

Interactive response surface approaches using computationally intensive models for multiobjective planning of lake water quality remediation

A. Castelletti,^{1,2} J. P. Antenucci,² D. Limosani,¹ X. Quach Thi,¹ and R. Soncini-Sessa¹

Received 16 February 2011; revised 3 August 2011; accepted 4 August 2011; published 30 September 2011.

[1] Recent work has demonstrated the utility of interactive response surfaces in integrating decision making and complex process-based modeling of environmental systems. Specifically, we focus on problems that require computationally expensive dynamic models where a limited number of simulations are possible and so traditional meta-modeling might be not an option for the optimization process. Within this constraint, the interactive global response surface method aims to develop a complete picture of how the system responds to a planning decision on the basis of a restricted number of simulations. The interactive local response surface method utilizes the current best alternative to determine improvements on the basis of relatively small changes in the current decision. We outline the use of both methods to a real example of remediation in a water supply reservoir, demonstrating the key advantages and disadvantages of each method and describe a new, interactive cascade approach that provides an improved solution to the problem.

Citation: Castelletti, A., J. P. Antenucci, D. Limosani, X. Quach Thi, and R. Soncini-Sessa (2011), Interactive response surface approaches using computationally intensive models for multiobjective planning of lake water quality remediation, *Water Resour. Res.*, 47, W09534, doi:10.1029/2011WR010552.

1. Introduction

[2] Inland waters are under increasing stress from changes in land use, rainfall and runoff patterns, and increasing temperatures. In many cases, this is leading to an increasing prevalence of cyanobacteria [Jönk *et al.*, 2008], and thus deleterious impacts on water quality.

[3] Remediation of lakes and reservoirs has traditionally been associated with reducing external nutrient loads by catchment management of point and nonpoint sources [Schindler *et al.*, 2008], however, in-lake strategies can complement these approaches for certain objectives. In-lake strategies usually involve the installation of some mechanical device to increase hypolimnetic oxygen concentrations [Singleton and Little, 2006], increase vertical mixing [Schladow, 1993], or increase horizontal transport [Morillo *et al.*, 2009]. Hypolimnetic oxygenation aims to reduce internal nutrient loading by changing the redox conditions at the sediment/water interface so that the sediments are no longer a source of nutrients. Artificial destratification has a similar aim, however, it also breaks down the seasonal thermocline and thus potentially impacts the light regime experienced by the phytoplankton in a lake environment [Huisman *et al.*, 2004].

[4] The scientific basis and engineering design aspects of these interventions are thus well known, however, the proliferation of available technologies now makes the

efficient assessment of the best alternative a complex task. We define an alternative as a combination of different technological interventions, and it is the potential for combining technologies across various spatial scales that makes the number of feasible alternatives large. In addition, there is the need to demonstrate the preferred alternative is optimal in some sense, and so multiobjective optimization adds an additional layer of complexity to the problem.

[5] An assessment of remediation options typically assumes a dynamic, process-based model is available to simulate the system behavior under each possible alternative. The traditional approach to this problem is the “what-if” analysis, where the model is directly used to evaluate the alternatives and to choose the most preferred by comparison with other feasible alternatives. Alternatively, the decision criteria for remediation can be formalized as a set of indicators to be optimized using some multiobjective technique. Indicators are much less informative than the full model output, yet they have the clear advantage of allowing for a rational and effective evaluation of the alternatives. In principle, by solving this multiobjective optimization problem one obtains the Pareto frontier and the associated set of Pareto-efficient alternatives, wherein the decision-maker (DM) can eventually choose the best compromise alternative (BCA) [see Castelletti and Soncini-Sessa, 2006; Castelletti *et al.*, 2008].

[6] Existing multiobjective optimization approaches work well when the computational costs of the dynamic model are relatively small [e.g., see Rodriguez *et al.*, 2011]. Where the computational costs are excessive, a common approach is to replace the process-based model with a fast-to-run surrogate (or meta-model) [e.g., Broad *et al.*, 2005]. To improve the efficiency of traditional meta-modeling

¹Dipartimento di Elettronica e Informazione, Politecnico di Milan, Milan, Italy.

