Environmental Modelling & Software 69 (2015) 222–225

Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

Preface

Thematic issue on Evolutionary Algorithms in Water Resources☆

Evolutionary Algorithms (EAs) and other similar optimisation approaches have become very popular in the water resources research literature over the last two decades. One reason for the emergence of EAs in the literature is that they use evolutionary principles found in nature, “evolving” to find better solutions to complex water resources problems. Another reason is that evolutionary optimisation provides a natural extension to the use of simulation models, as EAs simply “bolt onto” existing models. Consequently, the resulting optimisation process is very intuitive, as the way EAs try different solutions and then learn from the outcomes of these trials is similar to the process humans adopt when manually “optimising” or adjusting solutions to problems via a simulation based approach. The only differences when EAs are used are that the decisions as to which options to try are made with the aid of evolutionary operators, rather than human judgement, intuition and experience, and that the number of options considered is much larger. Moreover, outputs of the EA process are equivalent to outputs of trusted simulation models. Therefore, the optimisation results from EAs tend to have more credibility than those obtained using alternative approaches, such as mathematical programming, since the latter generally require gross simplifications of problem representation.

Another attractive feature of EAs is that they are not necessarily prescriptive in the sense of suggesting “the” optimal solution. This is because they work with populations of solutions and therefore produce a number of near-optimal solutions, which might be similar in objective function space, but quite different in solution space. This enables consideration of factors other than those captured in the mathematical formulation of the optimisation problem when selecting the solution to be implemented. As a result of the loose coupling between the optimisation engine, which decides which parts of the solution space to explore, and the simulation model, which evaluates how well the selected solutions perform in relation to the objectives and/or whether constraints have been violated, EAs can deal with discontinuities and nonlinearities with ease, as long as these have been captured appropriately in the simulation model. Another advantage of EAs is that they are well suited to multi-objective problems, as they can evolve optimal trade-offs between objectives (i.e. Pareto fronts) in a single optimisation trial.

Given the fascination and intrigue associated with the ability to use evolutionary processes to optimise water resources problems, the practicality and intuitiveness associated with being able to make use of existing simulation models and the advantage of being able to solve complex problems, it is not surprising that research involving EAs has received significant attention. This research has demonstrated the undoubted potential of EAs in the sense that they can be applied to and perform well in a wide range of application areas. In addition, significant research effort on the development and testing of different types of EAs, evolutionary operators and algorithm parameterisation has resulted in the ability to find better solutions with reduced computational effort. However, while there are pockets of research that continue to significantly push the boundaries of knowledge in this field, there is also a large amount of research that continues to re-visit the same themes. For example:

- There continue to be a large number of papers on using an ever increasing number of EA variants for solving an ever increasing number of water resources problems, with little focus on *understanding why* certain algorithm variants perform better for certain case studies than others. In addition, there is no consistency in algorithms, algorithm implementations, performance criteria and case studies in the papers. The above factors make it extremely difficult to draw conclusions that are applicable to the wider research field and enable meaningful guidelines for the application of different algorithms to be developed.
- There continue to be a large number of studies that use theoretical or very simplistic case studies. However, there are significant challenges associated with the application of EAs to real-world problems that need to be addressed in order to increase their uptake in industry.

In order to counteract potential repetition and stagnation in this field, Maier et al. (2014) identified a number of research questions that should be addressed. They suggest that the main areas in which research efforts should be directed include improving our understanding of algorithm performance and how to apply EAs to real-world problems, as summarised in Table 1. The 18 papers in this thematic issue begin to address some of these research questions, as summarised in Table 2 and discussed below.

Gibbs et al. (2015) develop a relationship between metrics that quantify fitness function characteristics and the number of generations needed for a genetic algorithm to converge in a predetermined number of generations for a large number of synthetically generated test problems with different attributes. This relationship is then validated on two water distribution system optimisation problems, including the Cherry Hill-Brushy Plains network, which is a commonly used test problem, and the optimal operation of the Woranora water distribution near Sydney, Australia, which is a real-world case study. The ability to select

☆ Special Issue on Evolutionary Algorithms.

http://dx.doi.org/10.1016/j.envsoft.2015.05.003

Table 1
Summary of key research questions identified in Maier et al. (2014).

