Planning and optimising organisational travel plans using an Evolutionary Algorithm

Neil Urquhart

Centre for Emergent Computing, Edinburgh Napier University, Edinburgh, UK
n.urquart@napier.ac.uk

Abstract. Commuting to the workplace is a highly individualistic experience, especially where the private car is the chosen mode of transport. The costs of using cars with low occupancy rates are significant in environmental terms as well as requiring the provision of parking space at the workplace. This paper examines the use of an Evolutionary Algorithm based problem solver to construct travel plans for three sites with 248,404 and 520 employees respectively at each site. Results presented suggest that a significant saving in overall distance travelled and parking spaces required is possible. The algorithm employed takes into account both hard constraints and soft constraints (such as work patterns and journey flexibility).

1 Introduction and motivation

Commuting to the workplace by private car is a major contributing factor to motoring related greenhouse gas emissions and to rush hour congestion in towns and cities. The provision of parking spaces can be a major cost to organisations who are based in areas with high land values. By having employees travel to work in groups, sharing a car, the pollution and congestion may be reduced and the number of parking spaces at the workplace also reduced. Planning car sharing requires the identification of groups of employees who live at set of addresses that can be served by one commute that is shorter than the collective individual journeys.

A major drawback to the automated planning of car sharing schemes is their inability to take into account the personal preferences of the individuals participating. The author previously undertook a study into such 'soft factors' in [2]. Some preliminary work on the concept of using an Evolutionary Algorithm to optimise car sharing was presented by the author in [8], this study presents a more in depth discussion of the algorithm and the results obtained. The algorithm used in this study minimises the toal distance travelled by cars, minimises the extent to which individuals deviate from their direct route and ensures that individuals share have similar work patterns.

2 Previous Work

There are many Vehicle Routing Problem (VRP) variants which have been identified and investigated by researchers. An overview of vehicle routing problems and the range of heuristics applied to them may be found in [7]. Disregarding soft preferences, the

problem under discussion may be formulated as a Capacitated Vehicle Routing Problem with Time Windows (CVRPTW). The principle difference being vehicles start from a "customer" and are routed to the central workplace, rather than starting and ending at the same central location. A notable recent approach to the Vehicle Routing Problem with Time Windows (VRPTW) is presented in [4], the problem is formulated as a multi-objective problem, and the EA employed uses ranking to evaluate the population. This work is further developed in [6] where it is applied to the problem of routing garbage collections. The ranking approach has the potential to produce a range of solutions, which collectively form a Pareto front. In this case a strategy is required to determine which solution should be adopted by the user.

There exists several approaches to the problem of car sharing. Naor [3] formulates the problem around groups of employees using their own cars to drive to a intermediate meeting point, and continuing to the place of work in one vehicle. This approach is significantly different to the problem being discussed in this document. Naor examines the possibilities of optimising the sharing of driving equally. The 2nd leg optimised such that from the entire pool each driver does a fair share of 2nd leg driving. Buchholz [1] presents a car-pool problem as an NP complete partitioning problem. The system formulates detours to individuals' journeys to allow them to pick up other individuals. No results are presented, not are the preferences of individual users taken into account.

A games-theory approach to public transport planning and integration is explored in [5]. The authors take a market based approach using games theory to establish how public transport artefacts such as bus services should be provided in order to secure the maximum benefit from them.

3   Problem Description

The problem under consideration here is the construction of a potential car-sharing plan for a UK-based University. Employee address data from the University payroll system may be combined with UK-based Ordnance Survey geo-spatial data (obtained for academic use under the Digimap Agreement) to allow distances between employees homes and work place to be estimated. The aim of any solution is to group employees into groups of up to 4 individuals, who may share the same car. The car is provided by the employee living furthest away who deviates from their journey to pick up other members of the group. In this case the system allows users to specify a work pattern constraint (start and end times of their working day) and degree to which they will deviate from their direct journey in order to share.

The prototype optimises with respect to three objectives:

– The total distance travelled by cars should be minimised.
– The additional distance travelled by any employee should not exceed the deviation constraint specified by that employee.
– Each employee may specify a time slot that they work in, employees should only share with others in that time slot.

4　The planning algorithm

The test bed system constructed utilises an Evolutionary Algorithm (EA) to produce a solution in the form of a travel plan. The solution must divide employees into groups of up to four, each group sharing one car. The solution must take into account the objectives outlined (see section 3). There may exist no ideal solution, but rather the challenge is to find a compromise solution that satisfies as many of the objectives as possible.

