23rd Australasian Transport Research Forum
Perth, Western Australia 29 September – 1 October 1999

Using Transport as a Measure of Access to a Public Health System

John Henstridge and S Boddington
Data Analysis Australia Pty Ltd, Western Australia

Abstract

The design of a public health system places great emphasis upon the location of services where they can be readily accessed by the public defined as both patients and other persons who interact with the health system. Traditionally health planners have used relatively simple models of proximity such as Euclidean distance and have neglected the importance of public and private transport systems in their formal assessment of various possible locations of health services. When setting up a data model for metropolitan health services in Perth, Western Australia, Data Analysis Australia saw a need to include a transport network model due to the importance of barriers such as the Swan River and major transport structures such as the single north-south freeway. The model included pedestrian and mixed mode journeys that resulted in a network model that included virtually all Perth streets. This remained manageable since relatively few destinations needed to be considered. The impact of the transport model was significant in choosing between a large number of otherwise very similar scenarios, whilst the optimisation process itself raised interesting issues of branch and bound algorithms applied to transport models.

Contact Author

John Henstridge
Data Analysis Australia Pty Ltd
97 Broadway
Nedlands WA 6009

Phone: +61 8 9386 3304
e-mail: john@daa.com.au

Fax: +61 8 9386 3202

Introduction

When locating outlets for government services, the principles of access and cost efficiency are critical. The need is to optimise the placement so that the best access for the community is achieved while costs are minimised. For centralised services where only a small number of outlets are planned, it is clear that siting on major transport routes is essential and often the precise choice of site is dominated more by issues of site availability. However, when the service is to be distributed across a number of outlets, typically greater than five, the optimisation becomes more complex and the precise quantification of access becomes essential. If the development of such outlets is to be staged, still more complexity is introduced since the sequence in which the outlets are to be built must also be optimised.

Many facets of the construction of the optimisation process are directly affected by transport issues. In addition the application of a travel model to an optimisation problem has influenced the structure of the model. Here we describe how transport issues have been incorporated into optimisation models developed for health and other services in Perth since 1994.

The experience of these models is that they add considerably to the more traditional health planning procedures as illustrated by, for example, the Toronto study of Ruth and DeBoer (1996) which used simple Euclidean or "as the crow flies" ideas of distance.

Models

Mathematical models have been commonly used to better understand the effect of location in the performance of retail outlets. These models are commonly based upon gravity concepts – attractiveness of shopping centres and the mobility of shoppers. The location of many government services faces similar issues but with a different emphasis:

- Shopping facilities are usually in competition with each other whilst different outlets of the one government service should work collaboratively. The aim of each service outlet is not to maximise share but to collectively offer the best coverage.
- Government services generally target a different client base – frequently lower socio-economic and older demographic groups are of least interest to most shopping centres.
- Mobility of clients is a greater issue due to both reduced car ownership and disabilities frequently associated with age and poor health.

The result is that models for government services need to have a much greater emphasis on achieving acceptable and equitable levels of access. It is impossible to achieve uniform access to services but it is necessary to minimise both variation in access and the worst case access. This requires simultaneously optimising the locations of many outlets.

Some other aspects of the models for services are virtually the same as for retail centres, with minor differences in how they are treated due to limitations in what information is available. Demographic factors dominate measures of demand but frequently government services have far more detailed information due to their near monopoly role. For example, in planning health services it is reasonable to expect that several years' historical measures of usage of services are available. This information will typically include age, gender and address, making it possible to relate the historical information to geographical databases such as the Census. The extra level of detail for demand makes it appropriate to have a similar level of detail for issues of access.

A final consideration is the greater call for objectivity or accountability with government services. Having a defensible methodology is expected and helps limit subjective criteria dominating decisions. At the same time, a model cannot fully represent every consideration in the decision making process. The models form an important, but not complete, component.

