Variable demand modelling

Thinking bigger with microsimulation

Boris Johansson
SIAS

Variable demand modelling (VDM) is the process used to predict and quantify changes in travel demand resulting from a change in the transport system and there has been a historical tendency to consider VDM as only being applicable to the modelling and assessment of larger scale transport schemes with anticipated impacts over wide areas. As a result, VDM has traditionally been undertaken using wider-area models that typically employ deterministic equilibrium modelling techniques to produce estimates of travel costs. More recent WebTAG guidance, however, has highlighted the need to consider VDM across a wider range of applications with less focus being placed on the size of the scheme being considered. This, suggests Boris Johansson, technical director at transport planning consultancy SIAS, will increase the diversity of VDM applications with respect to the scale and nature of schemes being assessed and the modelling software used in their assessment.

"As software and hardware technology has improved, the range of applications of microsimulation modelling within the transport planning community has increased," Johansson explains. "The abilities now exist to develop, apply and readily produce the relevant outputs from microsimulation models to assess virtually all forms of transport schemes, whether in a local or wide area context."

Microsimulation modelling can, he observes, provide a number of advantages in the assignment process including:

- Greater detail of vehicle origin/destination points within zones (e.g. on/off street parking areas, distinction between differing land uses, etc.)

- Detailed travel demand profiles (applied by five minute interval and by vehicle type/trip purpose and OD movement if required)

- Enhanced modelling of interactions between private and public transport vehicles within a single modelled network

- Improved operational modelling such as at junctions, overtaking and merge/weave areas through modelling of individual vehicles/drivers and their characteristics

- Dynamic re-routeing of vehicles due to delays experienced in the simulation

"Microsimulation modelling can also include intelligent transport systems (ITS) such as advanced traffic management (ATM) on motorways to estimate the effect this may have on journey times and their reliability and these key features provide the ability for microsimulation models to produce more accurate estimates of travel costs than traditional modelling methodologies," Johansson says. "Robust travel cost estimates form the cornerstone for VDM and ultimately for economic assessments and the use of microsimulation models in these processes therefore provides clear advantages."

Technology time

The software employed in the examples outlined below includes the S-Paramics microsimulation package, the DfT's DIADEM program and other bespoke modules. A key feature of microsimulation models is their use of 'Monte Carlo' simulation techniques, producing a range of outcomes as opposed to a single, repeatable answer. Johansson points out. This requires the outputs to be produced and averaged across multiple simulation runs. The S-Paramics microsimulation software used for the models referred to has an integrated 'batch farm' facility that exploits the available processing power of networked PCs, thereby minimising run-times and making VDM with microsimulation a practical proposition.

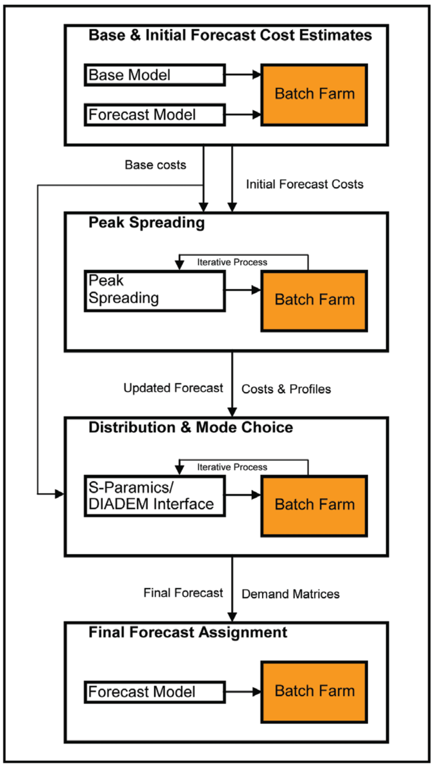

Figure one: VDM process using S-Paramics microsimulation

Figure one shows how the whole VDM process using microsimulation operates and demonstrates the central role played by the 'batch farm' facility.

