ON THE PERFORMANCE OF POPULATION-BASED METAHEURISTICS FOR THE SPACE ALLOCATION PROBLEM: AN EXTENDED ABSTRACT

E. K. Burke, P. Cowling, J.D. Landa Silva, S, Petrovic
Automated Scheduling, Optimisation and Planning Research Group
School of Computer Science and IT, Jubilee Campus, University of Nottingham
Email: {ekb , pic , jds , sxp}@cs.nott.ac.uk
http://www.asap.cs.nott.ac.uk

1 The Problem

In many real-world optimisation problems, we need to find a set of near-optimal solutions from which one can be selected [8][9]. For instance, multiobjective optimisation problems form an extensive class of problems where multiple criteria need to be satisfied. The multiple objectives are often conflicting and/or incommensurable and the goal is to find a solution which represents a set of good compromises between the objectives. Moreover, in some single-objective optimisation problems, it might be necessary to find several good local optima so that a decision maker may make the final choice.

The space allocation problem is a constrained multiobjective combinatorial optimisation problem related to the bin-packing and the knapsack problems [10]. It refers to the distribution of n objects of different sizes into m areas of space with different capacities. There are hard constraints that must be satisfied so that the solution is feasible and soft constraints that it is desirable to accomplish. The goal is to find a distribution of all objects into the areas of space that optimises the space utilisation, satisfies all hard constraints and satisfies as many soft constraints as possible. Sometimes, the conditions (constraints, number/size of objects, number/capacity of areas of space) of an existing allocation are changed and a reorganisation of the distribution is required. In this case, there is an additional objective: minimise the disturbance due to reorganisation [3].

The space allocation problem in academic institutions refers to the distribution of people (staff, research students, etc.) and other resources (laboratories, lecture rooms, etc.) to physical rooms. There are several types of constraints such as: proximity/adjacency, grouping and sharing constraints. Proximity/adjacency constraints exist when one person needs to be close to certain rooms or to another person. Grouping constraints refer to a group of people that need to be allocated close each other. Sharing constraints refer to the situation in which people should not share a room with another resource. In the real instances of this problem, there is an additional goal that is very difficult to accomplish and to evaluate: how satisfied is each person with the assigned room? When selecting a solution for a space allocation problem, we must consider multiple criteria from more than one decision-maker. It is desirable to select from among several possible solutions [5].

2 Solving the Problem

Our research is focused on solving the space allocation problem (not only finding a solution to a specific instance) using metaheuristics and in particular, multiobjective optimisation techniques. We have investigated the application of hill-climbing and some metaheuristics such as simulated annealing, tabu search and genetic algorithms to the space allocation problem in universities [4]. Recently we proposed a hybridisation of some of these methods that offers good solutions when constructing a new allocation for this problem [2]. Given the multiobjective nature of the space allocation problem, we incorporated a population of solutions into the hybrid metaheuristic [1]. The proposed approach is based on evolving an initial set of feasible solutions using self-adaptation and information-sharing between the individuals. This approach produces a set of good solutions that represents a compromise between the two objectives: minimisation of space misuse and minimisation of the violation of soft constraints. By preserving multiple possible solutions, the decision-maker can then chose a solution which gives the greatest possible satisfaction to the space users. In our experiments, it has been observed that the technique is capable of improving the initially constructed set of solutions while maintaining diversity in the population. In most of the cases, the quality of the solutions found by the population-based hybrid metaheuristic is better than the quality of the solutions produced by re-starting the single-solution version of the technique. In our experiments, we have used an aggregating function to combine the objectives into a single scalar value and a dominance relation to produce a set of non-dominated solutions. Several issues need more investigation in order to establish further conclusions about the effectiveness and efficiency of our approach for multiobjective optimisation.

3 Neighbourhood Structures

We use a local search heuristic that explores the neighbourhood using five different moves: allocate, remove, relocate, interchange and swap. Allocate is the allocation of a resource to a room. Remove is the removal of a resource from the assigned room. Relocate is the change in the assigned room for a resource. Interchange refers to the exchange of assigned rooms between two resources. Swap moves all resources allocated into a room to another room and vice versa.

It has been observed that during the evolution process of an allocation, the type and number of successful move attempts varies. The use of different neighbourhood structures in the various stages of the solving process might provide better results. This is particularly important in reorganisation problems where the aim is to provide feasible solutions that incorporate the requested changes, with as little disturbance to the allocation as possible. Defining different neighbourhood structures for reorganising an allocation may help solve other real-world combinatorial optimisation problems in which it is required to modify an existing solution [7]. In the full paper, we investigate the application of Variable Neighbourhood Search (VNS) proposed by Mladenovic and Hansen [12] to solve the space allocation problem.