Modelling access

The aspect of access to a facility considered here is that which is purely related to location. This is intimately linked to travel to a location and leads to measuring accessibility in travel terms - either time or cost. Cost is a complex function of mode, distance and time and requires modal choice models for the subset of the community at which the service is aimed. Furthermore many public transport ticket options have a low or zero marginal cost for extra distances and are not particularly helpful in optimisation - secondary criteria such as time must still be used. Hence it is more appropriate to use time directly.

Travel times require a model of the transport network. Since there is a need to consider both public and private transport consistently, this ideally should include all roads and lanes, all freeways, all train lines, all bus routes *and all linking pedestrian paths*. Our experience is that it is possible to maintain such models in a high level of detail, without resorting to forming aggregate zones and just major transport links. For a model of Perth the network had approximately 80,000 nodes corresponding to every intersection or curve in a road.

The transport networks overlap giving multiple links between nodes to which different speeds can be attached corresponding to different transport modes. The pedestrian links at stations between the train lines and the road system were critical to represent the limited access to trains – it would have been ideal to do a similar representation of bus stops. The only links not modelled were bicycle paths and ferries.

Calculating travel times with such a complex network is computationally feasible and is in fact easier to implement than developing a method of simplifying the network. A minimum cost flow algorithm based upon the original structure of Dijkstra (1959) can compute times from all network nodes to a specified node in less than one minute so there is no incentive to make it more efficient. As we have implemented it the limitations were:

- The absence of time penalties associated with controlled intersections and curves. These can be added with minor increase in computation.
- No modelling of public transport timetables. This can be justified through assuming that regular users of public transport choose their travel times to achieve minimal waiting times between stages. Realistic modelling of the impact of timetables also requires detailed knowledge of the time component of the demand for services. This detail is not generally available.

Geographical measures of demand for services are determined at the level of the Census Collection District (CD). They are primarily based upon the demographic profile, using age and gender specific rates of service utilisation. Other measures that can be used include socio-economic factors and ethnicity. It is also possible to use actual historical measures of utilisation but frequently these are affected ease of access and are not geographically detailed. (See for example Glover and Woollacott, 1992.)

The task of connecting CDs to the transport model was done by generating for each network node four virtual nodes displaced 30 metres from the node, one in each of the cardinal directions. Travel times to or from a CD were then calculated by averaging the travel times to each of the virtual nodes in the CD. This gave substantially better results than the more traditional method of creating a node for each zone and attaching that to the network at one point. The procedure also had the benefit of being possible to implement automatically in a standard Geographical Information System, Mapinfo.

One outcome of having this level of detail in the travel model is that travel times vary geographically in a relatively smooth fashion. Both theory and experience indicate that such smoothness of the criterion greatly assist the performance of optimisation procedures.

Catchments

Travel time models can be applied to hospital systems to define "natural" catchments. These can be a substantial improvement upon simple distance models used by Ruth and DeBoer that assign each point to its physically closest outlet, giving polygonal catchments that are frequently illogical.

Catchments are frequently a starting point for understanding the operation of an existing health structure. They also provide a suitable method of planning a new structure since their fundamental principle is that patients should go to the most convenient or closest facility. For this to work, it is essential that the best possible concept of "most convenient or closest" should be used.

Optimisation

The ultimate aim is to optimise the placement of outlets by selecting a combination of locations such that when customers travel to their closest one (as measured by travel times) the weighted average travel times for the whole population is minimised. If there are n possible locations and k outlets are being considered, then there are $\begin{pmatrix} n \\ k \end{pmatrix}$ possible configurations that must be evaluated. With a small problem where n maybe about 30 and k about 6, is it is tempting to try to eliminate some of the many (593,775) possible configurations by heuristic means. For Perth where there are about 2000 CDs in the metropolitan region, approximately one billion numbers are required. Attempts at reducing the volume of work by heuristic rules to eliminate certain configurations have only a small impact. It is more important to use an appropriate algorithm to span the lattice of all configurations.

An efficient algorithm makes use of the simple fact that adding an extra service outlet to an existing configuration can only decrease or leave travel times unchanged for every CD. Alternatively, taking away an outlet can only increase or leave travel times unchanged. Thus bounds on average of maximum travel times can be derived. This leads to branch and bound algorithms as have been used in other areas such as the travelling salesman problem. (See for example, Reingold *et al*, 1977.)

