Planning ahead in a harvesting variant of the travelling salesperson problem

J. Arundel, S. Winter

Department of Geomatics, The University of Melbourne, Australia
Telephone: +61 414 379 616
Fax: +61 3 9347 2916
Email: j.arundel@pgrad.unimelb.edu.au, winter@unimelb.edu.au

1. Introduction

As human beings, the choices we have available to us in the future often depend on the decisions we made in the past. This phenomenon can also apply to route choices, yet most research focuses only on the case of how good routes are chosen in a single time period. It is important to understand the qualitative and quantitative differences between routes chosen for a single time period and routes chosen over an extended period when developing realistic models of human behaviour. In this paper, we study these differences for the harvesting variant of the travelling salesperson problem (TSP).

The TSP is a well established combinatorial optimisation problem studied in operations research. In its general form, the TSP seeks the shortest tour given a set of points and the pairwise distances between these points. Among the many variations of the general TSP, there are a class of problems inspired by finding efficient tours for vehicles required to pickup and/or deliver goods at each of the points in the tour (Parragh, Doerner et al. 2008). There is also a more specific form of this problem that has applicability in the harvesting of natural resources.

In harvesting problems, there may be multiple sites at which harvesting can be performed, and multiple existing deployments of harvesting equipment. The operator must efficiently redeploy the harvesting equipment from their current sites to the new harvesting locations. Note that the concept of harvesting here is a general one; examples of harvesting processes include environmental data collection, stock grazing and commercial beekeeping. The common factor is that homogeneous objects are redeployed into an environment in which the harvested resource is renewed over time.

The most important distinction with harvesting problems from the standard TSP and its variants is the dependency between temporally adjacent solutions. Traditional pickup and delivery problems assume no dependence between the solutions for any two problems. However, in harvesting problems the delivery locations for any given solution assume the role of pickup locations when the problem is next solved. When there are more options for locating an object than there are objects to occupy the locations, this introduces the possibility that adopting the optimal solution in the current iteration may effectively penalise the next iteration and thus lead to a suboptimal solution overall.

A second consideration is the case where the delivery locations visited are not considered to be equally rewarding. Thus the attractiveness of a delivery site depends on both the travel costs to reach it and the profit gained from it (Feillet, Dejax et al. 2005). In this sense, tours are planned in such a way to maximise economic utility.

The hypothesis of this experiment anticipates the interaction of these two effects, and is thus:

- Solution performance for the harvesting variant of the TSP improves with increasing knowledge of future configurations, known as foresight.
- In the case where TSP locations have unequal attractiveness, the significance of foresight is reduced.

The logic behind the first of these points is intuitive; improved capabilities to predict the future and react accordingly are always associated with improved outcomes in many fields of human endeavour. However, the interpretation of the second point is less obvious; it raises the possibility that focussing on the more profitable (and ignoring the less profitable) locations in the current iteration can yield good performance without having to plan ahead.

2. Solution characteristics for the harvesting TSP

In order to examine the effects of foresight and heterogeneity on the harvesting variant of the TSP, a simulated Euclidean space was created and populated by randomly generated locations. Here, locations represent the sites at which harvesting is performed, and are occupied by notional objects representing the harvesting equipment to be relocated at each iteration. The locations are grouped into iterations as shown in fig 1, with the numbers within each circle indicating the iteration. A central depot, indicated by the circle labelled “0” is defined. Tours are defined to start at the depot, then pickup objects from the delivery locations chosen in the last iteration, deliver these to locations chosen in the current iteration and then return to the depot.

In the simulation shown there are four objects to be relocated over three iterations and their starting locations are given by the points labelled “1”. In each iteration (designated by r, where r = {2, 3, 4}), there are seven possible delivery locations for the four objects picked up in that iteration. These locations are labelled according to their r-value. As the vehicle is assumed to have unit capacity, each tour must interleave pickup and delivery locations. That is, in every solution a visit to a chosen delivery location will always be performed directly after a visit to one of the mandatory pickup locations. The temporal linking between iterations occurs because the chosen delivery locations in any particular iteration become the pickup locations in the next iteration.

