Environmental Modelling & Software 25 (2010) 1552–1564

Contents lists available at ScienceDirect

Environmental Modelling & Software

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

Visualization-based multi-objective improvement of environmental decision-making using linearization of response surfaces

A. Castelletti [a,b,*], A.V. Lotov [c], R. Soncini-Sessa [a]

[a] Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza L. da Vinci, 32, 20133 Milano, Italy
[b] Centre for Water Research, University of Western Australia, 35 Stirling Hwy, Crawley 6009 WA, Australia
[c] Dorodnicyn Computing Centre, Russian Academy of Sciences, Vavilova str., 40, 119333 Moscow, Russia

ARTICLE INFO

Article history:
Received 9 February 2010
Received in revised form
13 May 2010
Accepted 21 May 2010
Available online 7 July 2010

Keywords:
Multi-objective decision methods
Multi-criteria analysis
Response surface
Pareto frontier visualization
Environmental planning
Water quality

ABSTRACT

This study presents a new interactive procedure for supporting Decision Makers (DMs) in environmental planning problems involving large, process-based, dynamic models and many (more than two) conflicting objectives. Because of such features of the model, computationally-onerous simulations are the only feasible way of analysis, while the multi-objective nature of the problem entails the combined use of optimization techniques and appropriate tools for the visualization of the associated Pareto frontier. The procedure proposed is based on the iterative improvement of the current best compromise alternative based on interactions with the DM. At each iteration, the DM is informed about the Pareto frontier of a local multi-objective optimization problem, which is generated by linearizing the response surfaces that describe the objectives and constraints of the original planning problem. Interactive visualization of the multi-dimensional Pareto frontier is used to support the DM in choosing the new best compromise alternative. The procedure terminates when the DM is fully satisfied with the current best compromise alternative. The approach is demonstrated in Googong Reservoir (Australia), which is periodically affected by high concentrations of Manganese and Cyanobacteria. Results indicate that substantial improvements could be observed by simply changing the location of the two mixers installed in 2007 and adding another pair of mixers.

1. Introduction

Contemporary environmental decision-making is faced with a twofold challenge: the complexity of the physical domain in which decisions are taken, where most of the processes are dynamic, spatially-distributed and highly non-linear, and the heterogeneousness of the socio-economic–ecologic context they affect, which usually involves multiple (often many), conflicting objectives. Multi-objective decision methods have been recognized as an important tool to support Decision Makers (DMs) in environmental planning and management (e.g. Janssen (1992); Lahdelma et al. (2000); Soncini-Sessa et al. (2003)). In general, the overall goal is to determine the Best Compromise Alternative (BCA, see Castelletti and Soncini-Sessa (2006)) among a number of available decision options according to the DM's preference, that is to find the alternative that satisfies the DM the most, considering the multiple objectives that have been identified by relevant stakeholders (Castelletti et al., 2008) to be the (only) issues on which a decision should be made. Assuming a rational DM, the BCA must be a Pareto-efficient alternative and therefore its determination involves the solution of a Multi-Objective (MO) optimization problem.

A wide variety of methods exists to determine the BCA and these are usually classified (Hwang and Masud, 1979) as a-priori, interactive (or progressive) and a-posteriori methods, accordingly to the stage at which the DM is involved in the process.

1.1. A-priori decision methods

A-priori methods are based on the elicitation and articulation of the DM preference structure and its subsequent use to transform the MO problem into a Single-Objective (SO) problem, whose solution is assumed to be the BCA. They include both normative approaches, such as the Multi Attribute Utility (Value) Theory (MAU(V)T) developed by Keeney and Raiffa (1976), and heuristic approaches like the Analytic Hierarchy Process (AHP) proposed by Saaty (1980) and goal programming (Charnes et al., 1955). A-priori

* Corresponding author at: Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza L. da Vinci, 32, 20133 Milano, Italy. Tel.: +39 02 2399 9601; fax: +39 02 2399 9611.
E-mail address: castelle@elet.polimi.it (A. Castelletti).

1364-8152/$ – see front matter
doi:10.1016/j.envsoft.2010.05.011

methods have been extensively experimented on environmental decision-making, including water resources (e.g. Keeney and Wood (1977); Agha (2006); Arnette et al. (in press)), fishery (e.g. Pascoe and Mardle (2001); Herath (2004)), forestry (e.g. Tecle et al. (1995); Hayashida et al. (2010)), and conservation planning (e.g. Moffett et al. (2005)). A major drawback in their application lies in the intricacy of identifying the DM preference structure, which might suffer from DM contradiction (or intransitivity) and the existence of a non-unique preference representation (e.g. for group preference (Arrow, 1963)). For instance, the identification of the utility (value) function in the MAU(V)T requires the DM to answer an extremely large number (which usually grows with the number of objectives) of complicated questions, to which a human being can hardly give stable and logical answers (see Tversky and Kahneman (1974) and Larichev (1992)). As the preference structure is articulated a-priori, with no information on the objective tradeoff, a-priori methods are said to yield uninformed decisions (Hwang and Masud, 1979).