Optimisation over time

The development of a network of outlets must usually proceed by stages. This can greatly complicate the optimisation process since as well as deciding on the optimal final configuration, there is a need to decide the best path to reach it. If in the example above there are to be three stages each with two outlets, there are $\begin{pmatrix} 30 \\ 2 \end{pmatrix}\begin{pmatrix} 28 \\ 2 \end{pmatrix}\begin{pmatrix} 26 \\ 2 \end{pmatrix}$=53,439,750 possible development scenarios, 90 times the number of final configurations. Furthermore, if the optimisation criterion places weights upon the value of the development at each stage rather than only at completion, the optimal staged solution might not have the same end configuration that a single stage optimisation would give.

Such stage weights are common when the development takes place over an extended period such as ten to twenty years, reflecting the growth and spread of population. They provide the mechanism by which outlets in the outer metropolitan region are delayed until the demand is ready for them.

Again using travel time as the criterion in the optimisation permits an algorithm that is substantially more efficient. The most time consuming step is the evaluation for each configuration of what is the closest location to each CD. Once each CD has a location assigned, the evaluation of travel times, weighted by demand, is relatively straightforward. Hence an efficient algorithm must minimise the effort in assigning locations while at the same time maximising the utility of each such assignment.

The principle is again "branch and bound" but the implementation is the reverse of the normal. The assumptions are that having more locations at a given stage can only decrease travel times and that the development of the travel network over time only improves access. Hence the travel times evaluated with a configuration with (say) k

locations provides a lower bound on the travel times for all smaller configurations contained in it for that stage and earlier stages.

The first assumption is simply requires that no artificial restrictions be applied to prevent access to the closest outlet. The second relates to the development of the network and in some cases it might not hold. It can be relaxed slightly to permit almost uniform slowing of the transport network as may be brought about by speed restrictions.

The algorithm is illustrated with the example in Figure 1 where each node corresponds to a possible configuration at one of the stages. Suppose that the final configuration A has been fully evaluated together with the sequence leading to it as indicated by the heavy links. Further suppose that the configuration B has been evaluated for Stage 3 that the closest of the six locations to each CD has been assigned It is then possible to use these assignments for six locations with the Stage 2 weightings to give lower bounds for the travel times in each of the contained configurations for Stage 2 (shown in gray). If these lower bounds are not as good as that exactly calculated for the scenario leading to A then all the scenarios leading to B can be eliminated from further consideration This removes the two shaded nodes at Stage 2.

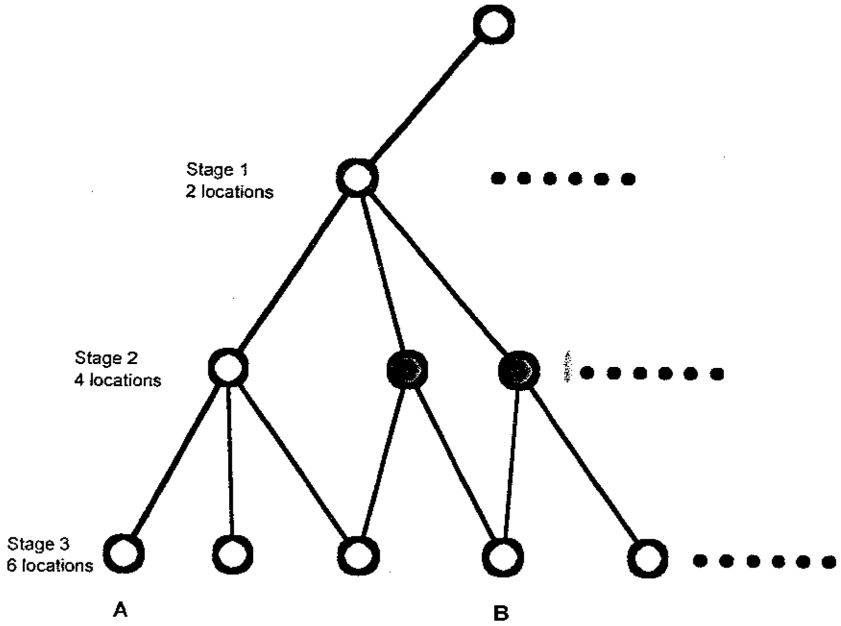

Figure 1. An example of possible scenarios for a three stage development. Each node corresponds to a configuration of locations and a path from the top level (corresponding to the initial conditions) through to the bottom level corresponds to a development scenario.

