CONVERGENCE OF THE SYDNEY STRATEGIC TRAFFIC MODEL – IMPLICATIONS FOR ECONOMIC ANALYSIS OF URBAN ROAD PROJECTS

Matthew WILSON, Xiaoping GUO & Marwan S DAIZLI
Roads & Traffic Authority of NSW

Abstract

The Roads and Traffic Authority of NSW maintains and operates an EMME/2 model of metropolitan Sydney, Australia. Outputs from the model are used, amongst other things, to estimate the economic benefits of road projects. The precision with which project benefits can be estimated is directly related to the convergence of the equilibrium assignment. A detailed investigation of model convergence has quantified the reliability of model-based economic analyses. Where single model runs are not able to provide sufficiently reliable results, alternative methods are suggested.

1. Background

1.1 The Sydney Strategic Traffic Model.

The Roads & Traffic Authority of NSW (RTA) uses EMME/2 as the platform for its Sydney Strategic Traffic Model (SSTM). The SSTM has its origins in the Sydney Area Transportation Study, carried out by consultants De Leuw Cather in 1971-73, and has since evolved using TRIPS, TRANPLAN and now EMME/2 software. The model currently covers metropolitan Sydney plus the outer suburban Central Coast and Blue Mountains districts, extending approximately 110 kilometres from north to south, and 70 kilometres from east to west. The model includes some 1,100 zones, 12,000 nodes and 32,000 links. It is mainly used to model a 2 hour morning peak period, but approximate demand models for other times of day have been prepared by a process of matrix inversion, factoring and adjustment. Models are maintained for the current year and for future years at 5 year intervals (2011, 2016, 2021, etc, representing future census years).

Macros have been developed using the network calculator to produce a range of output statistics including traffic noise, gaseous emissions, vehicle operating costs, vehicle kilometres and vehicle hours of travel. Aggregate measures of vehicle travel time, distance and operating cost are used as inputs to economic analysis of projects.

1.2 Model convergence

The EMME/2 traffic assignment module uses an iterative method to approach a true equilibrium assignment, in which no driver can find a better path through the network than the path they are assigned to. This process may not reach an absolutely stable equilibrium. Some sophistication is needed to ensure that the process converges towards a reasonable end result within a reasonable number of iterations. EMME/2 allows the user to define an acceptable level of convergence, in terms of the difference in mean travel times between successive iterations (termed the normalised gap), or the percentage difference between the mean travel time in the current iteration and an estimate of the value for a perfect equilibrium assignment (termed the relative gap). The assignment process is stopped when the specified point is reached.

The equilibrium assignment uses a linear approximation method, developed from an algorithm first proposed by Marguerite Frank and Philip Wolfe (Frank and Wolfe, 1956). In an interview with Irvin Lustig, published on www.e-optimization.com, Phil Wolfe himself described this as:

> "..a somewhat clumsy method of doing it, now known as the Frank-Wolfe algorithm. While it's a precursor to quadratic programming, it turned out to have a lot of applications in other areas where you have messy objective functions and linear constraints."

David Boyce et al note that this algorithm gradually shifts traffic from slower to quicker routes, but the assignment only converges slowly, if at all, and "when applying this method, link flows fluctuate substantially from iteration to iteration, or gradually drift up and down" (Boyce et al, 2004)

SSTM assignments do converge, but relatively slowly, especially after the first 10 iterations or so. The more congested future year models require more iterations than the less congested base year models to achieve a given level of convergence. Convergence to a relative gap of 1% is usually achieved within 15 to 20 iterations, which takes about 3 minutes of computer time at a rate of 6 iterations per minute. Convergence to 0.1% generally requires 40 or more iterations. There are a number of locations in the model network where cyclic switching of trips occurs between congested parallel routes, even after 50 or more iterations. However, relatively small numbers of trips are involved, and the model has no difficulty providing stable estimates of traffic volumes on links.

1.3 Application of the model to economic analysis

RTA uses the SSTM for more than just prediction of traffic flows. When a project is expected to have impacts on traffic flows over a significant area of the network, its economic performance is usually assessed using SSTM outputs. This assessment uses the network calculation facilities within EMME/2 to calculate total road user costs across the whole modelled network. These costs are composed of–

- The economic value of travel time
- Vehicle operating costs, based on modelled speeds
- The economic cost of traffic crashes, based on standard rates of crashes per million vehicle kilometres travelled, on three different categories of road (freeway, surface arterial and subarterial-local).

The economic benefit of a road project is assessed by comparing costs extracted from 'base case' and "improved case" networks, (ie networks without and with the project being analysed). The difference represents the benefit of the project in the modelled year. The benefit estimated from the 2 hour morning peak models is then factored up to produce an estimated annual benefit. This process is usually repeated for two or more modelled years, and linear interpolation and discounting are then used to produce an estimate of the present value of the project's benefits.