1. Research questions associated with improving our understanding of algorithm performance
1.1 Can we develop knowledge of the fundamental *characteristics of the problem being optimised* at the level at which optimisation algorithms operate?
1.2 Can we develop knowledge of the *underlying searching behaviour* of different search methodologies?
1.3 How can we rigorously *measure and improve the performance* of a selected search methodology?
2. Research questions associated with applying EAs to real-world problems
2.1 How do we best change the *formulation* of optimisation problems to cater to real-world problems?
2.2 What can be done to *reduce the size of the search space* for real-world problems?
2.3 How can *computational efficiency be increased* for real-world problems?
2.4 Which *searching mechanisms* are best for solving real-world problems?
2.5 What *termination/convergence criteria* are most appropriate for real-world problems?
2.6 What is the best way is to *convey the results* of the optimisation of real-world problems *to decision makers* and what is the *role of optimisation* in the *decision-making process*?
2.7 What is the best way to *take account of uncertainty* in the optimisation of realistic systems?
2.8 What is the best way to *implement* optimisation algorithms for realistic systems?

the population size that results in convergence for a given computational budget based on problem characteristics is likely to be very useful for solving real-world problems where computational issues are a problem, particularly in operational settings.

Zheng et al. (2015a) use a number of run-time behaviour analysis measures to better understand how a differential evolution (DE) EA explores the solution space and why it produces the solutions it does at various stages of searching for three water distribution system optimisation problems of varying complexity and different parameterisations of the DE. The ability to understand how algorithms and algorithm parameterisations navigate through the solution space throughout the search for different problems is vital in terms of the ability to select the most appropriate algorithms and their parameters, to design better algorithms and to dynamically adjust searching behaviour during an optimisation run in order to maximise performance.

Piscopo et al. (2015) address the important issue of problem formulation for real-world problems in the context of the application of the Borg multi-objective evolutionary algorithm to the problem of optimising engineered injection and extraction for groundwater remediation. While in the vast majority of optimisation studies in literature problem formulation is established *a priori* and treated as fixed, this is generally not the case when dealing with real-world problems. In this paper, a novel iterative optimisation approach is introduced, as part of which problem formulation is updated based on the results of prior rounds of optimisation.

Yang et al. (2015) tackle the issue of problem formulation for the real-world case study of optimising the hydropower reservoir operation of the Oroville–Thermalito Complex in California, USA. Particular attention is given to the impact of the simplification of the reservoir's highly non-linear storage-elevation relationship. In addition, the performance of a new multi-objective search technique (Multi-Objective Complex Evolution Global Optimization Method with Principal Component Analysis and Crowding Distance Operator) is compared with that of a number of other techniques, including the Multi-Objective Complex Evolution Global Optimization method, the Multi-Objective Differential Evolution method, the Multi-Objective Genetic Algorithm, the Multi-Objective Simulated Annealing approach and the Multi-Objective Particle Swarm Optimization scheme, in order to determine which searching behaviour performs best.

Fowler et al. (2015) introduce a formulation for the real-world problem of deciding which crop planting choices farmers should make when faced with competing revenue, water use and demand objectives. The formulation is tested on a hypothetical case study where MODFLOW-FMP2 is used as the simulation package and a Multi-objective genetic algorithm is used as the optimisation engine. An extensive sensitivity analysis is used to obtain a better understanding of the relationship between algorithm parameterisation, algorithm performance and problem characteristics.

Zheng et al. (2015b) address the issue of increasing computational efficiency of multi-objective optimisation problems by means of search space size reduction. This is achieved by decomposing the optimisation problem into a subset of smaller problems via graph theoretic approaches and optimising each of these subproblems independently. A novel approach is then used to

Table 2
Research questions addressed in papers in this thematic issue.

Paper	Research question										
	1.1	1.2	1.3	2.1	2.2	2.3	2.4	2.5	2.6	2.7	2.8
Gibbs et al. (2015)	X					X					
Zheng et al. (2015a)		X									
Piscopo et al. (2015)				X							
Yang et al. (2015)				X			X				
Fowler et al. (2015)	X			X			X				
Zheng et al. (2015b)					X	X					
Creaco and Pezzinga (2015)					X	X					
Li et al. (2015)					X	X	X				
Zimmer et al. (2015)					X	X	X				
Dumedah (2015)					X	X					
Hadka and Reed (2015)						X	X				
Bi et al. (2015)						X					
Broad et al. (2015)						X				X	
Tsoukalas and Makropoulos (2015)						X	X			X	
McClymont et al. (2015)	X						X				
Lerma et al. (2015)				X			X	X	X		
Mortazavi-Naeini et al. (2015)										X	
Stokes et al. (2015)											X