Experiments with a form of double iterated search for use on hard combinatorial problems with many objectives

Mike B. Wright

Department of Management Science
Lancaster University, LA1 4YX, UK
m.wright@lancaster.ac.uk

1 Introduction

1.1 Iterated Local Search

Iterated Local Search (Lourenço et al., 2003), henceforth to be referred to as ILS, is a metaheuristic that has been used successfully to solve many combinatorial optimisation problems, producing competitive results. The technique essentially consists of two phases: a local improvement phase which leads monotonically to a local optimum and another very short phase which may take one of a number of forms, often involving a small number of moves chosen partly or wholly at random. The technique then iterates between these phases until some stopping criterion is satisfied.

1.2 Problems with many objectives

Problems involving several distinct objectives are often formulated so that there is a single objective function formed by a linear combination of subcosts, each subcost relating to one of the many objectives. However, good metaheuristic methods for the solution of such problems can prove very difficult to achieve, partly because of the often highly complex nature of the solution space. Some metaheuristic techniques aim to address this issue by modifying the weights of the objectives at various times during the search, thus reshaping the new solution space. Such methods include Noising (Charon and Hudry, 1993) and SAWing (Eiben and van Hemert, 1999), which have achieved some successes.

1.3 This paper

This paper reports the results of experiments of a new technique which combines both of the ideas outlined above. The problem addressed is a sports rostering problem.

E. K. Burke, H. Rudová (Eds.): PATAT 2006, pp. 506–508. ISBN 80-210-3726-1.

2 A cricket umpire rostering problem

The problem addressed by this work is that of rostering cricket umpires for the Devon Cricket League in England. This has been solved for several years using a form of Simulated Annealing (SA). However, the problem is a useful test bed for new ideas, since it is sufficiently large and complex (with 14 separate objectives) to be an interesting challenge, yet not so large as to make experimentation excessively time-consuming. A full description of the problem has already been published (Wright, 2006).

3 The double ILS technique

Using intuitive methods for forming an initial solution and defining neighbourhoods, the solution method continues using a form of Double ILS, since there are two loops which use ILS in different ways and for different purposes. It proceeds according to the following pseudo-code.

```
repeat
  set counter = 0
  repeat
    undertake Local Improvement (LI) to a local optimum
    set counter = counter + 1
    if counter < N then make X perturbations
  until counter = N
  change subcost weights
until total number of iterations during LI phases > Z
return to best solution found, reset weights to
original values and carry out final LI
```

The inner loop is a simple form of ILS, while the outer loop is a very different form of ILS (suitable only for problems with many objectives). This is why the technique has been named "Double ILS".

An iteration is defined as the process of calculating the change in cost of a perturbation, whether or not that perturbation is accepted as a result.

The "best" solution means the best using the original weights.

Specifying Z ensures that fair comparisons between experimental runs can be made, since the value of Z effectively indicates the computational effort.

3.1 The Local Improvement phase

The LI technique systematically searches through all possible perturbations, accepting any solution found that has lower cost than the current solution. A device is used to identify perturbations that could not possibly improve the solution; such perturbations are not made, hence reducing the time taken.

3.2 Changing weights

There are two input parameters (H and L). Initially the weights used are the "real" weights and the "control parameter" C is set equal to H. When it is required to change weights, first C is set equal to H – ((H – L) * I) / Z, where I is the number of iterations to that point, and then the weights are changed as follows:

Method 1: for each subobjective separately, if its cost is greater than or equal to its value the previous time weights were changed, multiply its weight by R (a random number between 1 and C); otherwise divide its weight by R.

Method 2: decide at random (probability 0.5) whether to multiply or divide the weight by R.

3.3 Experimental results

Experiments to date have N = 1 (and thus X is immaterial) and Z = 500,000 or 2,000,000. They show that Method 1 outperforms Method 2 to a small but statistically significant extent; that values of H between 4 and 8, and L between 2 and 4, appear to work best; that the method is already fairly close to being competitive with SA; and that it is considerably better than repeated LI.

Full results will be presented at the conference for various values of N and X, and for different ways of choosing the X perturbations, including totally at random. The overall aim is to find robust values for H, L, N and X which obviate the need for tuning parameters anew for every application.

References

1. Charon, I. and Hudry, O. (1993) The Noising method: a new combinatorial optimization method. Operations Research Letters 14(3): 133-137.
2. Eiben, G. and van Hemert, J. (1999) SAW-ing EAs: Adapting the fitness function for solving constrained problems. In Corne, D. et al., editors, New Ideas in Optimization, pages 389-402. McGraw-Hill, London.
3. Lourenço, H.R., Martin, O.C. and Stützle, T. (2003) Iterated Local Search. In Glover, F. and Kochenberger, G.A., editors, Handbook of Metaheuristics, pages 321-353. Kluwer Academic Publishers, Massachusetts, USA.
4. Wright, M.B. (2006) Case study – problem formulation and solution for a real-world sports scheduling problem. To appear in the Journal of the Operational Research Society.