The EA uses an indirect representation, each individual is not itself a full solution, but a list of groups, each employee being allocated to one group, the groups are not ordered internally. The number of groups may differ between individuals (but the total number of employees will always remain the same) as solutions may vary the number of passengers in each car. The minimum number of groups being determined by $\frac{noOfEmployees}{4}$), the maximum being $noOfEmployees$ (the case where each individual travels to work in a separate vehicle). The ordering of the groups within the chromosome is not significant. The role of EA is to form such groupings and preserve useful groups through the generational cycle.

The algorithm used maintains a steady-state population of 25 individuals, during each generation a sub-population of 15 children is created. A child may be created from two parents, via recombination, or by cloning a single parent. The probability of recombination is 0.8, parents are selected using a tournament of size 2. Each child has a mutation applied to it, one of the three operators outlined above. The fitness of the child is calculated by building a solution from the encoding within that child as outlined below. Members of the child population are copied into the main population, replacing the looser of a tournament.

Recombination consists of creating a new individual by adding groups selected alternately from each parent. As each group is added employees are removed from the group if they have been added previously as part of a group from the other parent. This encourages the retention of group membership across generations.

The following three mutation functions are employed to modify groups within the selected individual:

- Swap two employees between two groups.
- Create a new group. An existing group is selected (must have a minimum membership of 2 persons) and half its members are moved to the new group.
- Two groups are selected at random and merged (the combined membership of the groups must be equal to or less than 4 persons)

In order to evaluate the fitness of an individual a complete solution must be constructed an evaluated. A solution is constructed by applying an ordering heuristic to each employee group, to order them by distance from the workplace.

A penalty fitness value is then calculated for each group based on the following:

$$gFit = d + (d * devP) + (d * tpP - 1)$$

where:

gFit = the fitness value for the current group of employees
d = total distance driven by the group in meters
devP = 1 if any of the individuals in the group have a journey length that violates their deviation constraint
tpP = the number of different timeSlots within the group

The total fitness value for a candidate solution is the sum of the group fitness values, the more constraints that are broken within a solution the higher the fitness allocated to it. The algorithm is executed until 100 generations have elapsed without any fitness improvement.

5 Experimental method and Results

The test bed system has been tested using data based on the payroll of a UK based University. Three university campuses were examined; the numbers employed at each site may be seen in table 1.

Location	Employees	Average direct distance to workplace(km)
Site 3	404	6.2
Site 2	248	6.5
Site 1	520	6.0

Table 1. The datasets used in this investigation.

For testing purposes each individual is allocated a random work pattern identifier in the range (1..4) to represent their work pattern and a random deviation value in the range (10-90%).

The reader should consider that this problem may be presented as a "design" type problem, where sufficient time is available to allow multiple runs to be made in order to take account of the stochastic nature of the algorithm. However in many cases it may be necessary to produce travel plans quickly (in order to respond to users' changing requirements) in which case there may only be time in which to build one solution. With this in mind, the results presented here represent the average over 10 runs.

Over the following series of runs the deviation constraint was altered through values of 10,30,40, 70 and 90% for each user. For instance if a user has a direct journey distance of 10 kilometres to their place of work, and a deviation constraint of 30% then the maximum acceptable journey distance to them when participating in the car sharing scheme would be 13 kilometres. The results obtained may be seen in table 2. Note how as the deviation constraint is relaxed the other constraints are met, this ability to "trade off" conflicting constraints is a well known feature of evolutionary algorithms.

From an environmental perspective, it is interesting to note the average occupancy of the cars arriving at the workplace. An optimal solution would present an occupancy

of 4, table 3 presents solutions close to this (an average of 3.9 and 3.8) when only 1 or two timeslot constraints exist.

Over the following series of the number of work patterns available was increased from 1 to 4 and results obtained may be seen in table 3. T-tests have been used in order to establish that varying the deviation constraint and the quantity of work patterns does result in statistically significant changes in results. Table 4 compares the fitness of the pupations produced with a deviation constraint of 10% with the results achieved with constraints of 50% and 90%. Table 5 makes a similar comparison between results obtained with only one work pattern and with 4 work patterns.

	Deviation (%)	Site 3	Site 2	Site 1
% Dist Saved	10	39.9	42.3	35.1
	30	44.9	45.3	39.7
	50	47.1	48.3	41.3
	70	48.7	48.5	43
	90	48.6	50.9	42.7
Parking spaces	10	44.9	39.4	53.3
	30	36.3	34.1	43.5
	50	32.7	31.8	39.5
	70	30.6	32.1	36.1
	90	30.61	29.7	35
Deviation Constraint Violations	10	4.8	4.9	4.3
	30	0.8	0	1.4
	50	0.5	0	0.6
	70	0	0.4	0.4
	90	0	0	0.1
Time slot Constraint Violations	10	0.4	0	0.6
	30	0.3	0	1.0
	50	0.3	0.1	0.8
	70	0.3	0.1	0.9
	90	0.3	0	1.2
Average Car Occupancy	10	2.2	2.5	1.9
	30	2.8	2.9	2.3
	50	3.1	3.1	2.5
	70	3.3	3.1	2.8
	90	3.3	3.4	2.9

Table 2. Results obtained whilst altering the deviation constraint.

