The DENVIL COOMBE Practice
Weybrook, Bunch Lane, Haslemere, Surrey, GU27 1ET
Tel/fax: 01428 658146
Email: denvilcoombe@ukonline.co.uk

Client:	Integrated Transport Economics Appraisal Division Department for Transport
Project:	Speed/Flow Relationships in Traffic Models
Document type:	Final Report
Date:	27 February 2007
Status:	Final version

INTRODUCTION

1 ITEA Division of the Department for Transport wishes to understand the capabilities of the transport modelling packages available to transport modellers in the UK to model speed/flow relationships separately by vehicle type and has commissioned me (Denvil Coombe) to conduct this investigation.

2 A specification of the functionality that the Department requires was sent to the known software providers with a request for evidence about the capabilities of their software to model speed/flow relationships separately for light vehicles and heavy vehicles. The required relationships are specified in Appendix E from Traffic Appraisal in Urban Areas, which is Volume 12, Section 2, Part 1 of the Design Manual for Roads and Bridges (DMRB). These relationships are the same as those used in COBA but are referred to in this report as the 'DMRB' speed/flow relationships.

3 This report is structured as follows:

- the **information provided by the software suppliers** is considered first, in alphabetic order of the software names; followed by
- my **summary of the findings from the software suppliers**, along with my **conclusions** drawn from the information obtained.

4 The statements in this Report which summarise the **information provided by the software suppliers** were sent, in draft, to the software suppliers for verification. All but Caliper responded and the suggested amendments have been reflected, where appropriate, in this final version. PTV and Peter Davidson Consultancy advised that they had changed their software since my first approach to them; however, no dialogue has taken place with these suppliers (as was the case with most of the other suppliers who said that their software could handle the DMRB relationships).

INFORMATION PROVIDED BY THE SOFTWARE SUPPLIERS

CONTRAM, supplied by Mott MacDonald and TRL

5 Mott MacDonald advised that CONTRAM allows modellers to specify different speed/flow curves for different assignment user classes. Like the DMRB curves, the CONTRAM curves are piecewise linear. Each curve can have up to 5 straight line segments, with the user specifying the free flow speed, minimum speed and up to four (flow, speed) breakpoints. These can be calculated outside CONTRAM according to the DMRB formulae, based on geometry, road type, the percentage heavy goods vehicles, etc. Once

the shape of the curve is defined, the user can choose whether it applies to one, some or all of the assignment user classes. The flow used to calculate the speed is always the total flow (in vehicles) on the link, i.e. summed over all user classes.

6 Mott MacDonald also explained that CONTRAM divides the OD flow for each user class into discrete packets of vehicles. For each packet in turn, CONTRAM calculates the current minimum cost route for that packet and, if different from the current route, switches that packet to the new route. The minimum cost route is calculated taking into account what the cost would be on that route **after** the packet has switched routes.

7 The process is essentially the same however many user classes are modelled. Multiple user classes just mean that the generalised cost function may differ between packets. Similarly, different speed/flow curves by user class will just mean that different packets will face different travel times on the same link. However, that is usually the case anyway due to the dynamic nature of CONTRAM - different packets face different times on a given link because they arrive on the link at different times and travel time/delay is a continuous function of arrival time.

8 Traditionally, convergence in CONTRAM has been based on a number of stability measures, although there is now a facility to calculate the gap statistic.

9 There is no proof of convergence for the method, and therefore no list of conditions for convergence that may be violated by having different speed/flow curves by user classes. Mott MacDonald also advise that, in practice, the method seems to work well and convergence problems can usually be solved by reducing the size of the packets.

10 **It therefore seems clear that CONTRAM <u>can</u> meet the DMRB specification for separate speed/flow relationships by vehicle type and that, in using speed/flow relationships which differ by vehicle type, there would be no special convergence problems to face**.

CUBE VOYAGER, supplied by Citilabs

11 Citilabs advised that CUBE VOYAGER can accommodate speed/flow relationships by different vehicle types. The modeller may include any number of user classes (vehicle types, journey purposes) and specify their own speed/flow relationships as a mathematical equation. These equations can be set up using any variables that can be found in the network, and the network format and number and types of variables are for the modeller to decide.

12 Citilabs also advised that convergence methods which rely on the minimisation of some kind of objective function, such as the Frank-Wolfe method, may not be appropriate for use with multi-class assignment. Citilabs suggested that the method of successive averages (MSA) can be used to achieve convergence with multi-class assignments but did not comment on whether the use of speed/flow relationships which differ by vehicle type would make the achievement of convergence either easier or more difficult.

13 **It therefore seems clear that CUBE VOYAGER <u>can</u> meet the DMRB specification for separate speed/flow relationships by vehicle type. The method of successive averages (MSA) offers a practical way of achieving convergence with multi-class assignments.**

Emme/2, supplied by INRO

14 INRO advised that Emme/2 can accommodate the DMRB speed/flow relationships using customizable volume-delay functions. Emme/2 provides flexibility for users to define the

functional forms that capture the performance of traffic flow on the network. These functions are designed by the modeller and input as expressions, restricted only by the requirements of the underlying assignment models in order to ensure the convergence of the solution algorithms towards equilibrium flows.

15 INRO also advised that the standard equilibrium assignment algorithm assumes that all vehicles on a facility, except perhaps transit vehicles, have the same speed. Accounting for differing speeds among light and heavy vehicle classes can be accomplished with asymmetric cost equilibrium assignment. This method uses an averaging scheme, a variant of the method of successive averages where the class flows (e.g. car, light goods vehicle, heavy goods vehicle) are updated at each iteration; auxiliary flows which correspond to all-or-nothing assignments for each vehicle class provide at each iteration the ability to generate improved solutions[1].

16 A meeting was held with Mike Florian of INRO on 6 October 2006 to discuss convergence of multi-class assignments. The following points emerged from that meeting:

- if an assignment is run in the same way, with the same input network and demands, to the same level of convergence, the results will always be the same; however, if the network is changed, because the multi-class assignment process does not necessarily produce a unique solution and because the with-scheme assignment uses a different network to the initial model, it is possible that the link flows by vehicle type may not be as sensible or plausible or realistic as those in the initial model, although this potential non-uniqueness is very rarely a problem in practice;

- under certain conditions, an objective function for Frank-Wolfe user equilibrium assignment may exist but may be complex, and the method of successive averages offers a more practical alternative;

- the method of successive averages works well in practice, even for ill-behaved problems, and especially if monotonic cost/flow relationships are used;

- in principle, multi-user class assignments which use separate speed/flow relationships for light and heavy vehicles and the method of successive averages for convergence can be programmed in Emme/2 (using the macro language), although speed/flow relationships which differ by vehicle type have not yet been used[2];

- the use of separate speed/flow curves for different user classes would not add to the theoretical problems of the convergence of multi-class assignments since similar models have been solved successfully; and

- validation by vehicle type is important for confidence to be gained that multi-user class assignment is likely to behave sensibly.

17 In subsequent correspondence, Mike also suggested that capping the speed at a minimum level, as his reading of the DMRB speed/flow relationships suggested, may cause difficulties for convergence.

[1] See Jia Hao Wu, Michael Florian and Shuguang He, An Algorithm for Multi-Class Network Equilibrium Problem in PCE of Trucks: Application to the SCAG Travel Demand Model, January 2005, appended to this Report.

[2] While this was the position at the time of the meeting, INRO have since advised that they have now successfully implemented a multi-user class assignment with speed/flow relationships which differ by vehicle type.

18 A related issue is that the passenger car equivalent of heavy goods vehicles has been shown to vary considerably by total traffic volume[3].

19 **It therefore seems clear that Emme/2 <u>can</u> meet the DMRB specification for separate speed/flow relationships by vehicle type. It also appears that the non-uniqueness of multi-user class assignments is not likely to be a problem in practice, and the use of speed/flow relationships which differ by vehicle type would not add to the problems already perceived by INRO in converging multi-user class assignments. Also, INRO are concerned about speeds being capped when assigned flows reach and exceed capacity.**

MEPLAN, supplied by WSP

20 WSP advised that the current version of MEPLAN implements its link-based speed/flow relationships in a form that does not entirely match that in DMRB. MEPLAN calculates speed based on the capacity of the specific link and on the number of passenger car units (pcus) loaded in the current iteration, rather than having an explicit percentage of heavy goods vehicles. In this way, account is taken of different pcu factors for a range of vehicle types for any given type of link.

21 There is a range of non-linear functional forms available within MEPLAN for which the modeller can input specified parameter values that enable them to fit the shape of the speed/flow curve for a specific type of road link. In this way, curves that are similar in shape to the piecewise linear forms of the DMRB relationships can be created within MEPLAN, based on the set of link supply descriptors for that link type. There is no limit on the number of link types that can be defined in MEPLAN so that, provided that the required link supply descriptor characteristics are available for each link, the DMRB formulae could be closely matched in shape for each road link considered in the network.