2. The problem – Stability of model-based estimates of benefit

The estimation of project benefits relies on calculating the difference in modelled road user costs between two different assignment runs. To obtain a reliable estimate of project benefits, the model must have converged sufficiently that the difference in road user costs calculated from successive iterations of the same model is small compared with the difference in costs calculated from the two different networks. The "signal" must be significantly stronger than the "noise".

If a project costing $AU100M has a benefit cost ratio of 2, it must generate benefits of $AU15 to $AU20 million per annum. The total road user cost across the whole of the modelled network is much larger than that – in the 2004 model, about $AU13 million in a single two-hour morning peak, equivalent to $AU23 billion per annum. The benefits of that $AU100M project therefore represent a little less than 0.1% of the total road user cost in the modelled network. For a $AU10M project with a BCR of 2, they would be less than 0.01% of the total.

When the model has been used for economic analysis of some relatively small projects, with capital costs in the range of $AU5M to $AU50M anomalous results have occasionally been observed, prompting the investigation reported here.

3. Observed model convergence

At a gross level, the assignment of SSTM appears to be very well behaved. Figure 1 shows the convergence of total road user cost in the 2004 model network, for the first 50 iterations.

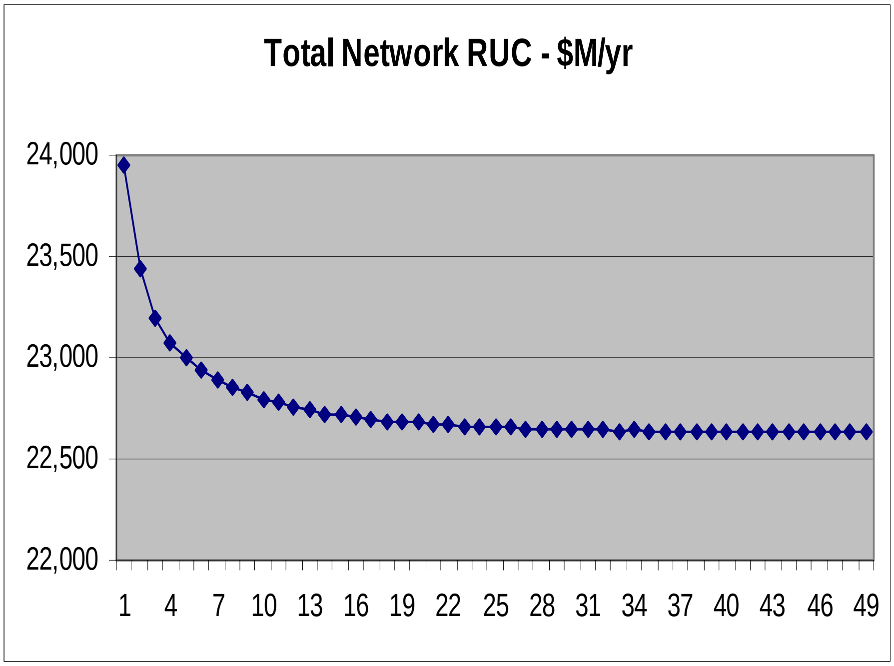

Figure 1 – Convergence of Total Road User Cost

However, the differences between successive iterations are larger and less well-behaved than Figure 1 might suggest. Figure 2 shows that, even after 40 iterations, the total road user cost is oscillating rather than converging smoothly, and the differences between iterations are of the order of $2M per annum.

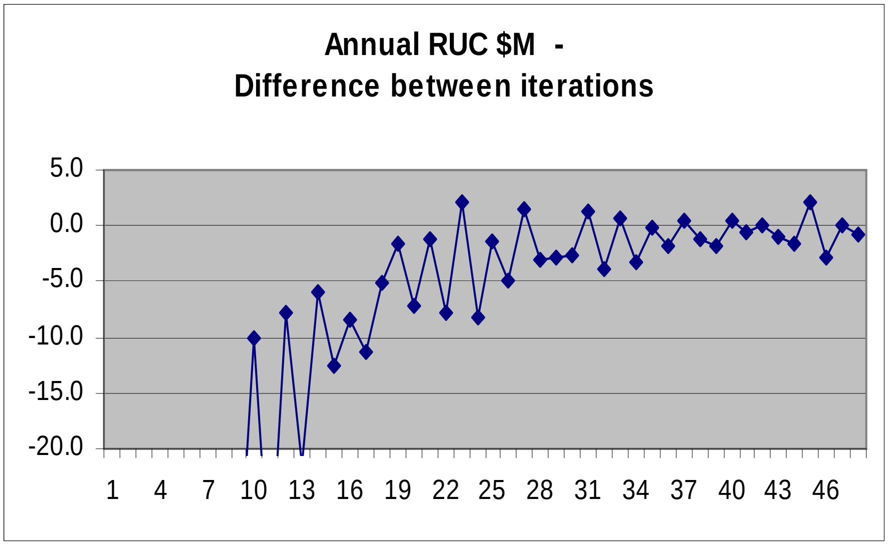

Figure 2–Differences between Iterations

Figure 3 shows the differences between the last 10 iterations in more detail.

