Evolution of Fitness Functions to Enhance Heuristics

Peter Cowling[1], Keshav Dahal[1], Stephen Remde[1] and Nic Colledge[1]

1. MOSAIC Research Group, University of Bradford, Bradford, BD7 1DP, United Kingdom
{p.i.cowling, k.p.dahal, s.m.remde, n.j.colledge}@bradford.ac.uk
http://mosaic.ac/

Keywords : Heuristics, Mutiobjective Optimisation, Variable Fitness Function, Workforce Scheduling, Travelling Salesman Problem, The Virus Game.

1 Introduction

Decision support systems which model real-world decision problems as combinatorial optimisation problems are increasingly used to increase productivity, improve service quality, and reduce cost and environmental impact in industrial, commercial and government applications. When modelling a real-world decision problem as a problem of combinatorial optimisation, it is usually assumed that that there is a single underlying objective (fitness) measure to allow automatic comparison between candidate solutions. In real problems, this objective measure is almost always a function of several underlying sub-objectives relating to revenue, cost, staff and customer satisfaction, sustainability etc. [1, 2], and solution heuristics are often highly tailored to deal with complex problem-specific decision rules. In commercial computerised decision support systems a weight is usually assigned to each sub-objective to reflect its relative importance, and the objective consists of a weighted sum of sub-objectives [2]. Multiobjective approaches [3] offer an interesting way forward that does not require the user to consider weights directly, although the issue of effectively comparing many different Pareto optimal solutions have prevented their application in problem areas where understanding of a single solution is already a significant challenge for the user [4, 5].

The weighted sum objective reflects the relative importance of sub-objectives in a "finished" solution. This is not an issue for exact approaches which guarantee to find an optimal solution (e.g. [6]), but since heuristic approaches are usually used for hard combinatorial optimisation problems [7], the objective weights for a "finished" solution may be a poor reflection of the objective weights which would provide the best local decision for a heuristic, working on a partial or poor-quality solution. Designing an objective function which varies from iteration to iteration, and takes account of the features of the current solution and the heuristic method(s) used to solve it appears to be a very difficult task. In this talk we will use Variable Fitness Function (VFF) methodologies which can provide an effective way of using additional CPU time (which is much cheaper than developer time) to find a way of varying the objective weights of a weighted sum objective function, tailored to the problem and the heuristic used to solve it.

2 Variable Fitness Functions (VFFs) and Related Approaches

A Variable Fitness Function (VFF) varies the weights of a weighted sum fitness function over the iterations of a search heuristic. In Fig. 1 we give an example of a piecewise linear VFF encoded using a Genetic Algorithm chromosome, although other representations are possible. Here the VFF representation simply encodes the value of the weight associated with each objective at certain iterations of heuristic search, and we linearly interpolate to obtain the weights at intermediate iterations.

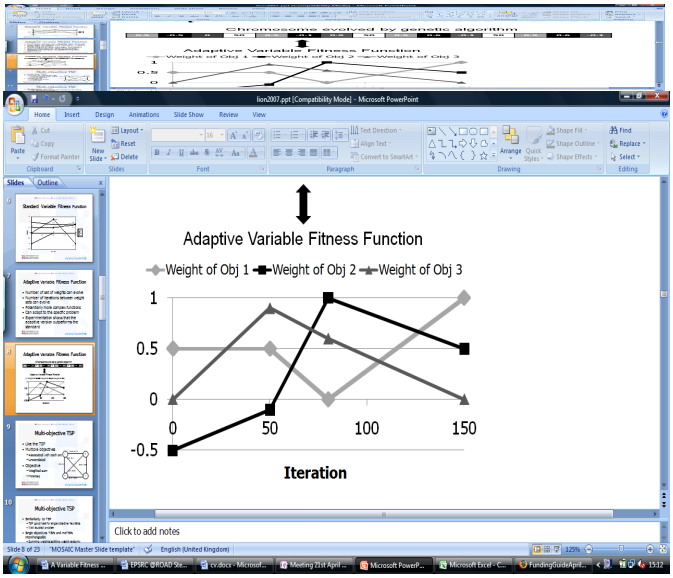

Fig. 1. Chromosome representation for a piecewise linear variable fitness function

Preliminary work [8, 9] shows that the VFF greatly outperforms simple random and multi-start approaches given the same amount of CPU time, and that VFFs evolved for a set of "training" instances yield instant, significant, improvements on an unseen "test" set of problems. VFFs in this preliminary work seem to have "discovered" rescheduling methods such as right-left shift, with many very effective VFFs having a somewhat "weird" form which would be very hard for a human to design.

3 Results

We discuss the VFF approach applied to three types of problem, a multiobjective travelling Salesman Problem (TSP), the board evaluation function of a 2-player board game (the Virus Game) and a very complex and messy real work workforce scheduling problem, achieving very good results in each case, and demonstrating that the VFF has "learned" useful information about the problem and the heuristics used in its solution.

4 References

[1] Viana, A., de Sousa, JP. (2000) Using metaheuristics in multiobjective resource constrained project scheduling. *Eur. J. Op. Res.* 120(2): 359-374.

[2] Thiagarajan S., Rajendran C. (2005) Scheduling in dynamic assembly job-shops to minimize the sum of weighted earliness, weighted tardiness and weighted flowtime of jobs, *Comp. Ind. Engrg.* 49(4): 463-503.

[3] Deb, K.: Multi-Objective Optimisation using Evolutionary Algorithms. Wiley, 2004.

[4] An Architecture for Exploring Large Design Spaces (1998) *in proc Nat. Conf. AAAI*, 143-150.

[5] Shackelford, M., and Corne, D. (2001) Collaborative Evolutionary Multi-Project Resource Scheduling. *in proc. IEEE CEC 2001*, 1131-1138.

[6] Albiach J, Sanchis JM, Soler D (2008) An asymmetric TSP with time windows and with time-dependent travel times and costs: An exact solution through a graph transformation, *Eur. J. Op. Res.* 189(3): 789-802.

[7] Kolisch, R., and Hartmann, S. (2006) Experimental investigation of heuristics for resource-constrained project scheduling: An update. *Eur. J. Op. Res.* 174 (1): 23-37.

[8] Remde, S., Cowling, P., Dahal, K. and Colledge, N. (2008) Evolution of Fitness Functions to Improve Heuristic Performance. to appear *in proc. LION II*, Springer LNCS.

[9] Dahal, K. , Remde, S., Cowling, P., and Colledge, N. (2008) Improving Metaheuristic Performance by Evolving Variable Fitness Functions, in *proc EvoCOP 2008*, Springer LNCS 4972, 70-181.