Preliminary Investigation of the 'Learnable Evolution Model' for Faster/Better Multiobjective Water Systems Design

Laetitia Jourdan, David Corne, Dragan Savic, and Godfrey Walters

Schoolof Engineering, Computer Science and Mathematics, Harrison Building,
University of Exeter,Exeter EX4 4QF, United Kingdom
jourdan@lifl.fr, {D.W.Corne, D.A.Savic, G.A.Walters}@ex.ac.uk

Abstract. The design of large scale water distribution systems is a very difficult optimisation problem which invariably requires the use of time-expensive simulations within the fitness function. The need to accelerate optimisation for such problems has not so far been seriously tackled. However, this is a very important issue, since as MOEAs become more and more recognised as the 'industry standard' technique for water system design, the demands placed on such systems (larger and larger water networks) will quickly meet with problems of scaleup. Meanwhile, LEM (Learnable Evolution Model') has appeared in the Machine Learning literature, and provides a general approach to integrating machine learning into evolutionary search. Published results using LEM show very great promise in terms of finding near-optimal solutions with significantly reduced numbers of evaluations. Here we introduce LEMMO (Learnable Evolution Model for Multi-Objective optimization), which is a multi-objective adaptation of LEM, and we apply it to certain problems commonly used as benchmarks in the water systems community. Compared with NSGA-II, we find that LEMMO both significantly improves performance, and significantly reduces the number of evaluations needed to reach a given target. We conclude that the general approach used in LEMMO is a promising direction for meeting the scale-up challenges in multiobjective water system design.

1 Introduction

Fast optimization (in terms of using as few fitness evaluations as possible) is more and more essential when faced with real-world problems in which each fitness evaluation involves running an time-expensive simulation. However, there is of course a compromise between speeding up the optimization method, and ensuring good quality in the obtained solutions (see figure 1). One area in which this issue is paramount is in the design of large scale water distribution networks. This area contains various difficult and complex optimisation problems which invariably requires the use of time-expensive simulations within the fitness function. These problems are typically multi-objective (often trading off financial cost against requirements for pressure, speed of flow and other aspects of the

C. A. Coello Coello et al. (Eds.): EMO 2005, LNCS 3410, pp. 841–855, 2005.

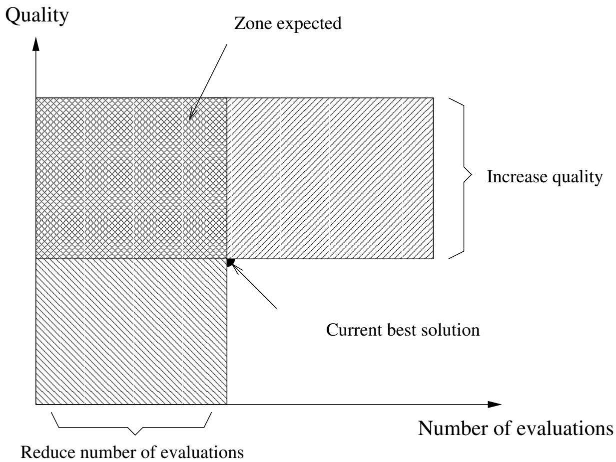

Fig. 1. Partition of the potential solutions

required design), and there is an increasing body of published research which addresses such problems, e.g.: [2, 26, 9, 23, 6, 4, 7, 1] However, the need to accelerate optimisation for such problems has not so far been seriously tackled. This is presumably because the problems involved are those of *design*, and sufficient time often exists to allow long optimisation runs before a final design is to be scrutinised and approved. However, as the quality of MOEA approaches causes them to be considered for ever larger problems, the time taken to find good solutions to such larger problems (e.g. a water design network with thousands of pumps, which is not uncommon) may turn out to be unacceptable with existing methods.

In this article, we address this scale-up problem by investigating a method based on the Learnable Evolution Model (LEM) [16, 17], which has appeared in recent years in the machine learning literature, but has so far been very little explored on real problems. The idea, which involves combining machine learning with evolutionary search, is a very generalised notion of which certain current trends in evolutionary computation (such as Estimation of Distribution Algorithms [15]) can be seen as specific instances. In particular, results to date on benchmark function optimisation problems show that LEM can save very significantly on the number of fitness evaluations needed to reach a certain target fitness. Our adaptation of this method to multiobjective optimisation, and in particular its application to water network design problems, is investigated here.

The remainder of the article is set out as follows. In section 2 we describe LEM and LEMMO (our version of this method). In section 3 we describe water systems design optimisation problems in general, and present the benchmark problems we use and briefly describe the associated fitness function simulator. In section 4, we describe the implementation of our genetic algorithm, the different variants of LEMMO and the metrics we use. Section 5 presents results from our LEMMO method, compared against NSGA-II (a well-known high-quality MOEA). Some concluding discussion is provided in section 6.