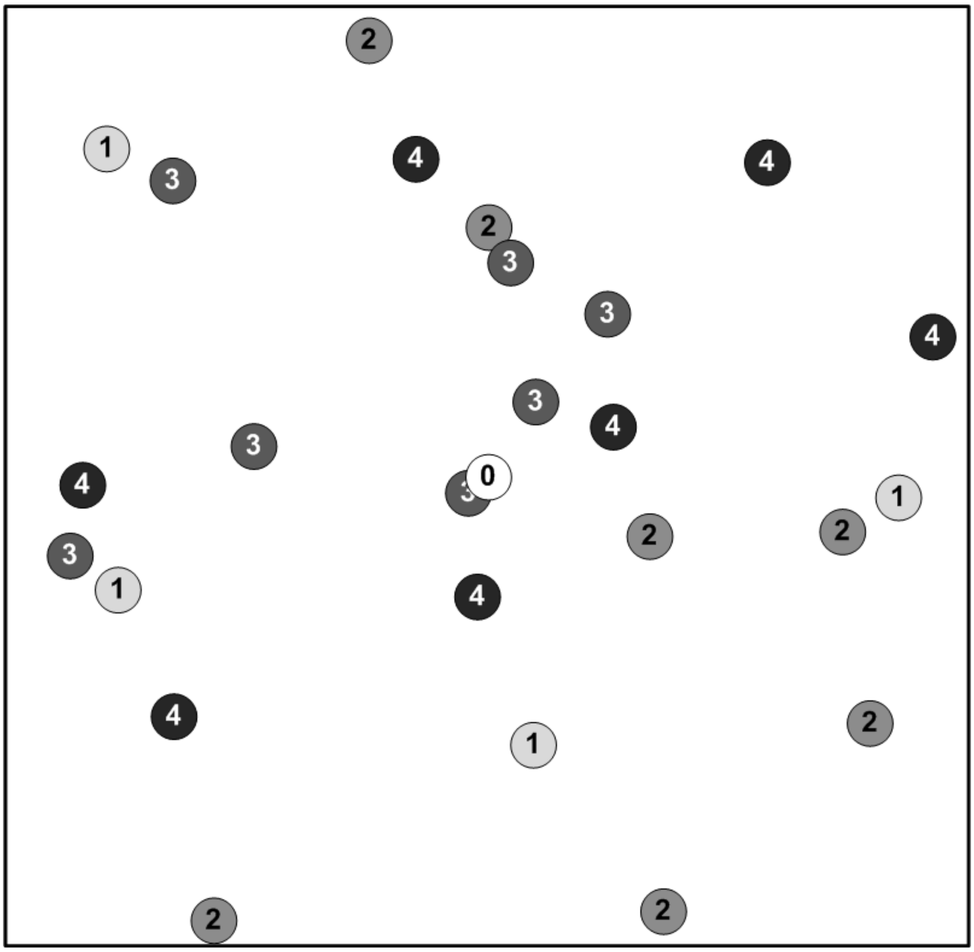

Figure 1: Random location distribution for three iterations of harvesting TSP

There are two categories of optimal tours; iteration optimal solutions and overall optimal solutions. In the iteration optimal solution case, the algorithm seeks the optimal solution for the current iteration only; this then determines the starting locations for the next iteration where the process is repeated. The overall optimal solution algorithm finds the single best solution across all iterations. Thus whilst in the first iteration the overall optimal solution will have a cost greater than or equal to the iteration optimal solution, it is guaranteed to have a cost less than or equal to the iteration optimal solution across all iterations. Iteration optimal and overall optimal solutions to the problem outlined in fig 1 are illustrated below in fig 2.

In the first iteration, the iteration optimal solution yields a path with length 2504 units. One of the delivery locations chosen is a point in the lower left corner of the simulation space. The iteration optimal solution algorithm, lacking any foresight, ignores the fact that this location is a comparatively large distance away from any possible delivery points in the next iteration. However, the overall optimal solution algorithm has complete visibility of all future iterations and the ability to determine the optimal solution across the entire duration of the simulation and thus outperforms the iteration optimal solution, in this case already by the end of the second iteration.