22 The congested speed that is calculated in MEPLAN is that of a representative vehicle. Each other vehicle type derives its congested speed by applying a multiplicative factor (a ratio specific to each combination of type of link and type of vehicle) to this congested speed of the representative vehicle. While this does enable the congested speed of each different type of vehicle to be represented separately on each type of link, it is different from the approach in DMRB.

23 WSP advised that they are currently considering extending the functionality of MEPLAN so that this multiplicative factor is replaced by a functional form that would allow more flexibility in adjusting this speed ratio to take account of the degree of congestion. In congested conditions the ratio would tend, say, to 1, whereas in free flow conditions it would represent the ratio of the speed limit (or typical free flow speed) of the actual vehicle type to that of the representative vehicle type on that specific type of link.

24 The MEPLAN speed/flow functions can also take account of the speed likely to be achieved when traffic reaches the capacity limit - extra traffic demand above this capacity is deflected in a manner that adjusts the value of time, rather than implying speeds that are unrealistically low when compared against those observed in practice on average on highly congested roads. WSP stated that they would not wish to match the shape of the DMRB speed/flow relationships close to congestion as they do not believe them to be particularly realistic for use in equilibrium models.

25 The MEPLAN assignment typically is iterated to either a stochastic or deterministic user equilibrium solution across all passenger and freight user classes, using the method of

[3] See footnote 1.

successive averages (MSA) rather than using a Frank-Wolfe method to determine the step size to adopt in individual iterations.

26 **It therefore seems clear that MEPLAN cannot meet the Department's specification exactly. It appears that MEPLAN can produce different speeds for different types of vehicle, albeit from an assignment of all vehicles together, although the nuances of the DMRB relationships do not appear to be fully reflected because the ratio of heavy vehicle to light vehicle speed varies only by link type rather than being dependent on all the link characteristics reflected in the DMRB relationships. Also, WSP are concerned about the use of speed cut-offs and about the shape of the DMRB relationships at flows above capacity**.

OMNITRANS, supplied in the UK by Minnerva

27 Minnerva advised that the form of the DMRB speed/flow relationships can be accommodated in OmniTRANS (OtTraffic). Within the calculation of speed or delay, several options are available to include the non-flow variables. If treated as discrete variables, this information can be added as 'type' information for all links, and 'families' of curves can be derived to accommodate the non-flow variables included in the predictive equations, such as the frequency of major intersections and the geometric characteristics. If, however, the data were not amenable to categorization, the non-flow variables could be stored as separate link attribute data and user-defined formulae applied which referenced these variables.

28 Minnerva also advised that convergence may be addressed relatively easily in OmniTRANS using volume-averaging methods which can take into account changes in costs on both links and at junctions. It is also reasonably simple to apply an approximation to an equilibrium solution whereby the 'lambda' value per iteration is constrained to be constant across all vehicle classes (lambda in this case is the proportion of the current iteration load that results in minimum network cost in each iteration of the equilibrium assignment).

29 Minnerva also advised that, although the application of a more rigorous equilibrium solution, which allows the 'lambda' to vary by user class per iteration, would require a more advanced procedure, this could be specified. However, while this approach would result in an equilibrium situation that is theoretically more correct, it is rarely used in practice for pragmatic reasons, especially when several user classes are required.

30 A meeting was held with Martin Bach and Miles Logie of Minnerva on 10 November 2006 to discuss convergence of multi-class assignments. The following points emerged from that meeting:

- multi-user class assignments do not necessarily yield unique flows by user class, but the use of generalised cost coefficients by user class will make non-uniqueness less of a practical problem, as would the use of speed/flow relationships which vary by user class (or any other distinction between the bases on which paths for the different user classes are built);

- Minnerva's understanding is that multi-user class assignments are not usual in The Netherlands – heavy goods vehicles are normally assigned first and then treated as a pre-load – and there is therefore little experience (of which they are aware) of multi-class assignments using OmniTRANS;

- the use of Frank-Wolfe user equilibrium assignment for the convergence of multi-class assignment is straightforward and its application using OmniTRANS is currently under development but not yet fully proven in a real application;

- the more usual approach to achieving convergence of multi-class assignment is the method of successive averages (MSA);

- a multi-user class test bed showed a significant degree of volatility in convergence using MSA and that this volatility would be reduced with different speed/flow relationships for each user class, although significant volatility remained in the example presented; this test bed was not presented, however, to make any particular statement about the efficiency of MSA but simply to demonstrate that multi-user class assignment could be implemented in OmniTRANS and show that, in this instance, convergence efficiency might be improved with the use of separate speed/flow relationships by user class;

- OmniTRANS allows pcu factors to vary by link type and flow level, although Minnerva are not aware of any empirical evidence for such variations; and

- normal practice is to assign matrices of pcus because junction capacities are modelled in these units, but it would be possible to conduct the speed change procedure (in the assignment process) in vehicles if the speed/flow relationships were constructed in these units (as are the DMRB relationships).

31 **It therefore seems clear that OmniTRANS (OtTraffic) can meet the DMRB specification for separate speed/flow relationships by vehicle type. It also appears that the non-uniqueness of multi-user class assignments is not likely to be a problem in practice, and the use of speed/flow relationships which differ by vehicle type would not impact adversely on, and might assist with, the convergence of multi-user class assignments.**

SATURN, supplied by WS Atkins and ITS Leeds

32 Dirck Van Vliet identified three ways of modelling different speeds for light and heavy vehicles, as follows[4]:

- heavy vehicles require a fixed extra time per kilometre relative to light vehicles over all flow conditions;

- heavy vehicles travel at a maximum speed until, with increasing flow, the speed of light vehicles drops below that value, after which both heavy and light vehicles travel at the same speed; and

- heavy and light vehicles have separate speed/flow curves with a maximum difference at zero flow but gradually converging until, at capacity (or, possibly, at a lower flow convergence point), both travel at the same speed.

33 The first option is currently available in SATURN 10.6 using the CLICKS facility and appears to be similar to the approach adopted in MEPLAN. The second method is equivalent to the DMRB relationships for road classes other than rural single carriageways and Dirck advises that this method will be included in release 10.7. The third method is similar to the DMRB relationship for rural single carriageways and Dirck advises that this may or may not be incorporated in the software depending on demand.

34 It should be noted that, in any case, SATURN uses a power flow/delay curve to approximate the two-stage linear DMRB speed/flow relationships.

[4] See On Variable Link Speeds by User Class within SATURN, by Dirck Van Vliet, dated 31 July 2006, appended to this Report.

35 It appears that one of the reasons for the choice of the CLICKS option within SATURN was that a fixed time (i.e., cost) penalty per user class per link fits very easily into the algorithm currently used to solve for multi-user class Wardrop equilibrium assignment and does not, for example, introduce any extra potential problems of non-convergence or multiple equilibria (which, in theory at least, already exist through the simulation model).

36 Dirck also appears to suggest that the differences between light and heavy vehicle speeds could be calculated from the DMRB speed/flow relationships at the outset of an assignment (that is, a series of loads and speed changes run through to convergence), but then kept fixed throughout any one assignment process. The resulting assigned flows could then be used to recalculate the difference between light and heavy vehicle speeds, as derived from the DMRB speed/flow relationships, following which the assignment would be repeated. By these means, account could be taken of the different DMRB speed/flow relationships by vehicle type.

37 A meeting was held with Dirck on 6 November 2006 to discuss the use of speed/flow relationships which vary by vehicle type. The following points emerged from that meeting (and some pre-meeting exchange of emails):

- although multi-user class assignments do not necessarily yield unique flows by user class, single user class assignments may not always be unique either if non-separable cost/flow curves are used, but, in practice, this issue of non-uniqueness is rarely problematic;

- the convergence of multi-user class assignments is best achieved by means of Frank-Wolfe user equilibrium methods, which, with some simplifications, are practical to implement with multi-user class assignments;

- SATURN could be amended to model the DMRB speed/flow relationships, albeit as flow/delay curves;

- the use of speed/flow relationships which vary by user class is unlikely to make the rare problems of non-uniqueness any worse or to affect materially the convergence of multi-user class assignments;

- SATURN currently fixes pcu factors by user class (although Dirck subsequently advised that he has algorithms in mind which could allow pcu factors to vary by link type or individual link); and

- normal practice is to assign matrices of pcus because junction capacities are modelled in these units, but it would be possible to conduct the speed change procedure in vehicles if the speed/flow relationships were constructed in these units (as are the DMRB relationships).

38 In subsequent correspondence, Dirck argued that speeds should continue to decline in some way when assigned flows exceed capacity, rather than being cut-off at some point.

39 **It therefore seems clear that SATURN <u>cannot</u> meet the Department's specification <u>exactly</u> at the time of writing. However, it appears that SATURN 10.7 <u>will</u> allow modellers to use DMRB relationships by vehicle type, for all but rural single carriageway roads. It also appears that the non-uniqueness of multi-user class assignments is not likely to be a problem in practice, and the use of speed/flow relationships which differ by vehicle type would not impact adversely on the convergence of multi-user class assignments. Also, concern was expressed about speeds being capped when assigned flows exceed capacity.**

TRANSCAD, supplied by Caliper Corporation

40 Caliper advised that there would be relatively little difficulty implementing the DMRB speed/flow relationships in TransCAD. A standard component of TransCAD is a multi-class, multi-modal assignment model for road traffic, which permits as many different types of vehicles as the user wishes, and also has provision for user-written volume/delay functions. As an integrated GIS and travel demand forecasting package, TransCAD also supports an unlimited number of link attributes so the various geographic characteristics specified in the DMRB speed/flow relationships can be easily accommodated. Indeed, some of them can be calculated directly from GIS files.

