equipped, the average benefit to guided drivers approached that obtained under UE routing.

Studying a subset of these scenarios, so that effects could be disaggregated further, route-based properties of the equipped drivers were firstly investigated. It was seen that the number of recommended routes per origin-destination pair tended to be higher under SO than under UE guidance, although at lower levels of take-up (10%) the difference was only marginal. It was seen that for both strategies, there tended to be a single dominant route (i.e. a route with greater than 5% of the O-D flow) on around 70% of O-D movements, with the great majority of the remainder using only two dominant routes. In practice, when new recommended routes need to be computed on-line in response to dynamic changes, the computation of a true equilibrium routing pattern may be infeasible; however, the fact that only a small number of dominant routes are used was seen as suggesting that strategies which are allowed to use at most, say, two or three recommended routes per O-D pair may be effective under high levels of take-up. Such strategies need to be developed, but could take the form of an heuristic method coupled with a small number of iterations of an equilibrium solution algorithm, or possibly some form of incremental assignment.

Returning to the comparison of UE and SO routing, it was seen that in all the scenarios considered, for around half of the origin-destination pairs, the route flows of guided drivers were similar under the two strategies; for the remainder, either the flows or the routes were substantially different.

In terms of the route travel times under the two strategies, it was seen that whilst 40%-50% of guided drivers were unlikely to perceive any difference between UE and SO routing (difference in travel time of less than 1%), of the order of 5% of equipped drivers lost out by more than 25% in following SO rather than UE advice, with a handful of drivers losing out by as much as 65% in being routed according to an SO pattern. On the other hand, a number of guided drivers (up to 20% of them, in the highest demand/level of take-up scenario) will actually be significantly better off in the SO than in the UE case, due to their co-operation in reducing congestion.

Of the link-based properties studied, the most interesting was a comparison of over-capacity links under the two strategies. It was seen that for both strategies and both demand levels, the total queues became smaller as the level of take-up increased. SO routing, however, always led to less over-capacity flow than UE routing, with the greatest

(percentage) difference at the lowest demand level, where it is still possible for the SO strategy to find less congested routes.

Drawing on findings from a parallel 'DRIVE' project into user acceptance of route guidance advice, and particularly its relation to the quality of the advice given, it was concluded that for around 95% of drivers, whose recommended route under SO guidance was no more than 15% longer than the minimum time route, the acceptance would be at worst 60%, but that for the remaining drivers who are given the worst advice under SO routing, the acceptance could fall to a third that under UE routing (that is, 25% of advice accepted, versus 80% acceptance if the advice were user optimal). However, taking into account the influence of the quality of previous advice, it was concluded that if SO routing were only to be used in favour of UE routing on fairly infrequent occasions (say, in a severe incident), then the user acceptance of the two strategies was likely to be quite similar (even for drivers losing out by as much as 70% in following SO as opposed to UE routing, whose acceptance would be around 65% if their previous advice had been good).

Finally, a short study was carried out on a 'compromise' strategy, which appeared to succeed in striking a balance between the system benefits of SO routing and the user benefits of UE routing and it was shown how this strategy could be incorporated in the multiple user class framework used previously.

In the short duration of this project, it has only been possible to compare user and system optimal routing in a fairly small number of scenarios - there is a need for further simulation, to verify that the effects observed here are reproduced in other networks/conditions. Likewise, the compromise strategy appears to offer a promising means for controlling the trade-off between system and user benefits, but requires more detailed analysis (e.g. of the route travel times relative to UE routing).

In terms of the real-life implementation of multiple route strategies of the kind proposed here in a real-time system, it may be computationally more practicable to restrict the number of recommended routes per movement to a small number - say 2 or 3 - and use a semi-heuristic equilibrium-like approach to determine them. The simulations carried out in this paper appear to indicate that such an approach would still be effective, since it has been demonstrated that even with a true equilibrium routing, only a small number of routes for each origin-destination movement carry a significant proportion of the flow.

The link between the modelling work carried out here and the behavioural research of activity B3.3 suggests that at the very least some refinements to current route guidance models are required in order to reproduce the likely effects of guidance (and particularly SO guidance). Whilst the fact that not all drivers will follow the route recommendations can be simulated to some extent by perturbing the costs as perceived by users (see the SUE routing strategy of Watling (1990)), this is clearly a great over-simplification, since as one would expect and as has been confirmed by the behavioural work, there is a relationship between quality of advice and acceptance. The influence of the quality of previous guidance advice is also clear from B3.3. Both of these issues appear to call for a new approach to the modelling of route guidance systems, in order to obtain reliable estimates of the potential benefits.