²Centre for Water Research, University of Western Australia, Crawley, Western Australia.

approaches, off-line sampling of the decision space has been replaced by more sophisticated techniques, such as sequential sampling [e.g., see *Yan and Minsker, 2006; Espinoza and Minsker, 2006; Jamshid Mousavi and Shourian, 2010*], that iteratively sample the decision space exploiting the information progressively gained during the optimization process. In most of these approaches the sampling strategy is driven by particular features of the structure of the approximate function (e.g., Kriging models) and/or of the optimization algorithm (e.g., genetic algorithms). While all these approaches ensure a good balance among efficiency, effectiveness, and reliability (e.g., they allow uncertainty analysis [*Kourakos and Mantoglou, 2008*]), they still require a remarkable number, in the order of thousands, of simulations of the process-based model. For our problem, this option is not available as the computational expense is extremely high, since the real-to-run ratio of the process-based model is in the order of 1:30. Hence, we require a multiobjective optimization procedure that must contend with a process-based model with extremely large computational cost, where less than 50 simulations are possible in total.

[7] Given the limited number of simulation runs possible, we must efficiently allocate these runs in the decision space in order to maximize the information they provide about the unknown Pareto frontier. Moreover, given the potential complexity of the decision space, resulting from the large variety of engineering options available and the strong nonlinearities characterized in the process-based model, we want to minimize the a priori assumptions on the regularities of the problem. To address this, a new a posteriori and interactive procedural approach was proposed by *Castelletti et al. [2010a, 2010b]*. This procedure iteratively computes an approximation of the Pareto frontier by combining the response surface methodology [*Myers and Montgomery, 1995*] and the interactive decision map technique [e.g., *Lotov et al., 2005* and references therein]. The first is used to obtain an approximation of the response surface (RS), which maps the alternatives into the indicator values; the second is to visualize the associated frontier and support the DM in his choice of the most “interesting” alternatives and, correspondingly, the new sample points in the decision space. A major difference with respect to the improved meta-modeling approaches mentioned above, is thus that the sampling strategy is not driven by a particular feature of the approximation scheme and/or the optimization algorithm but by the DM’s through an interactive approach. Interactive approaches have been used for genetic algorithm approaches in water management problems previously [*Babbar-Sebens and Mukhopadhyay, 2009; Singh et al., 2010*], however, the particular issues faced in those cases such as user fatigue are not considered here because of the restriction of the relatively few simulation runs possible.

[8] At any given iteration, this human-machine procedure consists of a learning and a planning phase (Figure 1), from which the name learning and planning procedure on the basis of a RS-approximation (LP-RS) is derived.

[9] In the *learning* phase, a sample of input (the engineering alternative to be evaluated) and output (the values of the indicators that result from this alternative) pairs is generated through a set of suitably designed simulation experiments conducted using the dynamic process-based

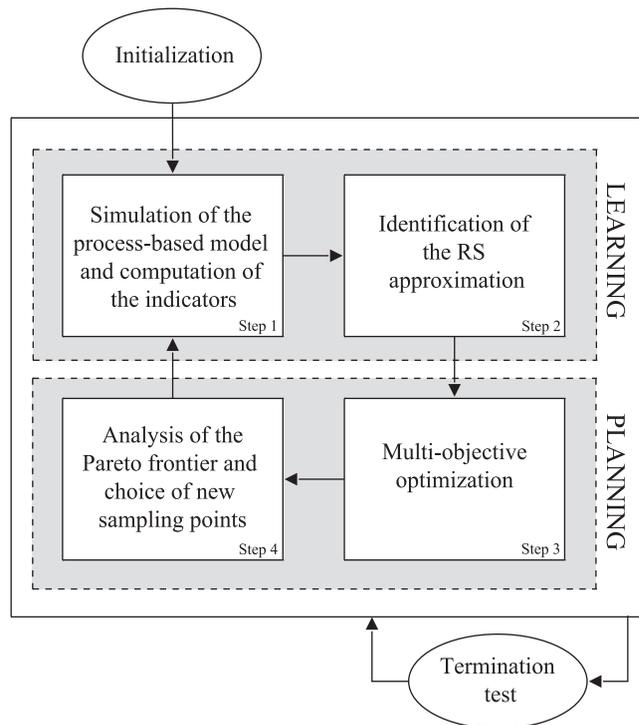


Figure 1. The learning and planning procedure based on RS-approximation.

model. The sample is then used to identify an approximation of the RS.