41 Caliper went on to advise that there is a serious problem of non-uniqueness of user equilibrium multi-class assignments: only the total link flows are unique, preferably in terms of total passenger car equivalents or some other combined measure, which means that the vehicle class breakdown by size by link can be rather arbitrary, although, if the class link flows match counts closely, this may be less of a worry. Caliper also advised that a second potential flaw would be the use of minimum speeds or speed cut-offs. These can be shown to break or greatly impair the convergence of equilibrium traffic assignment algorithms. Cut-offs also can invalidate measures such as the ‘relative gap’ that are used to assess convergence.

42 Caliper also advised that there are no difficulties with the use of Frank-Wolfe algorithms for achieving user equlibrium multi-class assignments, if the formulation of the objective function used is relatively simple.

43 Caliper also advised that the use of speed/flow relationships which vary by user class may not give rise to problems with convergence. For example, if lorries always travel at 80% of the speed as private cars there would be no difficulty, and, if the speed of cars varies with the percentage of lorries on a link, the problem may still be solvable. However, if the speeds are interdependent in a more complex way, there might be a problem.

44 **It therefore seems clear that TRANSCAD <u>can</u> meet the DMRB specification for separate speed/flow relationships by vehicle type. It also seems clear that Caliper have concerns about the practical implications of the non-uniqueness of multi-class assignments which do not appear to be shared by any other software supplier contacted in this research. Caliper also have concerns about the use of speed cut-offs.**

Visual-tm, supplied by Peter Davidson Consultancy

45 Peter Davidson Consultancy (PDC) advised that Visual-tm allows modellers to specify separate speed/flow relationships for each vehicle type. The modeller specifies the set of user classes being assigned (a user class being associated with one vehicle type). Each vehicle type assignment is converted to pcus with pcu factors defined by the modeller, summed to give total link flows, and then input to the speed/flow curves to give speeds for each vehicle type for each link. With speed/flow relationships by vehicle type, a Frank-Wolf convergence routine has not yet been implemented in Visual-tm, but user equilibrium can be achieved using the method of successive averages.

46 PDC also advised that, in Visual-tm, the modeller can apply speed/flow relationships as either two straight lines or equations. The equations can use link characteristics, which if they include the necessary variables (hilliness bendiness, etc.), can be used to compute the DMRB speed/ flow relationships.

47 In PDC's view, the ease or difficulty of achieving convergence depends more upon the individual circumstances than upon whether speed/flow relationships are specified separately by vehicle type.

48 **It therefore seems clear that Visual-tm <u>can</u> meet the DMRB specification for separate speed/flow relationships by vehicle type**.

VISUM, supplied by PTV

49 PTV advised that VISUM 9.6, which is scheduled for release early in 2007, will offer modellers the facility to model relationships of the DMRB type. According to the manual supplied by PTV, modellers may define their own functional forms for volume-delay functions which:

- include other link attributes in the calculation;
- calculate pcus in non-standard ways; and
- define separate volume-delay functions for different transport systems (that is, for light and heavy vehicles, for example).

50 **It therefore seems clear that, VISUM 9.6 <u>will</u> be able to represent the general two-regime linear form of the DMRB speed/flow relationship separately by vehicle type.**

SUMMARY OF THE FINDINGS FROM THE SOFTWARE SUPPLIERS

Speed/Flow Relationships by Vehicle Type

51 The following packages appear able to model separate speed/flow relationships for light and heavy vehicles, of the two-regime linear form specified in DMRB 12.2.1:

- CONTRAM;
- CUBE VOYAGER;
- Emme/2;
- OmniTRANS (OtTraffic);
- TRANSCAD; and
- Visual-tm.

52 The positions with the other packages appears to be as follows:

- MEPLAN can approximate the DMRB requirements;

- SATURN can represent the two-regime linear form of the DMRB speed/flow relationships by means of a power function flow/delay curve and can currently approximate the DMRB requirements using the CLICKS option, and there are plans to amend SATURN to model separate speed/flow relationships by vehicle type more closely (for all except rural single carriageways), albeit as power function flow/delay curves;

- VISUM 9.6, due for release shortly, will be able to represent the DMRB relationships.

Speed Cut-Offs

53 Four suppliers (INRO (Emme/2), WSP (MEPLAN), Dirck Van Vliet (SATURN) and Caliper (TRANSCAD)) interpreted the DMRB relationships as implying speed cut-offs at flows higher than capacity, and pointed out that the use of speed cut-offs could adversely affect convergence. (Note that the software suppliers were sent Appendix E from DMRB 12.2.1

and it is presumed that they based their interpretations on the advice given in paragraphs E1.8, E2.10, E3.8, E4.9, E5.6, and E6.11. Paragraph 4.4.15 in that appendix makes it clear that the Department does **not** recommend the use of speed cut-offs.)

Convergence

54 The position with regards to convergence to user equilibrium appears to be:

- given the heuristic nature of CONTRAM's assignment process, the use of multi-user classes and separate speed/flow relationships for different classes presents no special problems for convergence;

- INRO take the view that the form of the objective function required for a Frank-Wolfe user equilibrium convergence procedure for multi-class assignment is too complex and instead use the method of successive averages in Emme/2;

- the method of successive averages is recognised by Citilabs (CUBE VOYAGER), WSP (MEPLAN), Minnerva (OmniTRANS), and Peter Davidson Consultancy (Visual-tm) as being a practical method of converging multi-class assignments; and

- Caliper (TRANSCAD), Dirck Van Vliet (SATURN) and Minnerva (OmniTRANS) see few problems with using, and indeed the first two appear to prefer to use, Frank-Wolfe user equilibrium methods for convergence in multi-class assignments, albeit with some simplifications to the objective function.

55 All but one of the suppliers (of those that expressed a view) considered that the use of speed/flow relationships by vehicle class would not make convergence more difficult to achieve.

Non-Uniqueness of Multi-Class Assignments

56 From the discussions with the suppliers, my interpretation of the non-uniqueness of multi-class assignments is as follows. Assuming that we have established a well-validated multi-class assignment model, if we run the assignment in the same way, with the same input network and demands, to the same level of convergence, the results will always be the same. However, if the network is changed, because the multi-class assignment process does not produce a unique solution, and because the with-scheme assignment uses a different network to the initial model, is it possible that the link flows by vehicle type may not be as sensible or plausible or realistic as those in the initial model[5].

57 Caliper drew particular attention to the dangers of non-uniqueness of multi-class assignments. However, while, in principle, the potential for some degree of arbitrariness in the flows by vehicle class was recognised by other suppliers (INRO, Dirck Van Vliet, and Minnerva) none thought that the issue is of material significance. Indeed, Dirck argued that non-uniqueness could, in theory, arise with a single user class assignment, although this was not a common problem encountered in practice.

Passenger Car Equivalents (or Units) or Vehicles

58 The DMRB relationships are couched in terms of **vehicle** flows, while junction modelling, in SATURN or OmniTRANS, for example, generally requires flows to be in passenger car equivalents. SATURN, MEPLAN and OmniTRANS can accommodate variable pcu

[5] I understand that it is also possible for a base year validated model to yield forecast year assignments (that is, with changed demands) which appear to be less sensible even when there are no network changes between base and forecast year.

factors and could also carry out the speed change procedure (in the assignment) in vehicles as per the DMRB relationships.

59 However, while the passenger car equivalent of heavy goods vehicles has been shown in the USA to vary considerably by total traffic volume, there currently appears to be no substantial body of evidence on this topic relating to the UK.

Summary of Findings

60 The following table summarises my main findings, package by package.

Package (Supplier)	**Is it possible to use speed/flow relationships which differ by vehicle type?**	**Is convergence to user equilibrium likely to be affected with speed/flow relationships which differ by vehicle type?**	**Is the potential non-uniqueness of multi-user class assignments likely to be a problem in practice?**
CONTRAM (Mott MacDonald/ TRL)	Yes.	No.	Not asked.
CUBE VOYAGER (Citilabs)	Yes.	It may be possible to use Frank-Wolfe. MSA may be required.	Not asked.
Emme/2 (INRO)	Yes.	MSA method is required. Frank-Wolfe is not practical.	No.
MEPLAN (WSP)	An approximation is possible.	No.	No.
OmniTRANS (Minnerva)	Yes.	MSA is a practical approach. Frank-Wolfe can be used (with some simplifications).	No.
SATURN (Dirck Van Vliet/ WS Atkins/ ITS Leeds)	Yes, in Version 10.7 (forthcoming) for all road types except rural single carriageways. An approximation is currently possible.	Frank-Wolfe is used (with some simplifications).	No.
TRANSCAD (Caliper)	Yes.	Frank-Wolfe is used (with some simplifications).	It could be a serious problem.
Visual-tm (Peter Davidson Consultancy)	Yes.	MSA is used.	Not asked.
VISUM (PTV)	Yes, in Version 9.6 (forthcoming).	Not asked.	Not asked.