By modelling detailed travel demand profiles microsimulation models offer a unique opportunity to consider micro time period choice (peak spreading) effects, Johansson says. This is done by calculating appropriate adjustments to these profiles for the relevant demand segments, typically using five minute intervals. "Peak spreading effects, if considered at all, have traditionally been reflected in a relatively crude manner using methods that adjust the percentage of the overall peak period demand that occurs in the peak hour," he adds. "Current WebTAG advice identifies micro time period choice as the second most sensitive response after re-assignment in the hierarchy of demand responses and the development of more sophisticated approaches to reflecting the effects of peak spreading are therefore an important element within the overall VDM process."

Testing has been undertaken by SIAS using a congested urban/inter-urban S-Paramics model, covering approximately 70km2, to estimate peak spreading effects via differing methods (e.g. HADES, DMRB incremental logit, etc.). The demand profiles were calibrated for the 2007 base model using observed data and specified for different demand segments (e.g. commuting, education, shopping, goods vehicles etc.). Forecast demand matrices for a 2016 'do minimum' situation were derived and assigned to the 'do minimum' network using the calibrated base model profiles.

The peak spreading process calculates a new 2016 'do minimum' profile for the relevant demand segments (e.g. commuting) defined by the user. This process uses the change in average travel costs between the 2007 base and 2016 'do minimum' models for trips departing in each five minute period in combination with the base profile to calculate changes to the proportion of trips in each five minute period. The process is run iteratively until suitable convergence is achieved and the final, revised profile adopted.

Figures two and three respectively show the difference in travel cost and the implied demand profile for commuting trips after ten iterations of a DMRB incremental logit peak spreading process and compare these with the 2007 and 2016 'do minimum' values based on the calibrated base profiles.

These outcomes follow an intuitively correct trend with the increased travel costs in the 'do minimum' situation resulting in commuting trips departing earlier in the am peak. This example, Johansson says, demonstrates that the use of microsimulation techniques can ensure that the most sensitive choices of re-routeing and peak spreading are modelled in as much detail as possible to provide more robust cost estimates to feed into the choice models for other responses further up the hierarchy (e.g. mode, destination choice etc.).

"The use of microsimulation within a VDM context introduces additional considerations with respect to stability and convergence," Johansson points out. "Microsimulation models do not generally seek to achieve a pre-determined convergence but it is essential to demonstrate stability of the assignment process to ensure that stable transport costs are passed to the demand model which in turn calculates mode, distribution and possibly other demand responses."