As with most branch and bound algorithms, the efficiency, compared with simply evaluation all possible scenarios, increases with the size of the problem What is not so obvious but just as important in its application is that it relies for its efficiency upon substantial differences in the weightings from one stage to the next. Hence it is most effective when applied to areas where there are substantial demographic shifts over the time span being considered, for example where the staging is designed to handle a growing metropolitan area.

Practical Application

The application to Perth illustrates several benefits of this approach. Perth has a several features that make simplistic concepts of distance inappropriate. The Swan River provides an east-west transport barrier with limited crossings and non-trivial shape, the north-south freeway has a major influence upon private transport while the rail system with just four well-spaced routes imposes structure on public transport.

The initial application of the travel models was to the understanding of "natural" catchments of the three general teaching hospitals. For example, the closeness of Royal Perth Hospital to the hubs of the road, rail and bus networks gave it a large natural catchment whilst the barriers to private transport provided by the river and an adjacent railway limited the natural catchment of Sir Charles Gairdner Hospital.

Designing a system with substantially more outlets showed that there were frequently a number of configurations or scenarios with similar levels of overall access. It has been found particularly useful to be able to review a number of these scenarios since they provide an understanding of what factors control the optimality:

- Typically some locations will be contained in most of these sub-optimal configurations suggesting that they have a high degree of optimality in themselves.
- Some pairs of locations never appear together indicating that they are alternatives to each other
- Of still greater interest is the evidence they give of long range effects whereby the choice of a location on one side of the metropolitan region may, through a chain of intermediate choices, affect a choice on the other side of the region.

The availability of a number of sub-optimal models also allows consideration of other non-transport factors such as local site availability and adjacency to other services. Ideally these could be incorporated into the optimisation but this requires giving weights to the various components.

It should be noted that the identification of a number of models with similar optimality is in part a result of the smoothness of the travel time model. The earlier study of hospital catchments had shown that the differences in travel times to at least two of the three hospitals were for much of Perth relatively small - less than five minutes. As more locations are considered with the aim of improving access, small time differences become more important but average travel times are also less.

Summary

The application of travel time models to the assessment of access to health facilities has been shown to be a valuable addition to methods previously applied in health planning. Where a number of collaborating outlets are being planned, it provides a natural method of optimising their location so that access for the community as a whole is improved. One requirement is that the travel time models have a degree of smoothness greater than is commonly implemented.

Work done for the Perth Region has shown that the methods are feasible to implement provided that care is taken in the choice of optimisation algorithms. They can be applied to regions consisting of thousands of area units such as Census Collection Districts and to dozens of locations.

This approach to optimisation of locations does not balance the likely demands that will be placed upon each of the locations - balance is only a by-product of the methods tendency to space out optimal locations so that average travel times are minimised. Modifications to the optimisation would be required if this balance is to be improved.

References

Dijkstra, E.W. (1959) A note on two problems in connection with graphs, *Numerische Math.* **1**, 269-271.

Glover, J. and Woollcott, A. (1992) *A Social health Atlas of Australia*, Australian Bureau of Statistics, Cat. No. 4385.0.

Reingold, E.M., Nievergelt, J. and Deo, N. (1977) *Combinatorial algorithms*, Prentice Hall, Englewood Cliffs.

Ruth, A. and DeBoer, D.P. (1996) Patient origin and market share - tools to assist in hospital planning, in *Patterns of Health Care in Ontario*, ed. Goel, V. *et al.*