CONCLUSIONS

61 My overall conclusion from this research is that, were the Department to require the use of separate speed/flow relationships for light and heavy vehicles of the DMRB form, modellers could accommodate such a requirement.

ADDENDUM

62 In the course of the investigations, certain questions occurred to ITEA about current practice in the use of speed/flow relationships in traffic modelling in the UK. I was therefore asked to conduct a short **informal** survey to gain some insight into the issues of concern.

63 ITEA's questions were:

- Is it general practice to use a single speed/flow curve on any one link for all user classes?
- If so, what speed does the single curve generally represent – that of light vehicles or an average of all vehicles?
- Is it usual to define speed/flow relationships in terms of pcus or vehicles?
- In defining the speed/flow curves for use in a traffic model, is it common practice to measure and use geometric data specific to each link or is some more generalized approach usually adopted?

64 My approach was to ask modellers that I knew and who I thought would have something to contribute on this subject. Given this rather informal method of sampling, the views of the 17 respondents reported below should not therefore be taken as necessarily representative of the profession as a whole.

65 The answers were as follows.

66 **Is it general practice to use a single speed/flow curve on any one link for all user classes?**

- Yes (12 responses).
- Yes in SATURN models but curves which differ by user class are often used in CONTRAM models (1 response).
- Speeds for certain user classes, e.g. HGVs, are sometimes capped (4 responses).

67 **If so, what speed does the single curve generally represent – that of light vehicles or an average of all vehicles?**

- All vehicles or the average vehicle (7 responses).
- Light vehicles (4 responses without comment or explanation and 1 response with the additional advice that light vehicle speeds are used in journey time validation).
- Light vehicles, with multiplicative scaling factors, differentiated by vehicle type, applied to convert the light vehicle speed to other vehicle speeds (1 response).

68 **Is it usual to define speed/flow relationships in terms of pcus or vehicles?**

- Pcus (13 responses).
- Pcus in SATURN and vehicles in CONTRAM (1 response).

69 **In defining the speed/flow curves for use in a traffic model, is it common practice to measure and use geometric data specific to each link or is some more generalized approach usually adopted?**

- A generalised approach is commonly used with curves being defined for types of link (5 responses).

- A generalised approach is commonly used except where bends and hills are significant (1 response).

- A generalised approach is commonly used except where hills are significant (2 responses).
- A generalised approach is commonly used, especially for large-scale models (1 response).

- A generalised approach is commonly used, but subject to review and possible amendment during calibration (1 response).

- A generalised approach is used for large-scale models but geometric data are measured and used for detailed models (1 response).

- Curves are defined for link types, but these are then used with the capacity specific to each link, where available (1 response).

- Geometric data are measured and used for each link (1 response in the context of a specific model).

- It is not common practice to measure and use geometric data specific to each link (1 response).

- SATURN does not require the use of specific geometric information (1 response).

70 The following other points may be of interest to ITEA:

- SATURN was the assignment software most mentioned;

- COBA was the only source of speed/flow relationships mentioned;

- SATURN uses a flow/delay relationship defined by a power function to approximate the COBA relationships;

- a network-wide percentage of HGVs is generally used in the light vehicle rural road speed/flow relationships;

- the Highways Agency do not encourage or generally endorse amendments to the COBA speed/flow relationships unless good data are available to justify such amendments;

- standard pcu factors for HGVs are generally used (around 2.0);

- while link-specific pcu factors might be useful, SATURN does not currently provide the facility to allow this (1 respondent);

- multi-user class assignments may not be justified if there are doubts about the quality of the HGV trip matrix (1 respondent);

- multi-user class assignments are much less common than single user class assignments, with HGVs often being preloaded on all-or-nothing routes (1 respondent).

dc

An Algorithm for Multi-Class Network Equilibrium Problem in PCE of Trucks : Application to the SCAG Travel Demand Model[1]

Jia Hao Wu[2,3] Michael Florian[2,3] Shuguang He[2]

INRO Solutions Inc[2]
Centre for Research on Transportation, Université de Montréal[3]

January 2005 (Revision of August 2001 version)

[1] This project is supported in part by NSERC individual operation grants OGP0157735 and OGP0007406.

Abstract

The analysis of mixed traffic, which includes both private and trucks has received more attention recently due to the rapid development of urban truck traffic in many cities and regions of the world. In this paper, we consider a multi-class network equilibrium model where several classes of traffic, with their own travel times, interact on the links of the network. The volume/delay functions depend on the mix of trucks and cars and other factors like the slope of the links. Hence, the resulting cost functions are nonlinear, non-smooth and asymmetric. The problem is formulated as a nonlinear, nonsmooth, variational inequality model. A linear approximation type algorithm, which uses step size based on the method of successive averages, is used to solve the problem. Numerical results are reported for a large scale problem (6 classes of traffic, 3217 zones and 99867 links).

Key words: network equilibrium, multi-class traffic assignment, variational inequality problem.

Résumé

L'analyse du trafic mixte, qui inclut les voitures privées et les camions, a reçu plus d'attention récemment à cause de l'augmentation des déplacements urbains de camions dans plusieurs villes du monde. Dans cet article, nous considérons un modèle d'équilibre multi-classe où plusieurs classes de trafic, avec leur propre temps de déplacement, utilisent la même infrastructure routière. Les fonctions de volume/délai dépendent de la composition du trafic et de la pente du lien. En conséquence, les fonctions de coûts sont nonlinéaires, non-différentiables et asymétriques. Ce problème formulé comme une inégalité variationnelle est résolu par une méthode d'approximation linéaire couplée avec des pas calculés par la méthode de moyennes successives. Des exemples numériques sont présentés pour un problème de grande taille (6 classes de trafic, 3217 zones et 99867 liens.

Mots clés : équilibre de réseau, affectation multi-classe, problème d'inégalité variationnelle.

1. Introduction

The modeling of urban freight movements has received more attention recently due to the need of representing the increased truck traffic on urban road networks. The models used for transportation planning of urban areas are enhanced with the flows of trucks in the period of times studied. The traffic flow characteristics of mixed traffic of cars and trucks are quite complex since the delay depends on the proportion of the different vehicles on a given link, the grade of the link and the behavior of the drivers. The model presented in this paper is motivated by a detailed study carried out in the area covered by the jurisdiction of the Southern California Association of Governments (SCAG) by a consulting firm (Meyer, Mohaddes, [6]). The model that is developed in this contribution is a multi-class network equilibrium model with asymmetric costs, which is formulated as a variational inequality. It is solved by a linear approximation algorithm, which uses step sizes computed with the method of successive averages. The algorithm was implemented in a widely available transportation planning software package and was applied to the network used by SCAG for transportation planning purposes. This is a very large network consisting of 3,217 centroids, 25,430 nodes and 99,687 links.

The most difficult part of the model is that, as mentioned above the link travel times depend on the mix of trucks and cars on the link and on the slope of the link as well. The resulting link travel time functions are nonlinear and asymmetric. The volumes of trucks of different classes (that is size) are converted into passenger car equivalents, or PCE's (we will use this notation even though the reference to PCU's, passenger car units is common as well). The conversion factors depend on the percentage of each class of traffic on the link since the interaction between the various vehicle types is complex and quite variable.

An example of PCE conversions used may be found in the report of Meyer, Mohaddes [6] where the total PCE's on a link are a nonlinear function of three classes of trucks and three classes of private cars. The difficulty as well as the novelty of this problem arises from the fact that such a network equilibrium model was not previously studied in the literature. A new mathematical formulation as well as a solution algorithm must be developed for this model.

In this paper we focus on the model formulation, the solution algorithm and the presentation of the computational results obtained with the SCAG network. The only theoretical contributions in the literature that are relevant to the algorithm used in this model are asymmetric cost single class models. Marcotte and Zhu [5] and Magnanti and Perakis [4] use an LP based operator and a projection operator respectively and prove the convergence of these algorithms under certain conditions. We extend these types of algorithms for this multi-class of network equilibrium problem. The theoretical study of the convergence of the algorithm that is presented in the following is beyond the scope of this paper.

This paper is organized as follows. In Section 2, the mathematical model for the problem is defined. In Section 3, a solution method is developed for the mathematical model, while Section 4 is devoted into the application of the method in a real network. Section 5 concludes the paper.

2. Model Formulation

In this section, the notation used is introduced in order to state the mathematical formulation for the problem.

2.1 Notations

In this paper, the following notations are used. The links of the road network are designated by $a \in A$, where A is the set of links. The demand for travel by user class $m \in M$ for the origin-destination pair (i, j) is denoted as T_{ij}^m where M is the set of classes. This demands may use paths $R_{ij}^m \in R$ where R is the set of all routes $R = \bigcup_{(i,j),m} R_{ij}^m$. Table 1 summarizes the definition of sets and indices used in this paper.