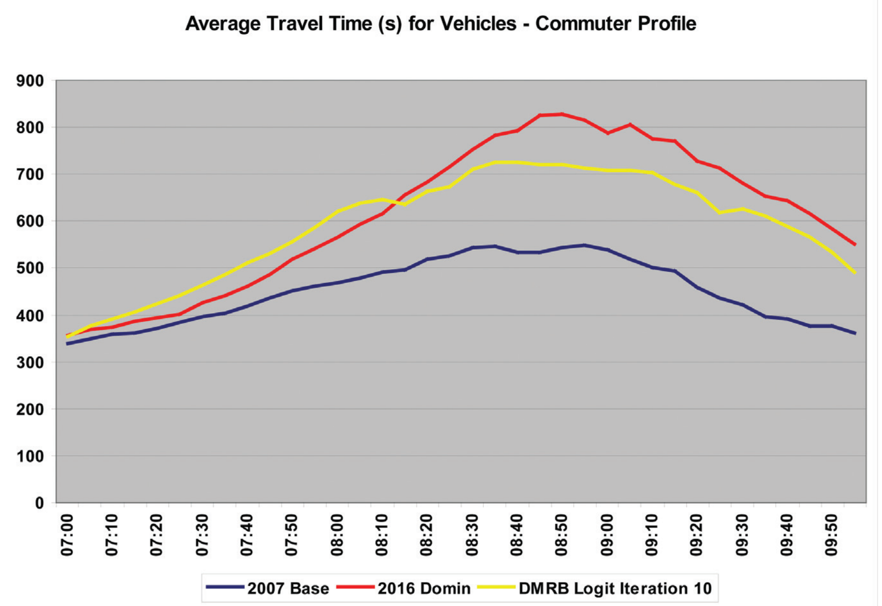

Figure two: Difference in travel costs – commuting trips

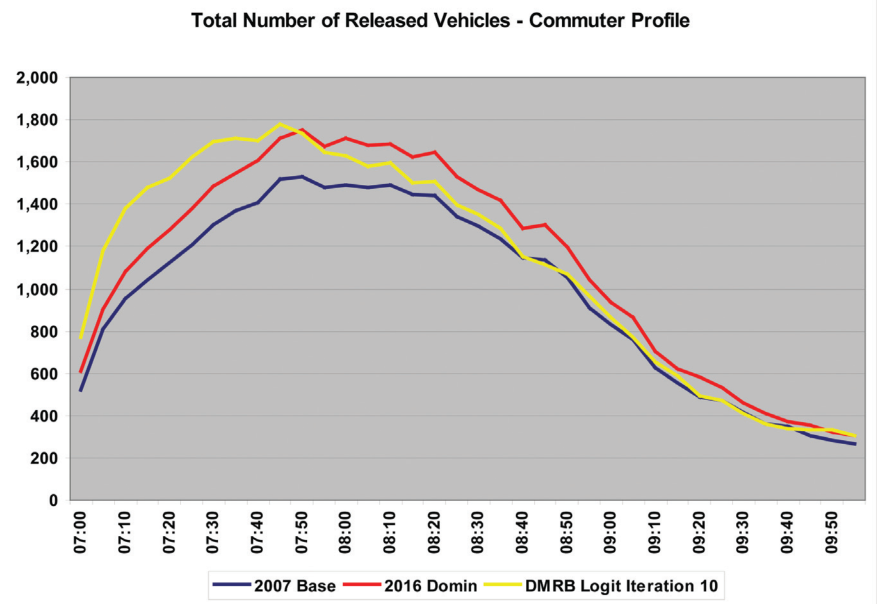

Figure three: Implied demand profile – commuting trips

The current stability measures outlined in DMRB, along with additional microsimulation-specific checks for very low flow links, can readily be used to demonstrate stability and robustness of the assignment process, he adds. The DMRB checks are largely specified as comparisons between flows and/or costs on successive iterations of a convergent assignment process. While microsimulation models do not typically employ iterative, convergent assignment processes the DMRB checks can be used to measure the stability of the rolling average flows and costs between successive different randomly seeded runs of the simulation.

Research was undertaken by SIAS on models ranging from a small model with 19 zones up to a city-wide model with 177 zones to establish the number of simulation runs required to achieve suitable stability in the assigned flows and costs. This demonstrated that the DMRB assigned flow and cost stability criteria can generally be met for even the largest microsimulation models within a relatively small number of simulation runs (e.g. less than ten).

The VDM process requires costs to be skimmed for both public transport and private vehicle trips. It is obviously essential that these costs are as robust as possible to ensure appropriate demand responses due to changes in the cost of travel by public transport and, by using microsimulation, additional detail can be included in the modelling, such as detailed dwell times by stop and service, reflection of schedule and frequency based services, detailed interaction between public transport and private vehicles in the network and detailed representation of public transport priority measures such as bus lanes, selective

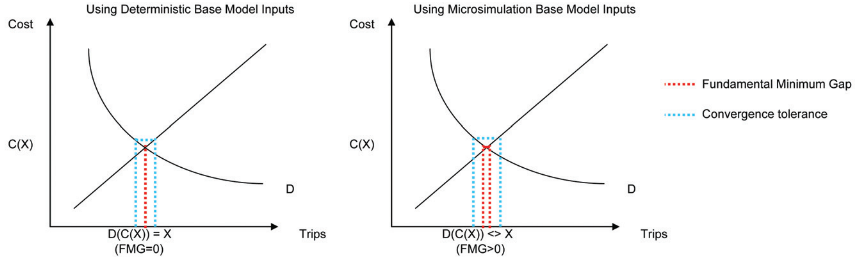

Figure four: Fundamental minimum gap

vehicle detection, etc. The effects of these more detailed aspects of public transport modelling are directly reflected in the transit times output from the completed assignment.

The process of skimming the public transport costs is similar to that adopted for traditional assignment models. Public transport stops are associated with the relevant zones in the model while walk links and interchanges are defined as appropriate. The costs are then skimmed by combining the relevant walk, interchange and waiting times with the on-board transit times output from the simulation. This ensures that robust estimates of the public transport costs, including the effects of the more detailed modelling afforded by microsimulation, are in turn reflected in VDM and any ancillary assessment processes.