4 Validating the Hybridisation

Adding a population of solutions that share information during the evolution process with no recombination, improved our initial hybrid metaheuristic. What is the real effect of the information-sharing mechanism? What are the real advantages of having a population over re-starting the hybrid metaheuristic? We will provide results that permit us to clarify these issues. Comparing the performance of a given single-solution technique with its variant including a population with a cooperation mechanism, allows us to establish whether this idea can be extended to other metaheuristics. Using a population of solutions in combination with variable neighbourhood search has a potential application to multiobjective optimisation. The various neighbourhood structures can be used by different members of the population, according to the various objectives.

Some researchers have found that evolutionary algorithms and in particular genetic algorithms require a large population size and number of generations to operate [11]. It is also true that these algorithms have to cope with two common problems: keeping diversity in the population and avoiding premature convergence [13]. The population size used in our approach can be quite small without decreasing the effectiveness of the technique. In addition, a considerable amount of improvement is achieved in the early iterations of the algorithm.

In some multiobjective optimisation problems, we are required to find the Pareto-optimal front or a near-Pareto-optimal front [6]. In this case it is not only the diversity of the resulting set of solutions that is important, but also the spread of the population over the desired front. Diversity means that a given set of solutions should not be clustered together. Spread refers to the portion of the front that is covered by the set of solutions. Diversity and spread can be defined in relation to the solution space or in relation to the objectives space. That is, do we require solutions that represent very different allocations, solutions that have different objective values or both simultaneously? Another measure to assess the effectiveness of a multiobjective optimisation technique is how far the obtained non-dominated set is from the Pareto-optimal front? We found that our population-based metaheuristic is capable of maintaining a reasonable diversity in both the solution space and the objectives space. Additionally, we present results using a genetic algorithm in our experiments. A comparison between this genetic algorithm and the other population-based approaches in all aspects mentioned here will be presented in the full paper.

5 Preliminary Results

In this section, we present some early experimental results for a medium-size space allocation problem. In this test problem, there are 77 human resources, 77 rooms and 30 constraints. A feasible solution must have all resources allocated and there are the two objectives: the minimisation of space misuse (wastage and overuse) and the minimisation of constraint violations.

Using our hybrid metaheuristic, we executed 30 runs distributed as follows: 10 runs minimising space misuse, 10 runs minimising constraint violations and 10 runs minimising both metrics simultaneously using an aggregating function. Details of the data used in our experiments and the evaluation function can be found in [3]. A CPU execution time of 120 seconds was given to each run.

During the improvement phase in our algorithm, three of the five moves are used to explore the neighborhood: relocate, interchange and swap. The success of each type of move to improve the solution varies according to the objective that is being optimised. We measured the contribution of each move to the total number of moves that produced solution improvements in our experiments. In Table 1, we observe that the contribution of the relocate and interchange moves is very similar in the minimisation of space misuse. When minimising constraint violations, the interchange move contributes considerably more than the relocate move. If both objectives are combined using the aggregating function, the contribution of these two moves is reduced. In all three cases, the most successful move is the swap operation. If minimising constraint violations, most of the improvement is produced by the interchange and swap moves. In the other two cases, the swap move contributes with almost 90% on the successful operations. Further analysis of the spread rate and the quality of the obtained front with respect to the Pareto-optimal front is discussed in the full paper. Similar experiments have been carried out for the other techniques such as hill-climbing, simulated annealing and tabu search.

	minimising space misuse	minimising constraint violation	aggregating function
Relocate	9.40	3.25	1.71
Interchange	7.45	30.42	8.71
Swap	83.15	66.33	89.58
total	100	100	100

Table 1. Average percentage of success for each type of move.

The disruption produced by these moves to the existing allocation is different in each case. Relocate is the operation that disrupts the least, while swap in the one that disrupts the most. The use of different neighborhood structures according to the type of problem and the objective being optimised appears to be useful in this multiobjective combinatorial optimisation problem.