Table 1 Sets and indices used

A	: link set of base network in period t	a	: link index of base network
W	: total OD pairs	w	: OD pair index
R_{ij}^m	: set of routes for pair (ij), class m	r	: path index
M	: class set	m	: class index

Given the sets and indices above, the following variables and given data are defined:

The variables include:

h_r^m : The path flow of class m on the route r;

f_a^m : The total link flow of class m on link a;

C_r^m : The path travel time of class m on route r;

v_a : The total link flow in PCE (passenger car equivalence) that include flows of all classes on link a (which will be discussed in next section);

v_a^m : Link flow in PCE of class m on link a, $\left(v_a = \sum_m v_a^m \right)$;

$c_a(v_a)$: The travel time on link a for total link flow v_a in PCE.

The given data consists of :

T_{ij}^m : The travel demand from origin i to destination j of class m.

2.2 Total link flow in PCE

The total link flow in PCE, that is, $v_a = f_a\left(f_a^m \middle| \forall m \in M \right)$ can expressed as a very general function of link flows (in vehicles) of all classes f_a^m, $\forall m \in M$, and the parameters associated with link a. It is not just simply be a linear combination of the link flows of all classes such as $v_a = \sum_m \beta^m f_a^m$ where β^m is a conversion factor into PCE from the total link flow of class m. In particular, the following nonlinear function is considered:

$$v_a^m = f\left(f_a^m, p_a^m, vc_a, k_a, adj_a^m, l_a, g_a \middle| \forall m \in M \right) \tag{1}$$

where

p_a^m is the percentage of the link flow of class m on link a

$$p_a^m = f_a^m / \sum_{m' \in M} f_a^{m'}, \tag{2}$$

vc_a is the link flow in PCE, v_a, over the capacity of link a, k_a, given as

$$vc_a = v_a / k_a, \tag{3}$$

k_a is the link capacity in PCE on link a

$adj_a^m(vc_a)$ is an adjustment factor of class m on link a based on vc_a,

l_a is the length on link a,

g_a is the grade on link a.

This type of nonlinear function has been calibrated and validated with observed data and can be formulated in the form of a look-up table in practice. Consider the following example with v_a expressed as

$$v_a = \sum_{m \in M} PCE_a^m \left(p_a^m, k_a, adj_a^m \left(vc_a \right), l_a, g_a \right) \times f_a^m$$

where $PCE_a^m \left(p_a^m, k_a, adj_a^m \left(vc_a \right), l_a, g_a \right)$ can be presented in the comprehensive look-up table (See Section 4 for an example).

2.3 Mathematical model

The feasible region Ω of the problem is defined as follows:

$$\sum_r h_r^m = T_{ij}^m, \quad \forall \left(ij \right), m \in M \tag{4}$$

$$h_r^m \geq 0, \forall m \in M, r \in R_{ij}^m \tag{5}$$

where

$$f_a^m = \sum_{r \in R} h_r^m \delta_{ar}, \quad \forall a \in A, m \in M \tag{6}$$

$$C_r^m = \sum_{a \in A} c_a^m \left(v_a \right) \delta_{ar}, \quad \forall m \in M, r \in R_{ij}^m \tag{7}$$

where (4) are the equations of conservation of flow, (5) is the nonnegativity of the path flows, δ_{ar} is 1 if link a is on route r and is zero otherwise and (7) is the definition of path travel time of class m on route r, C_r^m.

The multi-class network equilibrium problem can be formulated as a variational inequality problem. Find $h^* \in \Omega$ such that

$$\sum_{w \in W} \sum_{m \in M} \sum_{r \in R_w} C_r^m \left(h^* \right) \left(h_r - h_r^* \right) \geq 0, \forall h \in \Omega. \tag{8}$$

It is well known that while the total link flow v in PCE is unique if $\left(c_a^m \left(v \right) \right)$ is strictly monotone, the composed volumes $\left(f^m, m \in M \right)$ may not be unique with $c_a^m \left(v \right) = c_a \left(v \right)$ for all m. However the strict monotonicity conditions are difficult to verify for this model. It is clear that the solution of the problem satisfies the following equilibrium conditions:

$$C_r^m \begin{cases} = u_w^m & \text{if} \quad h_r^m > 0 \\ \geq u_w^m & \text{if} \quad h_r^m = 0 \end{cases}, \ \forall m \in M, r \in R_w \tag{9}$$

where $u_w^m = \min_{r \in R_w} \{C_r^m\}$ is the minimum travel time for pair w of class m, which are the well known Wardrop [8] user's optimal conditions. The derivation of the variational inequality formulation (8) from Wardrop's user optimal principle (9) is well known and may be referenced in the seminal work of Smith [7]. We may reference also the survey chapter of Florian and Hearn [1].

3. Solution Algorithm

In the literature, there are many recursive averaging schemes. Some were studied by Marcotte and Zhu [5] and Magnanti and Perakis [4]. These can be used to solve the variational inequality problem with only a single class and a mapping of either LP-based operator or projection operator. The evaluation of the projection operator is equivalent to a minimization of a quadratic convex optimization problem, while the LP-based (linear approximation) operator can be used for multi-classes of demand, which can be solved by solving a shortest path problem with a network loading (a decomposable linear program problem by class) instead of the quadratic optimization problem. We choose the LP-based operator for the development of the solution algorithm, which is given as follows:

Multi-class mixed flow network assignment algorithm

Step 0 : Initialization. Start with $l = 0,\ f_a^{m,l} = 0$.

Step 1: Compute percentage of link flow, v/c ratio and link flow in PCE.

$$p_a^{m,l} = f_a^{m,l} / \sum_{m' \in M} f_a^{m',l}, \forall a \in A, m \in M$$

$$vc_a^l = vc_a^{l+1} / k_a, \forall a \in A$$

$$v_a^{m,l} = PCE_a^m \left(p_a^m, k_a adj_a^m \left(vc_a^l \right), l_a, g_a \right) \times f_a^{m,l}.$$

Step 2: Computation of link cost. $c_a\left(v_a\right), \forall a \in A$.

Step 3: Computation of the shortest path flow problem for each class m.

$$\min \sum_i \sum_k C_k^{l,m} \bar{h}^m$$

$$\text{s.t.} \sum_{r \in R_w^m} \bar{h}_r^m = T_w^m, \ \forall w \in W, m \in M$$

$$\bar{h}_r^m \geq 0, \forall m \in M, r \in R_w^m, w \in W.$$

Step 4: Computation of link flow $\bar{f}_a^m = \sum_{r \in R} \bar{h}_r^m \delta_{ar}, \quad \forall a \in A, m \in M$.

Step 5: Computation of step size (MSA)

$$f_a^{m,l+1} = f_a^{m,l} + \left(\bar{f}_a^{m,l} - f_a^{m,l}\right) / \left(l+1\right), \forall a \in A, m \in M.$$

Step 6: Apply a stopping criterion.

Step 7: $l = l+1$. Go to Step 1.

4. Application

This algorithm was applied to solve the particular multi-class network equilibrium problem in PCE in the network used by SCAG for transportation planning, which was mentioned in Section 1. The algorithm was implemented in the EMME/2 transportation planning software by using its macro language and the judicious use of the multi-class assignment module. Step 2, Step 3 and Step 4 of the algorithm use the first iteration of a symmetric cost multi-class assignment module to compute shortest paths and load the demand for different classes, while Step 1 and Step 5 are implemented with network calculations. Step 6 is coded directly with a macro language procedure.

There are three classes of passenger vehicles and three classes of trucks. The well known BPR volume delay function was used in its original forms:

$$c_a\left(v_a\right) = \beta_a\left(1 + 0.15 \times \left(v_a / k_a\right)^4\right)$$

where β_a and k_a (practical capacity) are both parameters. The six classes of demand of traffic are listed in Table 2.

Table 2 Six classes of demands

Class ID	Descriptions
1	Passenger cars of one person demand that can not access HOV links
2	Passenger car of two person demand that can access whole network
3	Passenger cars of three person or more demand that can access whole network
4	Light-heavy duty trucks that can not access HOV links
5	Medium-heavy duty trucks that can not access HOV links
6	Heavy-heavy duty trucks that can not access HOV links

The trip demands for the six classes of traffic considered in the model are as follows:

Class 1 :	5,119,365	Passenger cars
Class 2 :	2,070,009	Passenger cars
Class 3 :	384,592	Passenger cars
Class 4 :	79,583	Light-Heavy trucks
Class 5 :	51,794	Medium-Heavy trucks
Class 6 :	34,809	Heavy-Heavy trucks

This is probably one of the largest multi-class assignment problems used in the practice of transportation planning.

The three classes of trucks are classified as

1. Light-Heavy : 8,500 to 14,000 GVW
2. Medium-Heavy : 14,000 to 30,000 GVW
3. Heavy-Heavy : over 30,000 GVW

In the computation, the various variables and data information for each link are stored, as shown in Table 3.