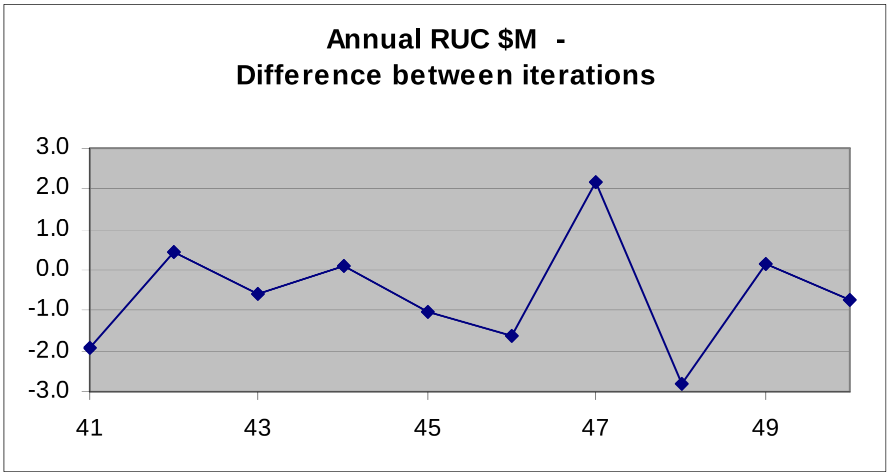

Figure 3–Differences between Iterations 41 to 50

After 40 iterations, the estimate of annual road user costs is stable to within about ± $3M. In fact, the standard deviation of the total costs from those last 10 iterations is $1.67M, or 0.007% of the average total cost. That is of the same order as the benefits of the hypothetical $10M project mentioned above, so its estimated benefit cost ratio (BCR) might be reliable within a range of ± 2–it probably lies between 0 and 4. This level of accuracy is insufficient for evaluating this project. The problem is likely to be more serious in models representing future years, because they will be more congested, and therefore less stable.

4. Possible solutions

Four possible solutions to the problem have been examined, using a 2016 model to evaluate a real project proposal–a bypass of a suburban main street, with an estimated capital cost of $AU80M.

4.1 Increase the number of model iterations

One might expect that more stable results should be obtainable simply by increasing the number of iterations, to obtain a higher degree of convergence. Since each iteration only needs about 12 seconds of computing time, there is ample scope for doing so. Figure 4 shows that the estimate of project benefits is distinctly less stable below 40 iterations than beyond that point. However, there is no indication of any further increase in stability between 40 and 70 iterations. Table 1 confirms this. The costs from single iterations are only reliable to within about $0.7M per annum.

There seems little prospect that the stability will improve if the model is run beyond 70 iterations.

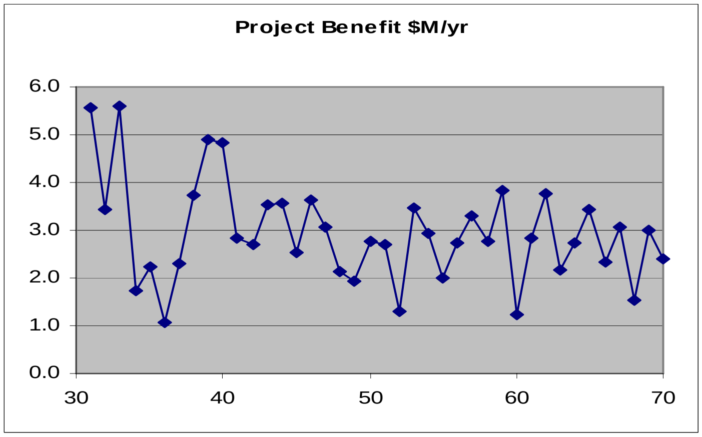

Figure 4–Project benefits by iteration

Iterations	Mean Estimated Benefit ($M/yr)	Standard Deviation
31 to 40	3.54	1.65
41 to 50	2.87	0.59
51 to 60	2.63	0.87
61 to 70	2.73	0.65

Table 1–Stability of Estimated Benefits
with Increasing Iterations

A variation in annual benefits of $0.7M is equivalent to a variation of about $9M in the present value of benefits over the life of the project. If the calculated BCR is to be reliable to within plus or minus 0.2, the capital cost of the project could not be less than $45M.