Number crunching

"Convergence within the supply/demand loop is an essential requirement to ensure that any demand changes predicted are as a genuine result of a change in transport costs and not influenced by model noise or lack of convergence," Johansson explains. "The issue of supply/demand convergence within a microsimulation context, however, requires additional consideration due to the inherent variability within 'Monte Carlo' simulation."

Traditional equilibrium models produce unique results that are exactly repeatable every time the model is run with a given set of inputs. Current WebTAG guidance suggests that the %Gap supply/demand convergence measure should ideally be less than 0.1% with less than 0.2% being accepted in problematic systems. These convergence thresholds were derived in the context of using traditional equilibrium models in the VDM process and can be thought of as the tolerance level above the minimum gap that is considered acceptable.

In the case of microsimulation models the simulation results are intentionally not repeatable every time the model is run even using the same inputs, although the stability of the assigned flows and costs from microsimulation models has been demonstrated. Tests were therefore undertaken by SIAS to isolate the inherent variability and assess its effect on convergence within the supply/demand loop using DIADEM.

The tests used the same input networks and matrices to represent both the base and reference cases. The average costs from a set of multiple runs of the base model were then input to represent the base costs, the only variable in the process being that the averaged costs from an independent set of multiple runs of the same base model were input as the reference case costs. A DIADEM run was then undertaken and the %Gap statistic inferred from the cost differences calculated.

This process, if run using a traditional equilibrium model, would produce a %Gap statistic of zero, as the costs from multiple runs of such a model would be identical for a given set of inputs. Applying this process using microsimulation models, however, produces %Gap values greater than zero due to the inherent variability. The %Gap produced using this process can be thought of as the fundamental minimum gap (FMG) which, for deterministic models, will equal zero and for microsimulation models will be greater than zero. Figure four demonstrates the concept of the Fundamental Minimum Gap using supply/demand equilibrium curves.

The current DfT guidance allows for up to a 0.2% tolerance around a FMG of zero for deterministic models. In applying microsimulation models in the VDM process, however, Johansson suggests that there may be a case for applying this tolerance level around an FMG value that is greater than zero in recognition of the variation inherent in the simulation outputs.

A series of tests was undertaken by SIAS using two models of very different scale and nature to identify the FMG. Model one covers a large urban and inter-urban area of approximately 70km^2 with 121 zones and where significantly congested conditions are experienced. Model two has a total of 19 zones and covers a small peri-urban area of less than 3km^2 where relatively uncongested conditions are experienced. The FMG tests were run as outlined above and results produced based on the average of differing numbers of successive simulation runs. The results demonstrated FMG figures as shown in table one.

The results demonstrate an intuitively correct trend with the FMG being greater than zero in all cases and the smaller, less congested model two having a much smaller FMG than the larger, more congested model one, Johansson says. As expected, the FMG is also generally shown to reduce as the number of individual simulation runs used in producing the average cost inputs to DIADEM increases.

In the case of model two, the FMG is less than the DfT specified upper %Gap acceptability threshold of 0.2% when five or more simulation runs are used to derive the average cost inputs. Applying the acceptability threshold to the FMG from 20 simulation runs would imply that a %Gap of 0.32% (0.12% FMG + 0.2% threshold) or less would be required to ensure that appropriate convergence was achieved within the supply/demand loop.

In the case of model one the FMG is shown to be larger and only begins to approach the DfT upper acceptability threshold after around 20 runs of the simulation. Again, applying the acceptability threshold to the FMG from 20 simulation runs this would imply that a %Gap of 0.41% or less would be required to achieve acceptable supply/demand convergence.

"The FMG should be minimised as far as possible by ensuring that measures are taken to address noise in areas of the model remote from the scheme being assessed and by undertaking the maximum practicable number of simulation runs possible," Johansson concludes. "In this regard it would be beneficial for the DfT's guidance to provide clear advice regarding the use of microsimulation in the VDM process." 🟩

Number of Runs	Fundamental Minimum Gap (%)	
	Model 1	**Model 2**
3	0.54%	0.32%
5	0.39%	0.17%
10	0.28%	0.18%
20	0.21%	0.12%

Table one: Estimates of fundamental minimum gap