Table 3 Partial link attributes

Notation	Description
v_a	Current total equivalent car flow
f_a^m	Successive average flow on link a of class m
g_a	Attribute as the grade for the link

Thus total link flow in PCE v_a on link a is a nonlinear function of three heavy duty trucks and three auto vehicles which is defined as:

$$PCE_a^m\left(p_a^m, k_a, adj_a^m\left(vc_a\right), l_a, g_a\right) = PCE_a^m\left(p_a^m, l_a, g_a\right) \times adj_a^m\left(vc_a, p_a^m\right)$$

where $adj_a^m\left(vc_a, p_a^m\right)$ and $PCE_a^m\left(p_a^m, l_a, g_a\right)$ $\left(PCE_a^m\left(p_a^m, l_a, g_a\right) = 1 \text{ for } m = 1,2,3\right)$ are defined in Table 4 and Table 5. Thus the nonlinear functions of (1) are defined in the form of look-up tables instead of the continuous functions. These functions have been validated with observed data, as reported in Meyer, Mohaddes Associates [6].

Table 4 adj_a^m **values** $m = 4,5,6$

vc_a	p_a^4			p_a^5			p_a^6		
	0-5	5-10	>10	0-5	5-10	>10	0-5	5-10	>10
<0.5	0.60	0.66	0.90	0.66	0.77	0.93	0.90	0.77	0.93
0.5-1.0	0.77	0.89	1.15	0.89	1.01	1.20	1.15	1.01	1.20
1.0-1.5	1.10	1.20	1.30	1.20	1.25	1.34	1.30	1.25	1.34
1.5-2.0	1.00	1.05	1.22	1.05	1.22	1.25	1.22	1.22	1.25
>2.0	1.19	0.66	1.26	1.05	1.24	1.29	1.26	1.24	1.29

Table 5 Heavy Duty Truck PCE values $PCE_a^m\left(p_a^m, l_a, g_a\right)$

p_a^m	l_a	$m=4$			$m=5$			$m=6$		
		g_a								
		0-2	3-4	>4	0-2	3-4	>4	0-2	3-4	>4
0-5	≤1	2.0	4.2	6.4	3.4	6.9	8.8	4.3	8.0	11.3
	>1	2.0	5.5	7.5	5.2	8.4	10.7	6.7	10.5	13.5
5-10	≤1	2.0	3.4	4.8	3.1	5.0	6.4	3.5	5.8	8.8
	>1	2.0	4.2	5.3	3.9	5.9	7.8	4.8	7.8	13.5
>10	≤1	2.0	3.3	4.1	2.8	4.5	6.0	3.2	5.1	8.3
	>1	2.0	3.5	5.0	3.7	5.3	7.5	4.0	7.5	12.5

Two measures of the convergence of the algorithm are provided. The first measure (M1) is the relative difference d^l at iteration l between the volume at iteration l and successive average volume at iteration l :

$$d^l = \sum_a \sum_m \left| \bar{f}_a^{m,l} - f_a^{m,l} \right| \Big/ \left(\sum_a \sum_m \bar{f}_a^{m,l} \right)$$

The second measure (M2) is the relative gap $rgap^l$ at iteration l computed with the flow $\bar{f}_a^{m,l}$ which is the "all-or-nothing" assignment on shortest paths and the last flow $\bar{f}_a^{m,l-1}$ weighted by the current travel time $c_a\left(\bar{v}_a^l\right)$:

$$rgap^l = \sum_a \sum_m c_a\left(\bar{v}_a^l\right)\left(\bar{f}_a^{m,l-1} - \bar{f}_a^{m,l}\right) \Big/ \left(\sum_a \sum_m c_a\left(\bar{v}_a\right) \bar{f}_a^{m,l-1} \right).$$

It is well known that if $d^l \to 0$ or $rgap^l \to 0$ as $l \to \infty$, then one can obtain an equilibrium solution. This solution procedure was implemented in the EMME/2 software package as a macro. The algorithm converges on this network, as shown in Figure 1, where M1 and M2 are the two convergence measures mentioned above.

The computations require about 40 min per iteration on a SUN SPARC ULTRA 5 workstation and about 7 minutes on an IBM Thinkpad T40 based on an Intel Centrino 1.6 Mhz. If the algorithm is used in a travel demand forecasting model, one may need only 4-6 iterations for a reasonable convergence of an inner loop. The assignment results are shown in Figures 2 and 3 where private vehicles (that is, $m = 1, 2, 3$) are not shown.