4.2 Use average values from a number of successive iterations

Table 1 also shows that, beyond 40 iterations, a relatively stable estimate of benefits can be obtained by taking the average of 10 iterations. This average lies somewhere between $2.6M and $2.9M per annum–that is, it is reliable to within plus or minus about $0.15M per annum. That would provide an acceptable estimate of BCR for a project costing $10M.

Some time is required to extract the data from the model manually, however, so there is a significant effort in extracting data from 10 separate iterations. Averaging a smaller number of iterations might give a sufficiently reliable estimate of benefits. Average benefits were calculated from five groups of 2 iterations, starting with iteration 41–that is, iterations 41 & 42, 43 & 44, 45 & 46, 47 & 48, and 49 & 50. The standard deviation of those five results was then calculated, to give an indication of the variability of the two-iteration average. This process was then repeated for five groups each of 3, 4 and 5 iterations. The results are shown in Table 2.

Averages of -	Mean Benefit $M/yr	Std Devn
2 iterations	2.9	0.46
3 iterations	2.7	0.42
4 iterations	2.7	0.36
5 iterations	2.8	0.23

Table 2–Estimated Benefits Averaged from
Increasing Numbers of Iterations

This suggests that projects costing $15M to $30M could be evaluated reliably by averaging 2 to 5 successive iterations. It is apparent that the effort involved in evaluating such projects will be inversely proportional to the cost of the project! A limit would need to be placed on such a process, unless it could be automated (perhaps using EMME/2 macros) to reduce the manual effort involved.

4.3 Extract road user costs from a limited portion of the model network

It seems reasonable to expect that if costs extracted from a large network show a certain variability, a smaller sub-network would show a lower level of variability. More reliable estimates of benefits should be possible for smaller projects by restricting the calculation of road user benefits to a portion of the network surrounding the project. That was most conveniently done by Local Government Area (LGA), since every link in the Sydney model has an LGA code attached.

To test this method, benefits for the test project were extracted from links within the five surrounding Local Government Areas. This area accounts for about 10% of the total road user costs in the model network. It would be reasonable to expect that it would account for about 10% of the total variability in the model. Table 3 compares estimates of the benefits of the project obtained from iterations 51, 52 and 53, from the full network and from the five LGAs listed.

Iteration	51	52	53	Mean	Std Devn
Whole Network Benefits $M/yr	2.7	1.3	3.5	2.5	1.09
Benefits in five LGAs $M/yr	1.6	1.0	3.0	1.9	1.03

Table 3–Estimates of Benefits from Full and Partial Networks

Surprisingly, the restriction to five LGAs hardly reduces the variability of the estimate at all. Perhaps the five LGAs selected, being in inner and middle suburban areas, are more congested (and therefore less stable) than the average across the network. Perhaps it is a random occurrence that this particular area is especially badly behaved in those specific iterations. Further investigation would be needed to determine a workable system for applying this method.

4.4 Set a minimum capital cost of projects suitable for evaluation using the SSTM.

It has been demonstrated that, after 40 iterations of the model, the results from a single iteration should be stable enough for the economic analysis of projects costing $45M or more. The average of 2 to 5 iterations should be stable enough to evaluate a project costing $15M to $30M. Further testing could establish the number of iterations needed to evaluate projects for different ranges of capital costs, and the minimum cost of a project .

5. Conclusions

If the SSTM is allowed to converge to 0.1%, total road user costs calculated from single model iterations provide a valid basis for economic evaluation of projects with a capital cost exceeding $AU45M.

Higher precision can be achieved by averaging the results of a number of model iterations. The number of iterations would need to be determined in each individual case to achieve the necessary precision. Indications are that averaging up to 5 iterations should provide sufficient accuracy to evaluate projects costing as little as $15M.

If an excessive number of model iterations would be required, it would indicate that the SSTM is too coarse a tool to evaluate that particular project. Consideration should be given to an alternative modelling approach, perhaps using a simulation model such as Dynameq to examine a smaller local network, or using an intersection-based model such as Transyt, with fixed traffic volumes.

The results reported here apply only to the Sydney Strategic model. However, the method used to evaluate the stability of the results would be applicable to any EMME/2 model, and would be a worthwhile test if it to be used as the basis for economic evaluation of road projects.

References–

M. Frank and P. Wolfe. An algorithm for quadratic programming. Naval Research Logistics Quarterly 3, 1956

D Boyce, B Ralevic-Dekic and H Bar-Gera. Convergence of traffic assignments: How much is enough? ASCE Journal of Transport Engineering, Jan/Feb 2004