Multiobjective Optimization on a Budget of 250 Evaluations

Joshua Knowles[1] and Evan J. Hughes[2]

[1] School of Chemistry, University of Manchester, Faraday Building, Sackville Street, PO Box 88, Manchester M60 1QD, UK
j.knowles@manchester.ac.uk
[2] Cranfield University, Shrivenham, Swindon SN6 8LA, UK

Abstract. In engineering and other 'real-world' applications, multiobjective optimization problems must frequently be tackled on a tight evaluation budget — tens or hundreds of function evaluations, rather than thousands. In this paper, we investigate two algorithms that use advanced initialization and search strategies to operate better under these conditions. The first algorithm, Bin_MSOPS, uses a binary search tree to divide up the decision space, and tries to sample from the largest empty regions near 'fit' solutions. The second algorithm, ParEGO, begins with solutions in a latin hypercube and updates a *Gaussian processes* surrogate model of the search landscape after every function evaluation, which it uses to estimate the solution of largest expected improvement. The two algorithms are tested using a benchmark suite of nine functions of two and three objectives — on a budget of only 250 function evaluations each, in total. Results indicate that the two algorithms search the space in very different ways and this can be used to understand performance differences. Both algorithms perform well but ParEGO comes out on top in seven of the nine test cases after 100 function evaluations, and on six after the first 250 evaluations.

Keywords: multiobjective optimization, expensive black-box functions, ParEGO, DACE, Bin_MSOPS, landscape approximation, response surfaces, test suites.

1 Introduction

The vast majority of research effort in developing modern multiobjective evolutionary algorithms (MOEAs) has concentrated on improving algorithm performance and efficiency on runs, typically, of ten thousand function evaluations or more. In this paper, we consider multiobjective problems where a 'budget' of at most 250 evaluations is imposed because of the expensive nature of evaluating candidate solutions. More specifically, we are interested in problems where most or all of the features described in Fig.1 are true.

Features 1–4 limit the numbers of function evaluations possible, while features 5–8 make it reasonable to apply global search techniques rather than either random search or hillclimbing. Problems exhibiting these features include various combinatorial biochemistry and materials science applications [6, 26], as well as instrument set-up optimization in analytical chemistry [20, 24]. In [20], a standard MOEA, PESA-II, was successfully used to substantially improve the settings of a GC-MS spectrometer, using

C. A. Coello Coello et al. (Eds.): EMO 2005, LNCS 3410, pp. 176–190, 2005.

1. the time taken to perform one evaluation is of the order of minutes or hours,
2. only one evaluation can be performed at one time (no parallelism is possible),
3. the total number of evaluations to be performed is limited by financial considerations,
4. no realistic simulator or other method of approximating the full evaluation is readily available,
5. noise is low (repeated evaluations yield very similar results),
6. the overall gains in quality (or reductions in cost) that can be achieved are high,
7. the search landscape is multimodal but not highly rugged,
8. the dimensionality of the search space is low-to-medium,
9. the problem has multiple, possibly incommensurable, objectives.

Fig. 1. Features exhibited by problems of interest

just 180 evaluations. However, it is clear that given such a restricted number of evaluations, and no particular restriction on computational overhead (since each experiment requires 20 minutes), a search strategy that more carefully considers each evaluation would be more appropriate.

Scanning the optimization literature reveals that a sparse but varied array of different techniques (that were proposed or could be used) for economizing on evaluations in multiobjective optimization has already been examined. One strand in this focuses on the use of *neural networks* for modeling the search landscape during optimization, in order to replace some real function evaluations with approximated ones [19, 8, 9], or to replace standard variation operators with adaptive ones [1]. The simpler concept of *fitness inheritance* has also been investigated in multiobjective optimization to economize on function evaluations [2, 5]. And a third strand is to use Bayesian network and/or other probabilistic model-building algorithms in a multiobjective scenario, e.g. [17].

However, while the above methods may offer some performance gains over standard MOEAs when function evaluations are expensive, not one of the studies above has demonstrated a significant performance advantage within the challenging evaluation budget we are interested in here. In this paper, we present and compare two recently proposed algorithms that take very different approaches to this challenge. The first algorithm, Binary-MSOPS, which is summarized below, is based on two separate pieces of work previously published by the second author [11, 12]. The second algorithm, ParEGO, was first described in a recent technical report [16], and is described here again in some detail. We evaluate these algorithms over a range of problems and, unlike in other studies, we focus explicitly on the first 250 evaluations only.

The rest of the paper is organized as follows. Sections 2 and 3 describe the two algorithms, while 4, 5 and 6 detail the test functions, performance assessment methods and parameter settings of the algorithms, respectively. Section 7 presents results and section 8 discusses findings and concludes.

2 Binary-MSOPS

The Binary-MSOPS algorithm is based primarily on the Binary Search Algorithm [11], summarized below. This method, which can be combined with almost any fitness assignment scheme, attempts to improve decision space sampling to ensure that promising