6 Conclusions and future work

From the results presented it may be seen that total commuter millage is reduced by 50%, and on average less than 30% of employees actually have to park at work. Given the constraint of limiting individuals to a maximum of 4 persons per vehicle, the algo-

	Work Patterns	Site 3	Site 2	Site 1
Distance saved	1	64.9	65.1	64.0
	2	58.1	59.7	54.6
	3	55.7	55.4	49.7
	4	51.6	52.5	46.5
Parking spaces	1	25.4	25.5	25.4
	2	26.3	26.5	26.4
	3	27.0	27.3	27.0
	4	27.6	28.2	30.0
Deviation Constraint Violations	1	0	0	0
	2	0	0	0
	3	0	0	0
	4	0	0	0
Time slot Constraint Violations	1	0	0	0
	2	0.2	0	0
	3	0	0	0.4
	4	0	0	1.2
Average Car Occupancy	1	3.9	3.9	3.9
	2	3.8	3.8	3.8
	3	3.7	3.7	3.7
	4	3.6	3.5	3.3

Table 3. Results obtained whilst altering the number of work pattern.

Comparison (Deviation)	Site 3	Site 2	Site 1
10% - 50%	0.0001	0.0001	0.0001
10% - 90%	0.0131	0.0001	0.0001

Table 4. T-Test results comparing the fitness of the individuals that comprise the final populations produced with the deviation constraint set at 10% and 50% and then between 10% and 90%. The values returned suggest that the fitness produced with 1 and 4 work patterns are statistically significant.

Comparison (Groups)	Site 3	Site 2	Site 1
1 - 4	0.0001	0.0001	0.0001

Table 5. T-Test results comparing the fitness of the individuals that comprise the final populations produced with the work pattern variable set at 1 and 4. The values returned suggest that the fitness produced with 1 and 4 work patterns are statistically significant.

rithm manages to reduce the number of cars at the workplace to less than 1% more than the 25% minimum.

In every case individuals' desires for deviation distance were met and only in a few cases at the largest campus were some individuals not placed in groups compatible with their timeslot. This system produces a plan within approximately 10 minutes of CPU time, although this time will differ depending on hardware , software implementation and on the dataset being used. Future work, involves allowing a wider range of user variables to be taken into account, the constraining nature of such variables and their potentially random nature should create a search space that may be successfully explore using the evolutionary algorithm. The survey work undertaken in [2] suggests that there is potential for modelling of soft constraints. This may be achieved by allowing users to feedback into the system their satisfaction level with the arrangements proposed. Such feedback would allow individuals to be allocated reputational scores indicating their tolerance of sharing. It would be possible to build up a graph structure of employees with weighted arcs indicating that individuals have shared previously and the success of that share.

References

1. F. Buchholz. The carpool problem. Technical report, Institute of Computer Science, University of Stuttgard., 1997.
2. R. Holden N. Urquhart, T. McEwan and T. Vogogias. Co2y: the intelligent green solution: minimising carbon emissions by maximising shared travel opportunity. Scottish Transport Applications Research Conference, Glasgow, UK, 2009.
3. M. Naor. On fairness in the carpool problem. *Journal of Algorithms*, 55(1):93–98, 2005.
4. Beatrice M. Ombuki, Brian Ross, and Franklin Hanshar. Multi-objective genetic algorithms for vehicle routing problem with time windows. *Appl. Intell.*, 24(1):17–30, 2006.
5. Athena Roumboutsos and Seraphim Kapros. A game theory approach to urban public transport integration policy. *Transport Policy*, 15(4):209 – 215, 2008.
6. Andrew Runka, Beatrice M. Ombuki-Berman, and Mario Ventresca. A search space analysis for the waste collection vehicle routing problem with time windows. In Franz Rothlauf, editor, *GECCO*, pages 1813–1814. ACM, 2009.
7. Paolo Toth and Daniele Vigo. *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics, 2002.
8. Neil Urquhart. Carbon-friendly travel plan construction using an evolutionary algorithm. In Hod Lipson, editor, *GECCO*, page 2269. ACM, 2007.