[10] In the *planning* phase, an approximate Pareto frontier and the associated set of Pareto-optimal alternatives are obtained by using the current RS-approximation and, possibly, shrinking the set of alternatives considered. The DM selects one or more “interesting” Pareto-optimal alternatives.

[11] Two approaches are possible for the approximation of the response surface: a global approach where the RS approximation is identified over the entire set of the feasible alternatives [*Castelletti et al., 2010a*] or a local approach, where the response surface is approximated in a neighborhood of the current best compromise alternative (BCA) by the tangent plane in the BCA [*Castelletti et al., 2010b*].

[12] In this paper these two approaches are compared in a case study involving the installation and operation of mechanical mixers to increase vertical mixing in a water supply reservoir. The set of feasible alternatives is quite large, as it involves combinations of the number of mixers to install, the position of each mixer, the maximum installed power, and the real-time control of mixers’ thrust. In this paper, we consider only the first two degrees of freedom (number and position) and assume the installed thrust as given with the mixers always operating at that thrust. The goal of our planning exercise is to determine the BCA, that is to find the alternative that satisfies the DM the most. The novel contribution of this paper is a new approach that combines the above RS approximation procedures in an interactive cascade (the interactive cascade approach), where the suboptimal solution obtained with the global approach is subsequently refined by applying the local one.

[13] The paper is organized as follows: We briefly review the approaches in section 2, focusing on how the alternative

definitions of the decision vector affect the problem formulation, and outline the case study to which they were applied. Each of the approaches is then applied to the case study to elucidate the key advantages and disadvantages, which are then discussed.

2. Methods

2.1. Global and Local RS Approximation

[14] We distinguish between a local and a global LP-RS procedure depending on the approach adopted for the RS approximation. With the global procedure [LP-gRS; Castelletti et al., 2010a] the RS approximation is identified over the entire set of feasible alternatives (see Figure 2a) and as a consequence, the associated Pareto front approximation can be used to evaluate trade-offs. With the local procedure [LP-lRS; Castelletti et al., 2010b] the RS is approximated in a local neighborhood of the current BCA (see Figure 2b), and the Pareto front obtained is only used to drive the DM in choosing an alternative that dominates the current BCA. As the local approach is based on linearization, it provides a good approximation of the real Pareto front only in a neighborhood of the current BCA, where the linear approximation is assumed to be acceptable. The size of this neighborhood is a priori unknown.

[15] Both the local and global procedures consist of an initialization step, four iterative steps, and a termination test (Figure 1), as follows:

[16] The initialization step involves selection of the alternatives to be simulated for the first iteration. With the global approach, design of experiments (DOE) can be employed for effectively sampling the alternative space either using statistical techniques or based on physical considerations and a priori knowledge [see Castelletti et al., 2011 and references therein]. For the local approach, the business as usual (BAU) alternative, or an equivalent suggested by an expert, can be used as initial current BCA.

[17] At the k -th iteration, follow the following steps:

[18] In step 1, for the local approach, the number of simulation runs is the same as the dimension of the decision

vector plus one, since each component of the current BCA, as selected at the previous iteration, is perturbed by a small amount plus the simulation of the BCA itself. For the global approach, the number of simulations to be run is equal to the number N of “interesting” alternatives selected by the DM at the previous step (or, when $k = 1$ to the number of samples designed by the DOE). This figure is affected by computational resource limits but is generally greater than the number simulated for the local approach.

[19] In step 2, for the local approach, the sample obtained in step 1 and the performance of the current BCA are used to identify a *linear* approximation of the unknown RS. In contrast, the global approach uses input/output data (i.e., alternatives and the associated value of the indicators) generated *up to* the current iteration to identify a global *nonlinear* approximation of the unknown RS.

[20] In step 3, using a suitable optimization algorithm, the Pareto frontier associated to the current RS approximation is computed. For the local approach, this optimization requires the solution to be within an a priori defined neighborhood \mathcal{D} of the current BCA such that the local linear approximation to the RS can be assumed to be acceptable. Note that the quality of the linear RS approximation quickly degrades as we move away from the current BCA (Figure 2b). As for the optimization algorithm, the choice is driven by the characteristics of the problem. Given the MO objective nature of the problem, evolutionary algorithms are generally preferred over gradient-based methods [see Castelletti et al., 2010a].

