The Combative Accretion Model – Multiobjective Optimisation Without Explicit Pareto Ranking

Adam Berry and Peter Vamplew

School of Computing, University of Tasmania, Private Bag 100,
Hobart, Tasmania, Australia
{Adam.Berry, Peter.Vamplew}@utas.edu.au

Abstract. Contemporary evolutionary multiobjective optimisation techniques are becoming increasingly focussed on the notions of archiving, explicit diversity maintenance and population-based Pareto ranking to achieve good approximations of the Pareto front. While it is certainly true that these techniques have been effective, they come at a significant complexity cost that ultimately limits their application to complex problems. This paper proposes a new model that moves away from explicit population-wide Pareto ranking, abandons both complex archiving and diversity measures and incorporates a continuous accretion-based approach that is divergent from the discretely generational nature of traditional evolutionary algorithms. Results indicate that the new approach, the Combative Accretion Model (CAM), achieves markedly better approximations than NSGA across a range of well-recognised test functions. Moreover, CAM is more efficient than NSGAII with respect to the number of comparisons (by an order of magnitude), while achieving comparable, and generally preferable, fronts.

1 Introduction

As the artificial intelligence community realises the importance of multiobjective optimisation in real-world problem domains, research attention continues to grow, with a majority of the effort being focussed on the development and investigation of Multi-Objective Evolutionary Algorithms (MOEA) [1]. At the core of much of this research rests Pareto-ranking – a concept that has been prevalent since Goldberg's early work [2] and features in a host of techniques (such as NSGA [3], MOGA [4], NSGAII [5], SPEA [6] and SPEAII [7]). Such popularity is grounded on the assumption that "Pareto ranking is the most appropriate way to generate an entire Pareto front" [1], and results investigating its use certainly support such a theory. However, despite garnering both popularity and legitimately impressive results, Pareto-ranking is not without considerable limitations. The approach carries a significant complexity cost due to its reliance on population-wide comparisons and is typically accompanied by a diversity controlling parameter that is both difficult to tune and generally expensive to use. Such is the level of computational burden that populations are fundamentally limited in size and the potential for MOEA use in high-dimensional or difficult real-world problem domains is restricted.

C. A. Coello Coello et al. (Eds.): EMO 2005, LNCS 3410, pp. 77 – 91, 2005.

While a minority of contemporary algorithms endeavour to address the complexities introduced by the Pareto-ranking approach (see section 2.1), most second generation MOEA techniques extend the procedure through the inclusion of archiving, elitism and minor variations in the selection procedure (such as SPEA, SPEAII, PAES [8] and, to a lesser extent, PESA [9] and PESAII [10]). The inference that can be drawn from such a trend is that archiving and elitism are beneficial inclusions for general MOEA design, though results supporting such a claim are limited and lacking theoretical rigour. Moreover, given that the inclusion of an active secondary population generally incurs increased complexity, it is worth considering that such archiving need not be a pre-requisite for contemporary MOEA systems at all.

Consequently, this paper presents a model that moves away from active-archiving, while adopting an adaptable, inexpensive and implicit Pareto-ranking scheme that is grounded in pair-wise comparisons and simple diversity control. Furthermore, the population life-cycle is continuous rather than discrete (akin to Artificial Life systems) and agent generation is largely accretion – based on a consolidation of genes from an adaptive gene pool. Thus, the Combative Accretion Model (CAM) proposed herein represents a particularly novel approach to MOEA design that focuses on reducing complexity whilst maintaining high levels of performance.

2 Background

2.1 Pareto Ranking

Since the aim of all multiobjective optimisers is to develop an approximation of the Pareto optimal front, it is not particularly surprising that both contemporary and traditional efforts largely favour population-based Pareto dominance as a measure of fitness. By promoting those solutions that are non-dominated with respect to the current population, selection pressure favours exploration of potentially promising areas of the search space and focuses investigation on the current non-dominated front.

While it is apparent that measuring Pareto dominance is valuable in determining the direction of search, it is also significantly expensive. Even in the simplest case, where the population is divided into just two classes, the complexity[1] is $O(n^2)$ and infers a limiting bound on feasible population sizes. Such expense is only exacerbated as ranking becomes more fine-grained and continuous subdivision of dominated fronts is required (as in NSGA and NSGAII).

The complexities inherent in population-wide ranking have led to a number of algorithmic alternatives in the literature. Perhaps the most obvious approach, and the one adopted by Horn and Nafpliotis [11], is to reduce the percentage of the population under consideration when assessing dominance. In the Niched-Pareto Genetic Algorithm (NPGA), a tournament selection procedure is used, where the victor is determined by a single layer ranking process based on only ten percent of the total population (with ties broken through diversity estimation). By limiting the size of the

[1] This paper measures complexity in terms of the number of solution comparisons per evaluation (as per [1]) – objective comparisons are an equally valid measure and can be obtained by increasing comparison complexity by a factor of k, where k is the number of objectives.