ACKNOWLEDGEMENTS

This work was carried out under the funding of DRIVE project V1011 "Integration of dynamic route guidance and traffic control systems", activity B2.5.

I would like to thank Tom van Vuren for his help and suggestions during this project, and Tim Parry and Peter Bonsall for their discussions with me on the results of the IGOR simulation game.

REFERENCES

- Bonsall, P.W. (1991a) "Response to Route Guidance", unpublished paper presented at Universities Transport Study Group annual conference, University of Nottingham, 2-4 January, 1991.
- Bonsall, P.W. (1991b) "Using an Interactive Route-Choice Simulator to Investigate Drivers' Compliance with Route Guidance Advice", paper to be presented at Transportation Research Board Annual Meeting, Washington DC, 14-17 January, 1991.
- Bonsall, P.W. & Parry, T. (1990), "A Computer simulation game to determine drivers' reaction to route guidance advice", Proc. of PTRC Summer Annual Meeting, Seminar H (Transportation Planning Methods), University of Sussex, Brighton, September 1990.
- Boyce, D.E. (1986). "Combining Communication and Transportation Technology to Improve Urban Travel: A New Look at the Urban Transportation and Urbanisation Problem", International Symposium on the Development of Information Technology and its Impact on the Urban-Environmental System, Toyohashi University of Technology, Toyohashi, Japan, November 1986.
- Braess, D. (1968), "Über ein Paradox der Verkehrsplanung" *Unternehmensforschung* 12, pp. 258-268.
- CAR-GOES (1990a) DRIVE V1011 "Integration of dynamic route guidance and traffic control systems", deliverable 9.
- CAR-GOES (1990b) DRIVE V1011 "Integration of dynamic route guidance and traffic control systems", deliverable 20.
- CAR-GOES (1991) DRIVE V1011 "Integration of dynamic route guidance and traffic control systems", deliverable ?
- Sheffi, Y. (1985) "Urban Transportation Networks", Prentice-Hall, Englewood Cliffs, New Jersey.
- 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.
- 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).
- Watling, D.P. (1990) "Route guidance algorithms effective for all levels of take-up and congestion", Institute for Transport Studies, University of Leeds. Working Paper No. 315.

Figure 1: Individual travel times, Weetwood demand level 1.

Figure 2: Individual travel times, Weetwood demand level 2.

Figure 3: Individual travel times, Weetwood demand level 3.

Figure 4: Individual travel times, Barcelona demand level 1

Figure 5: Individual travel times, Barcelona demand level 2

Figure 6: Individual travel times, Barcelona demand level 3

Figure 7: Number of recommended routes, Weetwood demand level 1, 10% take-up

Figure 8: Number of recommended routes, Weetwood demand level 1, 30% take-up

Figure 9: Number of recommended routes, Weetwood demand level 2, 10% take-up

Figure 10: Number of recommended routes, Weetwood demand level 2, 30% take-up

Figure 11: Comparison of route flows under UE and SO routing, Weetwood network

Figure 12: Comparison of principal route flows under UE and SO routing, Weetwood network

**Figure 13: Comparison of route travel times under UE and SO routing,
Weetwood demand level 1, 10% take-up**

**Figure 14: Comparison of route travel times under UE and SO routing,
Weetwood demand level 1, 30% take-up**

**Figure 15: Comparison of route travel times under UE and SO routing,
Weetwood demand level 2, 10% take-up**

**Figure 16: Comparison of route travel times under UE and SO routing,
Weetwood demand level 2, 30% take-up.**

**Figure 17: Comparison of link flows under UE and SO routing,
Weetwood demand level 1, 10% take-up.**

**Figure 18: Total travel time under combined strategies,
Weetwood demand level 1, 10% take-up.**

**Figure 19: Total travel time under combined strategies,
Weetwood demand level 1, 30% take-up**

**Figure 20: Total travel time under combined strategies,
Weetwood demand level 2, 10% take-up**

**Figure 21: Total travel time under combined strategies,
Weetwood demand level 2, 30% take-up**

**Figure 22: Travel times for guided drivers under combined strategies,
Weetwood demand level 1, 10% take-up**

**Figure 23: Travel times for guided drivers under combined strategies,
Weetwood demand level 1, 30% take-up**

**Figure 24: Travel times for guided drivers under combined strategies,
Weetwood demand level 2, 10% take-up**

**Figure 25: Travel times for guided drivers under combined strategies,
Weetwood demand level 2, 30% take-up.**