Additional experiments were conducted to compare the performance of a given metaheuristic using the population-based and the single-solution variants. Using different population sizes, our hybrid population-based metaheuristic was executed using the test problem described before. For each population-size, the single-solution variant of the metaheuristic was re-started to produce the same number of solutions. Table 2 presents the statistics for the set of solutions produced in these experiments. The diversity metric measures the difference between allocations within a specific set or population. The higher percentage value, the more diverse the set of solutions is. A value of 100% would indicate a set of solutions in which all the allocations are different, i.e. no resource has been allocated to the same room in any two different solutions. We observe that with both population-sizes, the population-based metaheuristic outperforms the multi-start single-solution approach in terms of the solution fitness maintaining a diverse population too.

	population size = 5		population size = 20	
	spmh ps=5	pbmh ps=5	ssmh ps=20	pbmh ps=20
Best Fitness	7687.09	3566.93	5457.68	3566.93
Worst Fitness	10030.79	9123.49	15408.93	11151.12
Average Fitness	8662.21	6795.61	9226.98	7578.04
Diversity Percentage	77.92	72.73	56.04	54.50

Table 2. Comparison of the single-solution and population-based variants with different population sizes.

6 Conclusions

We do not intend to establish our approaches as an overall winner, because as is stated by some researchers in the community [9], it is not clear which technique is the best for multiobjective optimisation. However, with our experiments and results, we aim to achieve four goals. To validate the effectiveness/efficiency of our approach for multiobjective optimisation. To extend the idea of population-based variants to other known single-solution methods. To offer an insight into the application of variable neighborhood search in multiobjective optimisation. Finally, we intend to present a fair comparison between different techniques to solve the multiobjective space allocation problem.

Bibliography

[1] Burke E.K., Cowling P., Landa Silva J.D., Petrovic S., Combining Hybrid Metaheuristics and Populations for the Multiobjective Optimisation of Space Allocation Problems, accepted for publication in the Proceedings of the GECCO 2001, Genetic and Evolutionary Computation Conference 2001, San Francisco, USA, 7-11 July 2001.

[2] Burke E.K, Cowling P., Landa Silva J.D., Hybrid Population-Based Metaheuristic Approaches for the Space Allocation Problem, accepted for publication in the Proceedings of the CEC2001, Congress on Evolutionary Computation 2001, Seoul, Korea, 27-30 May 2001.

[3] Burke E.K., Landa Silva J.D., The Space Allocation Problem, [Online], Available at: http://www.asap.cs.nott.ac.uk/ASAP/space/spaceallocation.html, [2001, March 1st].

[4] Burke E.K., Cowling P., Landa Silva J.D., McCollum B., Three Methods to Automate the Space Allocation Process in UK Universities, Proceedings of the 3rd International Conference on the Practice and Theory of Automated Timetabling, PATAT 2000, Konstanz, Germany, pp. 374-393, 2000.

[5] Burke E.K., Varley D.B., Space Allocation: An Analysis of Higher Education Requirements, The Practice and of Automated Timetabling II: Selected Papers from the 2nd International Conference on the Practice and Theory of Autometed Timetabling, PATAT 97, University of Toronto, 20-22 August 1997, Lecture Notes in Computer Science, Vol. 1408, pp. 20-33, Springer-Verlag 1998.

[6] Coello Coello C.A., A Short Tutorial on Evolutionary Multiobjective Optimization, Proceedings of the 1st International Conference on Evolutionary Multi-Criterion Optimization, EMO 2001, Lecture Notes in Computer Science Vol. 1993, pp.21-40, Springer, Zurich, Switzerland, 2001.

[7] Cumming A., Paechter B., Ranking R.C., Post-Publication Timetabling, Proceedings of the 3rd International Conference on the Practice and Theory of Automated Timetabling, PATAT 2000, Konstanz, Germany, pp. 107-108, 2000.

[8] Dasgupta P., Chakrabarti P.P., Desarkar S.C., Multiobjective Heuristic Search: An introduction to Intelligent Search Methods for Multicriteria Optimisation, Computational Intelligence - Vieweg, 1999.

[9] Horn J., Multicriteria Decision Making and Evolutionary Computation in T. Bäck, D.B. Fogel, Z. Michalewicz (Eds.), Handbook of Evolutionary Computation, Institute of Physics Publishing, Bristol, UK, 1997.

[10] Martello S., Toth P., Knapsack Problems, Algorithms and Computer Implementations, Wiley-Intersciense Series in Mathematics and Optimisation, 1990.

[11] Man K.F., Tang K.S., Kwong S., Genetic Algorithms, Springer, 1999.

[12] Mladenovic N., Hansen P., Variable Neighbourhood Search, Computers Operations Research, Volume 24, No. 11, pp. 1097-1100, 1997.

[13] Zitzler E., Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications, Thesis submitted to the Swiss Federal Institute of Technology Zurich, Shaker Verlag, Germany, 1999.