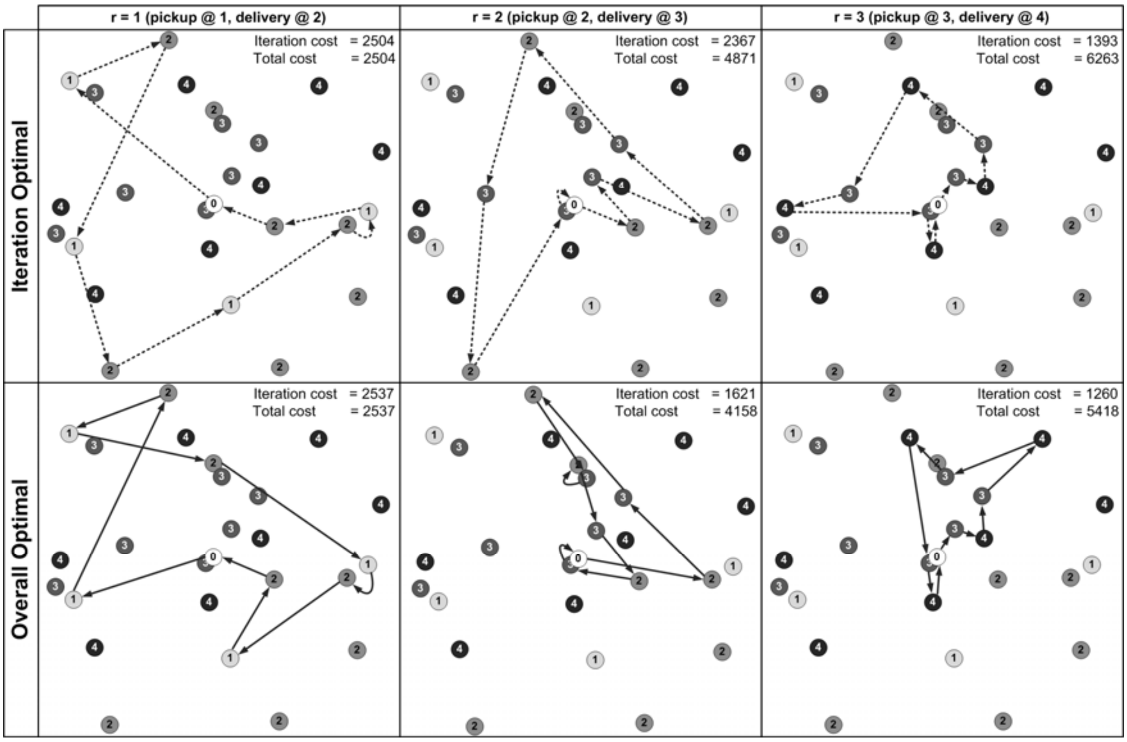

Figure 2: Iteration optimal and overall optimal solutions for a particular random configuration across three iterations

3. Comparison of homogeneous and heterogeneous location attractiveness

To investigate the average behaviour of both strategies, we ran large numbers of random simulations. As shown in table 1 and fig 3, a solution optimised across multiple iterations significantly improves upon that optimised for a particular iteration alone. This result applies both for the cases of homogeneous and heterogeneous location attractiveness. Whilst the magnitude of the performance improvement grows with increasing foresight ability, the rate of growth slows.