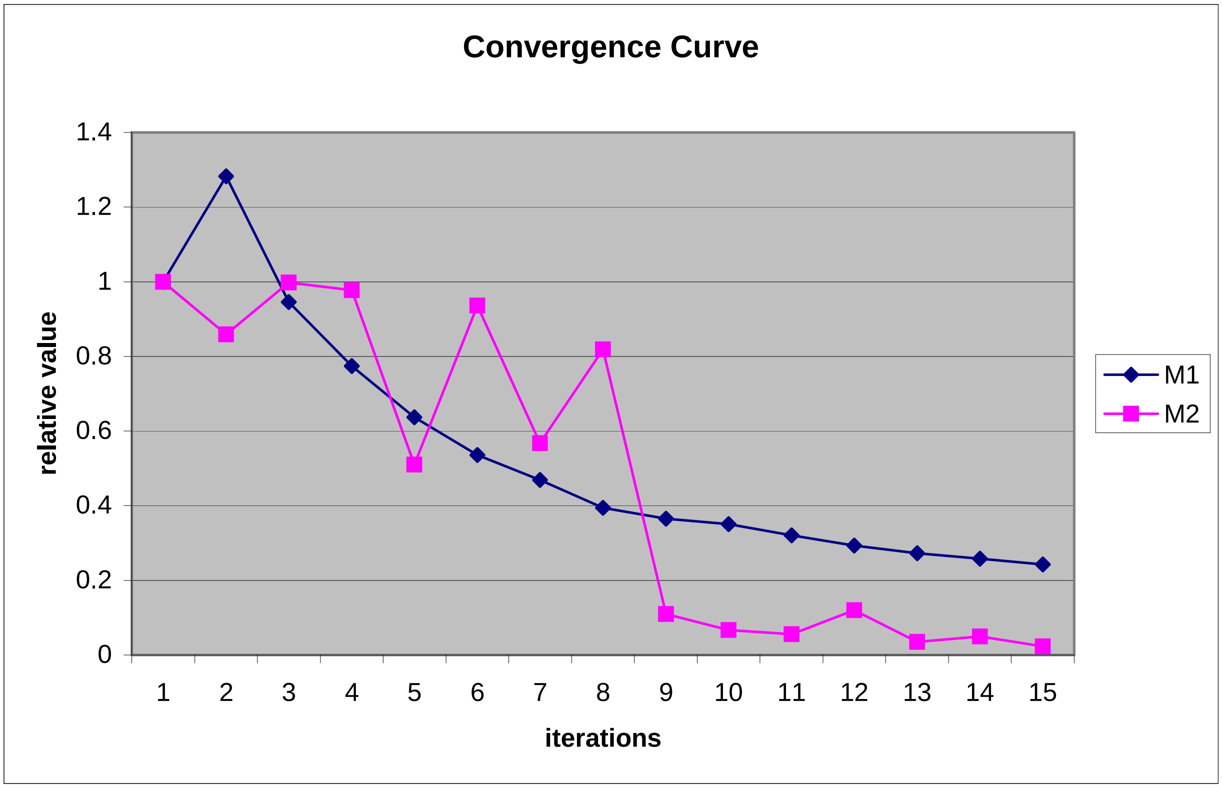

Figure 1 Convergence of the solution procedure

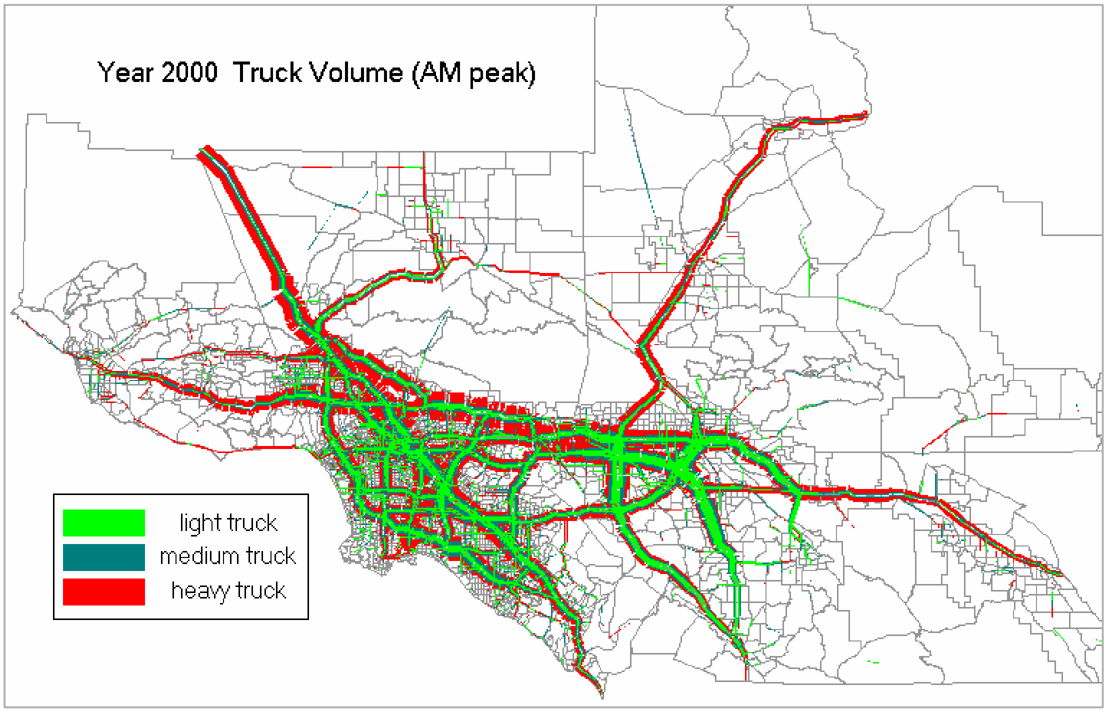

Figure 2 Three classes of trucks – general view

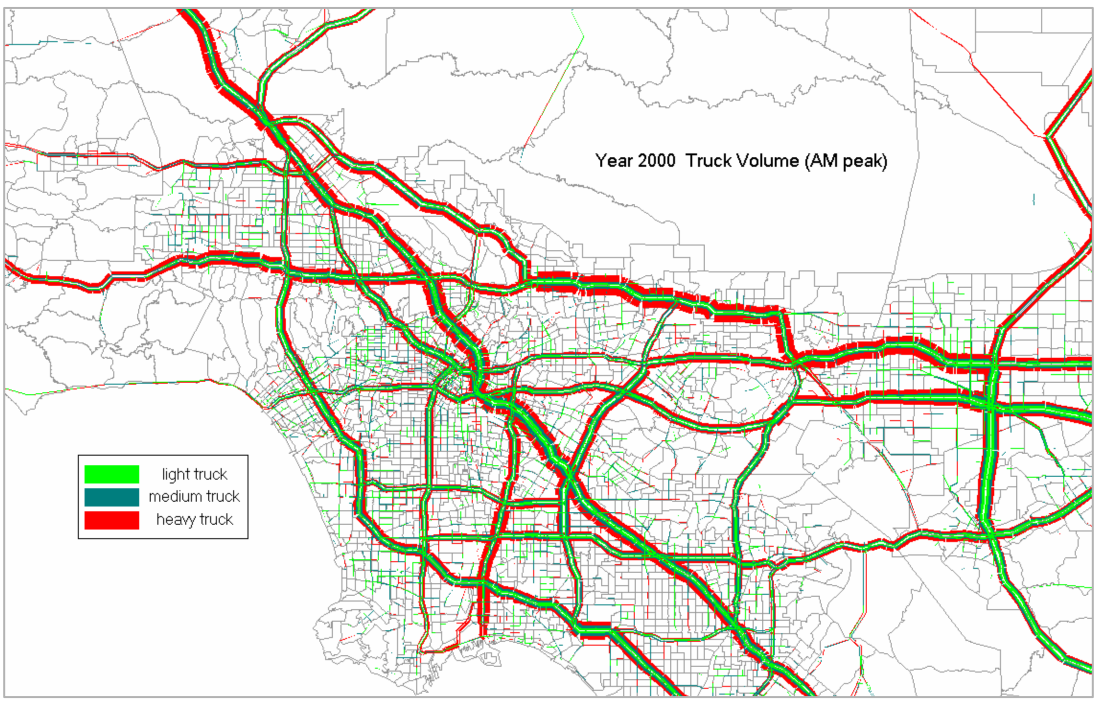

Figure 3 Three classes of trucks vehicles – a window

5. Conclusion

In this paper, the issue of a multi-class mixed car and truck traffic network equilibrium problem was considered. Based on the analysis of the problem, a variational inequality problem was formulated and an adapted successive averaging method was developed for solving a six-class variational inequality problem. Then the linear approximation based method was applied to a real network with good numerical results. This model is actually used in practice.

6. References

[1] Florian, M. and Hearn, D., "Network Equilibrium Methods" Chapter in Handbook on Operations Research – Network Routing, 8, 485-550, 1995.

[2] INRO Consultants Inc. EMME/2 User's Manual, Montreal, 1998.

[3] Magnanti, T.L. and Perakis, G., "Averaging schemes for variational inequalities and systems of equations", Mathematics of Operations Research, 22, 568-587, 1994.

[4] Magnanti, T.L. and Perakis, G., "Solving variational inequality and fixed point problems by averaging and optimizing potentials", OR 324-97, MIT, 1997.

[5] Marcotte, P. and Zhu, D., "Convergence properties of feasible descent algorithms for solving variational inequalities in banach spaces", CRT-96-12, University of Montreal, 1996.

[6] Meyer, Mohaddes Associates Inc. Heavy Duty truck Model and VMT Estimation. Report to Southern California Association of Governments, 1999.

[7] Smith, M.J., "The existence, uniqueness and stability of traffic equilibria", Transportation Research, 13B, 295-304, 1979.

[8] Wardrop, J.G., "Some theoretical aspects of road traffic research", Inst. Civil Eng., Part II, 325-378, 1952.

On Variable Link Speeds by User Class within SATURN

Dirck Van Vliet
31 July 2006

1. Introduction

This note is intended as a form of internal discussion document for circulation between myself and various people in Parsons Brinkerhof (PB), the HA and/or DfT. The objective is to review the "state of the art" of modelling variable link speeds by vehicle/user class and to suggest possible improvements to be implemented within SATURN.

Following the first "release" in May 2006 I added (14 July) various new thoughts in the light of comments received to date and Marcus Chick's' report on CLICKS. I had also been sent excerpts from the COBA Manual dealing with speed-flow curves.

The "third" release, 31 July, 2006, expands slightly on possible alternatives under Option (3) in section 3 in the light of a closer inspection of the COBA recommended curves for rural single Carriageway roads (Class 1), reformats section numbers and tries to emphasise the "major" questions to be resolved..

I would prefer to initially limit the circulation if only to keep the discussion within limits but I have no objection to anybody who receives this circulating it more widely. Once we seem to have reached some agreement and are at the stage of making changes within SATURN it could be more widely circulated.

So, the problem under discussion is how to deal with the obvious fact of life that on certain roads not all "vehicle types" (generally referred to as "user classes (UC)" here) travel at the same **average** speed. In particular on motorways HGV's may have a maximum speed limit of 60 mph whereas cars have a 70 mph maximum. Whether they keep to these limits is another question of course but the fact remains that they do tend to travel at different speeds. (N.B. We are not therefore concerned here with the **within-class** variability of speeds, day-to-day variability, etc, etc.; only the between-class differences.)

More specifically I am obviously primarily interested in how to model these effects within SATURN although it may be best to initially think about the issues in more general terms and, having devised an appropriate modelling strategy, think about implementation within SATURN later on.

Differences in speeds by vehicle class may have implications for:

(a) Assignment (route choice)
(b) Economic assessment of travel time
(c) Second-order affects such as emissions, noise, fuel consumption, vehicle operating costs, accidents, etc., all of which may be speed-dependent.

A further complication is the fact that average speeds are not fixed but depend on the flow so what we really need, first and foremost, is a **procedure** (e.g., equations) to determine the speed by UC as a function of flow(s); i.e., class-dependent speed-flow curves.

To keep the subsequent discussions simple I shall just concentrate on the simplest case of two user/vehicle classes, HGVs and LGVs/cars, although any methods discussed will be more generally applicable.

2. MUC Speed-Flow Curves: "The X-axis"

The first question to be resolved in terms of MUC (multiple user class) speed-flow curves is whether or not it is acceptable to formulate them as essentially functions of a single independent variable along the X-axis, the total pcu-weighted flow, or whether it is necessary to formulate them as an n-dimensional function of **all** UC flows. In other words, are the speeds on a motorway with 1,000 HGVs/hr - rated at 2 pcus per HGV – and 2,000 cars/hr, ie 4,000 pcus/hr, the same as with 500 HGVs/hr and 3,000 cars/hr? Or, to put it another way, are pcu factors flow independent?

Alternatively, it is also possible to make the X-axis simply "vehicles" as is apparently recommended by COBA (although there may be an implicit assumption of 15% HGVs).

It certainly makes life simpler if pcu factors **are** constant so, unless there are strong objections and/or evidence to the contrary, I shall assume that we can consider UC speeds as functions of total pcu flow only and that pcu factors will be constant throughout the network..

The first major question to be considered is what is the COBA/DfT view in this regard.

3 Available Options

There appear to be three different potential differential-speed models "on the table" at the moment:

(1) HGVs require a fixed extra time per kilometre relative to cars over all flow conditions;

(2) HGVs travel at a maximum speed until, with increasing flow, the speed of cars drops below that value, after which both HGVs and cars travel at the same speed;

(3) HGVs and cars have separate speed-flow curves with a maximum difference at zero flow but gradually converging until, at capacity (or, possibly, at a lower flow convergence point), both travel at the same speed.