1.2. Interactive decision methods

In interactive (or progressive) decision making the DM gives an initial guess on her preferences, which is used, following the a-priori approach, to convert the original MO problem into a SO problem. Based on the solution to this problem, the DM is allowed to refine her preference and the process is re-iterated until a satisfactory solution is found. Interactive methods include the Geoffrion-Dyer-Feinberg method (Geoffrion et al., 1972), the Tchebycheff method (Steuer and Choo, 1983), STEP (Benayoun et al., 1971) and the Reference Point (Wierzbicki, 1980). Also the well-known ELECTRE methods (Roy, 1991) can be classified as an interactive approach, which involves the weighting of the multiple objectives for constructing a binary relation (a feature of an a-priori method) and its subsequent modification by iteratively changing some threshold values. Also interactive approaches have been widely applied in environmental decision-making (e.g. Monarchi et al. (1973); Duckstein and Gershon (1983); Tecle et al. (1994); Cai et al. (2004)). With interactive methods, the DM is eventually able to make an informed decision, at least partially, and thus a major limit with the a-priori approaches is mitigated. The complexity and high number of questions to be posed to the DM remain an unsolved problem (Larichev, 1992), which is made even worse by the large number of iterations often required to come to an acceptable BCA.

1.3. A-posteriori decision methods

In contrast to the a-priori and interactive decision methods, in a-posteriori methods the DM preferences for the alternatives are expressed after the Pareto frontier (or an approximation of it) and the associated set of Pareto-efficient alternatives are identified, by solving the MO problem, and presented to the DM. The alternatives are initially assumed all to have the same preference. Then, the DM expresses her preference by analyzing objective tradeoffs and selecting a point on the Pareto frontier, whose associated alternative is the BCA. Basically, the main advantage of a-posteriori methods lies in the fact that the searching and decision processes are separated and used in sequence: thus the DM has full insight into her decision preference, the final decision is taken in a more informed way, and the decision-making process is more transparent. A-posteriori methods were proposed by Zeleny (1974), Cohon (1978), Chankong and Haimes (1983) and Steuer (1986), and, despite their nominal potential advantage over the other two groups of methods, have been only rarely applied to environmental problems (e.g. Cohon et al. (1979); Chankong and Haimes (1983); Lotov et al. (2004); Bekele and Nicklow (2005); Lotov et al. (2005a); Kennedy et al. (2008)). Indeed, a-posteriori methods suffer from two main weaknesses, which turn out to be particularly critical in dealing with environmental issues:

(1) The solution to the MO optimization problem for determining an approximation of the Pareto frontier and the associated set of Pareto-efficient alternatives requires a high computational effort, which usually grows superlinearly with the complexity (number of state variables) of the model adopted to describe the underlying physical process and with the number of objectives considered. As a result, most of the distributed-parameter, process-based, simulation models traditionally employed in environmental modelling can hardly be combined with the optimization algorithms available to solve MO problems. As an example, the ELCOM-CAEDYM model used in this study to simulate hydrodynamics and ecological processes in a 1.21×10^8 m^3 volume reservoir comprises 4.7×10^5 spatially distributed state variables (14 state variables in each cell) and has an associated real-to-run time ratio of 30:1.
(2) While in bi-objective problems the DM can easily explore the Pareto frontier and the associated set of Pareto-efficient alternatives by visual inspection of a graph or a simple table, with the growth in the number of the objectives considered, it might become very hard, or even impossible, for the DM to understand the properties of the Pareto frontier; especially to assess the objective tradeoff rate, which is key-information in learning the decision-making problem. Problems having three or more objectives, the so-called high-order Pareto optimization problems (also known as many-objective problems (Fleming et al., 2005)), are quite common in environmental decision-making (e.g. Lotov (1998); Lotov et al. (2004); Reed and Minsker (2004); Bekele and Nicklow (2005); Tang et al. (2007)). Without a proper representation of the Pareto frontier, especially the objective tradeoffs, the potential of a-posteriori methods for supporting more informed decisions is only a theoretical possibility.

1.4. Model reduction

An effective approach to overcoming the computational limitations mentioned at point 1 above relies on the mathematical reduction of the computationally-onerous, process-based model available to a simplified, computationally-efficient empirical (I/O) model, identified over a data set produced via simulation of the original model. In planning problems, this reduction can be performed by using the Response Surface (RS) methodology, which was first proposed by Box and Wilson (1951). This methodology involves the direct approximation of the multi-dimensional function (Response Surface) that maps the alternative decisions into the objectives of the planning problem. The approximate RS is a low-order regression model that can be more efficiently used in place of the original process-based model within any MO optimization framework. The approach has been extensively adopted in many modelling applications and optimization problems (see Myers and Montgomery (1995) and references therein), but has received little attention in environmental problems, except for several traditional environmental engineering studies (Fen et al., 2008; Fu et al., 2008). Lately, Castelletti et al. (2010) have proposed a novel approach in which the concept of RS approximation is re-interpreted in a interactive/iterative way as a sequential learning and planning process. An initial, small set of alternatives is simulated through the process-based model and the corresponding values of the objectives computed. Using an appropriate class of functions (e.g. linear interpolators, neural networks), a first approximation of