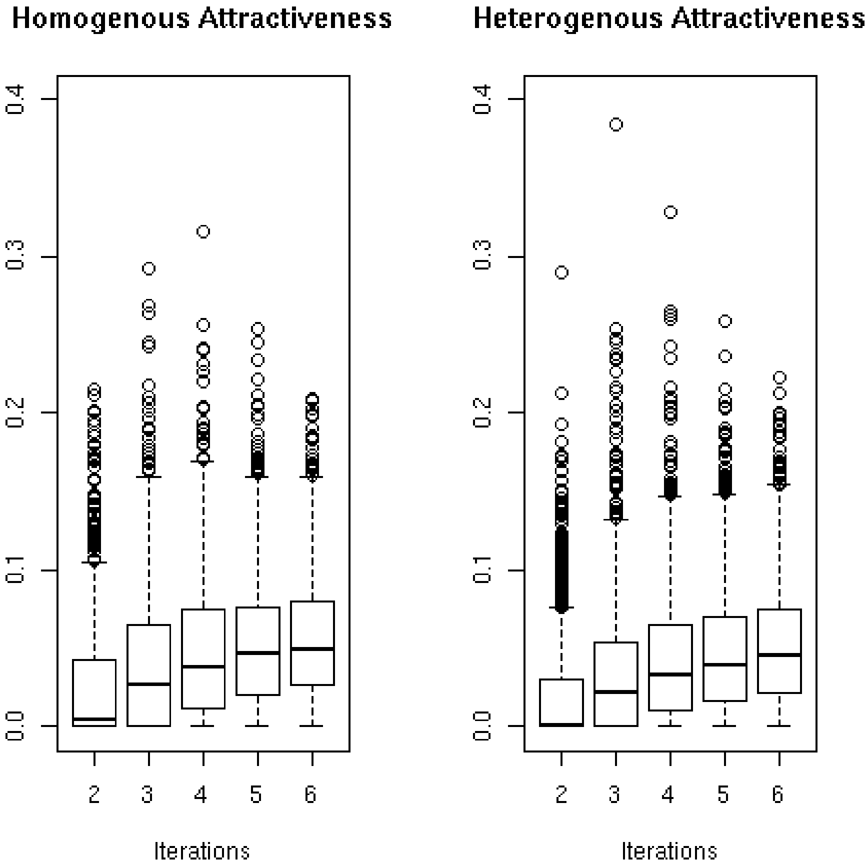

Figure 3: Boxplots for 1000 random distributions and differing numbers of iterations

r	Homogeneous attractiveness mean	Heterogeneous attractiveness mean	Two-sample t-test			Two-sample Wilcoxon test
			p	95% CI		p
2	2.67%	2.22%	0.008	0.12%	0.80%	< 0.001
3	4.09%	3.61%	0.02125	0.07%	0.88%	0.009537
4	4.90%	4.47%	0.03866	0.02%	0.83%	0.03444
5	5.37%	4.74%	<0.001	0.26%	1.00%	< 0.001
6	5.60%	5.17%	0.01392	0.09%	0.79%	0.01018

Table 1. Results for 1000 randomly generated location distributions

The results reject the second hypothesis that when locations are given heterogeneous attractiveness that the performance difference would greatly reduce. A decrease of the mean performance difference was observed for all numbers of iterations, and this decrease was significant at the 5% level. However, as the range of the mean difference between the two samples was less than 1%, the impact of heterogeneous attractiveness was small in absolute terms.

4. Conclusion

It would appear from these results that heterogeneous location attractiveness simply transforms the Euclidean problem space into a non-Euclidean (and non-two-dimensional) space where no fundamentally new efficiencies (such as a focus on the most attractive locations) can be exploited. As heterogeneous location attractiveness is more characteristic of the application domains, this raises a further challenge to determine how humans find good tours as no evidence for the validity of previously investigated heuristics or other information discarding heuristics has yet been found. In particular, techniques that diminish the need to plan ahead whilst still yielding acceptable performance will be a focus for further research.

6. Acknowledgements

Support of the Rural Industries Research and Development Corporation (RIRDC) under project PRJ-002809 is acknowledged.

7. References

Feillet, D., P. Dejax, et al. (2005). "Traveling Salesman Problems with Profits." Transportation Science 39(2): 188-205.

Parragh, S. N., K. F. Doerner, et al. (2008). "A survey on pickup and delivery problems: Part II: Transportation between pickup and delivery locations." Journal fur Betriebswirtschaft 58(2): 81-117.