[21] In step 4, for the global approach, the DM analyzes the Pareto frontier obtained at step 3 with a suitable visualization tool and chooses N “interesting” points on the Pareto front according to his judgment. The alternatives that correspond to these points are obtained by inverting the RS approximation and constitute the new input sample that will be simulated at step 1 of the subsequent iteration. For the local approach, only one point is selected as the tentative new BCA.

[22] The procedure can be terminated using different criteria. With the global approach, the procedure can be stopped when the RS approximation converges between

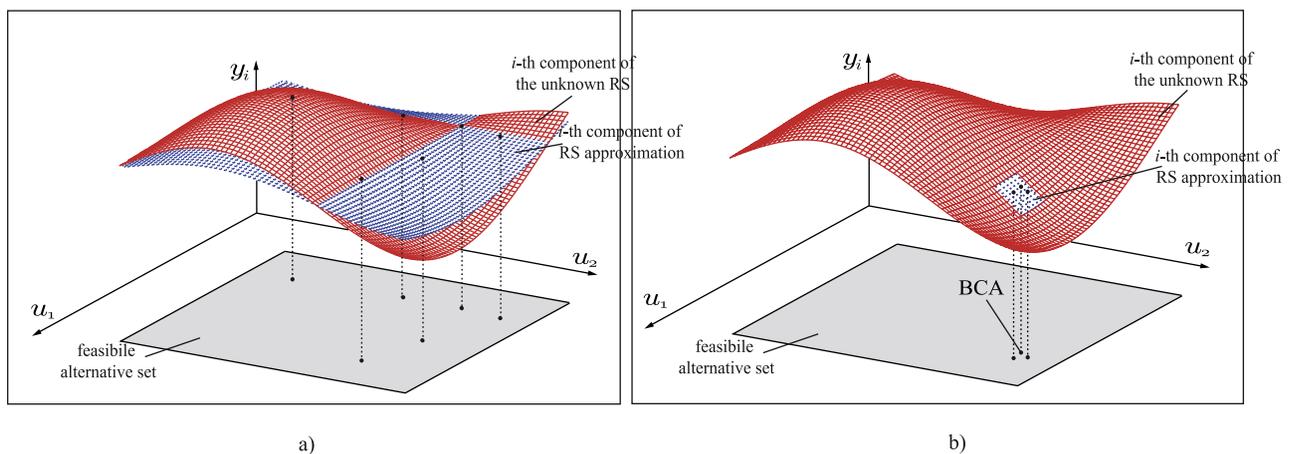


Figure 2. (a) Global and (b) local approximation of the i -th component of the unknown RS (say $i = 1, \dots, m$). In both the panels, the red surface is the unknown surface; the blue surface is the approximated one. (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

two iterations. Alternatively, the procedure can be terminated when the average distance between the available output vectors simulated via the numerical model and the corresponding output computed through the RS approximation is below a given threshold. When the procedure is over, the DM is provided with an approximation of the Pareto frontier that includes the “interesting” alternatives lastly chosen and among which the final BCA will be selected. By construction, this approximation is more accurate in the region corresponding to the “interesting” alternatives.

[23] For the local approach, if the performance obtained via simulation of the tentative new BCA is preferred over the performance of the current BCA, then the tentative BCA becomes the current BCA for the subsequent iteration. The procedure then continues from step 1 and terminates when either the DM is not able to find a point that improves the current BCA or is satisfied with it.

[24] At each iteration, the local procedure approximates the RS locally and linearly in a neighborhood of the current BCA. Therefore, the set of alternatives explored at each iteration is not the whole feasible set as in the global procedure, but a subset \mathcal{D} centered in the current BCA. The size of \mathcal{D} is key to the effectiveness of the procedure: too small might require a large number of iterations before getting to a Pareto-efficient alternative; too large might prevent us from determining an efficient alternative when the unknown RS is far from being linear within the neighborhood defined.