Option (1) is currently implemented in SATURN as “CLICKS”. Thus, if HGVs are UC 2, CLICKS(2) = 100.0 implies a maximum free-flow speed of 100 kph for HGVs on **all** f links. (But see point 4 in Section 6.) If, say, on a motorway the free-flow speed for cars is 120 kph then this implies that HGVs take 3.6 seconds per km longer than cars not only under free-flow conditions but under all flows. Thus, if at capacity the speed of cars drops to 40 kph then the corresponding speed of HGVs (calculated by adding 3.6 seconds to the time/km of cars at 40 kph) becomes 37.5 kph.

Option (2) was initially suggested to me by Andy Stoneman based, as I understand it, on his personal travel experience / “engineering judgment” but, having now read the relevant sections of COBA, it appears that this is in line with the COBA recommendations for HGV speed-flow curves (in particular for Classes 2-6). It certainly satisfies the criteria set by option (3) that both cars and HGVs have the same speed at capacity. How about we call this “speed capping”?

Option (3) seems an obvious option if one wishes to impose the same “crawl speed” on all vehicles at capacity. There are, however, a potentially large number of methods by which the two curves could be made to converge.

Firstly, assume that the point of convergence of the two curves is to be at capacity. Thus, for example, one could simply take the “standard” car speed-flow curve and construct the HGV speed-flow curve by subtracting a fixed differential speed component (e.g., 20 kph as per option (1)) and multiplying it by (1 – V/C) such that at V = C the difference is zero. Or, and this might be more “natural” within an assignment model, one could work in terms of time/km vrs. flow and add 3.6(1 – V/C) seconds/km. Or, instead of 1 – V/C, take any other function of flow that goes from 1 to 0 as V goes from 0 to C.

Alternatively, one might wish to have two distinct forms of light and heavy speed-flow curves which converge at a flow below capacity. An example of this is the COBA specification for rural single-carriageway roads (Class 1) where both LGV and HGV speeds decrease linearly (from different free-flow speeds and with different slopes) until the two curves cross (below capacity), beyond which point the LGV curve applies to both.

So, the second major question, is there any empirical evidence to suggest which, if any, of the above options is best or should we be looking elsewhere?. I would hope that this is an area that the DfT/HA would have a view on. And, if COBA is anything to go by, the view would seem to be in favour of Option 2 and/or 3 – both of which can be dealt with in the same manner within the assignment as discussed next.

4 Implications for Assignment Algorithms

One of the reasons for the choice of the CLICKS option (i.e., option (1) above) within SATURN was that a fixed time (i.e., cost) penalty per user class per link fits very easily into the algorithm currently used to solve for MUC Wardrop Equilibrium Assignment and does not, for example, introduce any extra potential problems of non-convergence or multiple equilibria (which, in theory at least, already exist through the simulation model).

On the other hand, and having thought about the alternatives to fixed penalties a bit more, it does seem to me that it should not be that difficult to deal with "**variable**" flow-dependent penalties (e.g., options 2 and 3 above, including the various sub-options under 3) in the same way that the simulation sub-model introduces "variable" cost-flow curves for turns in the simulation network. In other words, at the start of any assignment stage within the SATURN simulation-assignment loops, set any UC penalties based on the final V/C ratios from the **previous** assignment and treat them as fixed within that assignment. The loops continue normally – as required by the simulation anyway – until the V/C ratios and the corresponding penalties converge to constant values.

For example, if for a certain road the required speed-flow curves for LGv and HGVs predict that, given the current assigned flows, LGVs take 360 seconds and HGVs take 390 seconds then the **fixed** HGV penalty in the **next** assignment would be 30 seconds. If that assignment then, say, assigns more flows to that link such that the LGV/HGV times become 380 and 405 seconds then the penalty would be changed to 25 seconds. By contrast, if the free-flow travel times were 300 and 360 seconds respectively then the "standard" CLICKS approach, option (1), would be to fix the HGV penalty at 60 seconds independent of the assigned flows.

Technically this approach is known as "diagonalisation" and is no different from that used to deal with the interactions introduced by the simulation. In principle it could lead to non-convergent looping with V/C ratios and penalties oscillating ad infinitum but, as long as the penalties are relatively insensitive to V/C ratios, e.g. in comparison to the normal dependence of costs on V/C, then I would not expect this to be a problem and certainly no worse than any current problems of non-convergence between assignment and simulation. In effect, the variable UC penalty would be calculated by "simulation".

The third major question is whether a diagonalisation option is required in SATURN.

5. Skimming O-D Times for Assessment

Once UC-specific times/speeds have been set up in the assignment model it is clearly straightforward to skim UC-specific O-D times for use in, e.g., evaluation software such as Tuba. What is perhaps less obvious is that it is equally possible to introduce UC-specific times at the skimming stage which **differ** from those used in the assignment proper.

Thus, normally, if you wish to skim "time" for a certain UC, SATLOOK selects the definition of time as used during the assignment, i.e., including any effects such as CLICKS used within the assignment. (Strictly speaking CLICKS may be optionally included or excluded as desired.) However, it is possible to create **new** link-based "attributes" (i.e., times) and to skim those attributes. So, even if CLICKS were not introduced at all at the network-building / assignment stages, it is quite feasible to define a new link property which equals "time" plus extra CLICKS-based penalties and to skim that property over the paths assigned to a particular UC. Or do the same thing with "speed-capping" or converging speeds – options 2 and 3 above.

Technically the new variables could be calculated internally to SATURN using, e.g., SATDB, or they could be set externally using, say, EXCEL and re-imported into SATURN for skimming.

Equally it is quite possible to introduce a different mathematical form of, e.g., speed-flow curves. Thus we could assign using the traditional SATURN "power law" equations but calculate times for skimming using the COBA 2-stage linear speeds vrs flow form post assignment. See also point 1 in Section 7.

Indeed there may be good pragmatic reasons for **not** wishing to include extra time penalties for, say, HGVs on motorways since very often the problem is to get **enoug**h HGVs onto the high capacity roads since the advantages of those roads to HGV drivers may not be adequately represented by the "correct" economic weighting of time and distance. Therefore you can achieve both a better assignment and take on board differential motorway speeds by introducing an alternative definition of HGV speeds at the skimming stage.

6. Secondary Effects

Secondary effects such as fuel consumption, emissions etc. are (to my mind!) much more naturally calculated at the level of individual links and, if you then wish to use an O-D based evaluation package such as Tuba rather than a link-based approach such as COBA, skimmed by O-D pair. In this way the impact of, say, differential UC speeds on a motorway link can be taken on board in the calculation of UC emissions, etc.

It is perhaps important to make the point here that, within SATURN, it is quite feasible to calculate **almost** any quantity you want to at the link level and to skim that over forests to produce o-d matrices for evaluation. For example, you could do accident calculations that way.

My understanding is that, at the moment, vehicle operating costs (VOC) in Tuba are calculated by skimming time and distance per O-D pair, using those matrices to calculate an **average** O-D speed and then calculating O-D VOC using that speed. To my mind this is not very sensible given that it is perfectly straightforward – as with the case of differential times above – to calculate a VOC per link and to skim that by O-D pair to give a "correct" O-D VOC.

7. Further Questions

A number of further issues have arisen since the first copy of this note was written and these may be summarised as follows: (N.B. See also the three "major" questions posed at the ends of sections 2, 3 and 4.).

1. If a user wished to use the **standard form** of COBA speed-flow curves (which, in terms of speed vrs. flow, have two linear segments) as the basis for their network modelling is it acceptable to approximate the 2-stage curves by a power law of time vrs flow as is standard in SATURN. My personal strong preference and recommendation is to use the power law on the basis that: (a) minimum change; (b) discontinuities make life difficult in calculating marginal cost curves; (c) the fit is generally extremely good; and (d) on the basis of empirical data it would be difficult to distinguish one from the other.

2. How, following section 2, does one define "capacity": in units of vehicles or pcu's per hour?

3. How should times for V > C be set. SATURN uses a simple linear queuing model d = (LTP/2)(V/C – 1. Advice Note 1A proposed the same general form but with LTP fixed at 15 minutes.

4. Should the value of CLICKS be disaggregated further? Or, more accurately, how much further disaggregation is required. One option would be to define a 2-D set of CLICKS value by user class and by capacity index (say). On the other hand it might be better to go the whole hog and define CLICKS for each link individually as well as by user (or vehicle) class. The disadvantage of defining CLICKS by road classification only is that that does not necessarily take into account factors such as hilliness or bendiness which feature in COBA equations for HGV's.

5. Should we be thinking of defining CLICKS by **vehicle** class in SATURN rather than user class on the basis that a model might conceivably wish to define up to 36 user classes in some circumstances (e.g., toll studies) but the number of classes with distinctly different vehicle characteristics is likely to be only 2 or 3 at most. It's basically a question of (a) simplifying user input and (b) reduced RAM.

31//07/06