2.2. Interactive Decision Maps

[25] As for the effective representation of the multidimensional Pareto frontier in step 4 and the associated set of Pareto-efficient alternatives, there are an increasing number of studies [see Lotov *et al.*, 2005; Kollat and Reed, 2007; Madetoja *et al.*, 2008, and the review in chapters 8 and 9 of Branke *et al.* 2008] demonstrating that visualization tools can provide DMs with a graphical representation of the Pareto frontier in problems with three to seven objectives. Interactive decision maps (IDMs) [see Lotov, 1989 and references therein] are particularly effective and are used in this paper. IDMs are based on the approximation of the feasible set in the objective space and on displaying various decision maps, i.e., collections of bi-objective slices of the approximated frontier (an example is shown in Figure 9). For three objectives, the IDM is a prepared collection of several bi-objective trade-off curves, while the value of the third objective changes. The tradeoff curves are the frontiers of the slices. For more than three objectives various three-objective decision maps are constructed and displayed on request of the DM for different values of the fourth through the seventh objectives.

2.3. Problem Formulation and Suboptimal Approaches

[26] The optimal formulation of the problem would allow us to determine the optimal number and the locations of mixers (the *original multiobjective problem*). There are, however, operational difficulties in the resolution of this problem [see Castelletti *et al.*, 2010a], and thus we are forced to take a suboptimal approach and solve either one or the other of the following problems: we fix the number r

of mixers and determine their best location (problem P1), or we determine the best number \bar{r}^* of mixers with the location decided a priori by a suitable positioning rule (problem P2). From these two problems, a third option emerges: We first solve problem P2 and subsequently a problem (problem P3), which is problem P1, where \bar{r} is assumed equal to the optimal number \bar{r}^* of mixers determined from problem P2. In other words, problem P3 optimally reallocates the \bar{r}^* mixers.

[27] In detail, a planning alternative is described by the decision vector $\mathbf{u} = |r, \mathbf{Z}_1, \dots, \mathbf{Z}_r| \in \mathcal{U}$, whose components are the number r of mixers and the n -component vectors $\mathbf{Z}_1, \dots, \mathbf{Z}_r$ defining their positions, where n is the dimension of the set \mathcal{L} of feasible positions, obtained by considering the existing physical or technical constraints on the mixer positioning (e.g., minimum lake depth, minimum distance between mixers, distance from recreational or navigational zones). Assuming the maximum feasible number of mixers is R , the feasibility set \mathcal{U} is

$$\mathcal{U} = \{\mathbf{u} : \mathbf{u} \in \mathbb{R}^{nr+1}, r < R \text{ and } \mathbf{Z}_i \in \mathcal{L}, i = 1, \dots, r\}. \quad (1)$$

[28] A dynamic, process-based model \mathcal{M} is available to simulate the system behavior under each alternative. \mathcal{M} transforms an alternative \mathbf{u} into a collection \mathcal{B} of trajectories of spatially distributed outputs, i.e., $\mathcal{B} = \mathcal{M}(\mathbf{u})$.

[29] The DM’s decision criteria can be formalized in terms of m quantitative indicators y_1, \dots, y_m and the set of Pareto-optimal alternatives is obtained by solving the following multiobjective optimization problem:

$$\min_{\mathbf{u}} [y_1(\mathcal{B}, \mathbf{u}), \dots, y_m(\mathcal{B}, \mathbf{u})], \quad (2a)$$

subject to

$$\mathbf{u} = |r, \mathbf{Z}_1, \dots, \mathbf{Z}_r| \in \mathcal{U}, \quad (2b)$$

$$\mathcal{B} = \mathcal{M}(\mathbf{u}), \quad (2c)$$

Within this set the DM selects the best compromise alternative (BCA).

[30] Since the output \mathcal{B} can be computed for each alternative $\mathbf{u} \in \mathcal{U}$ via simulation, the indicator $y_i(\mathcal{B}, \mathbf{u})$ can be seen as a function $f_i(\mathbf{u})$ of \mathbf{u} , which encompasses both direct relations between \mathbf{u} and y_i and relations mediated by the model \mathcal{M} through \mathcal{B} . Since $\mathcal{U} \subseteq \mathbb{R}^{nr+1}$, the functions $f_i(\cdot)$, with $i = 1, \dots, m$, constitute the components of a vector function $\mathbf{f} : \mathbb{R}^{nr+1} \rightarrow \mathbb{R}^m$ called the response surface (RS). Equation (2) can be rewritten as

$$\min_{\mathbf{u}} [f_1(\mathbf{u}), \dots, f_m(\mathbf{u})], \quad (3a)$$

subject to (2b) and (2c).

[31] The definition of \mathbf{u} is an implicit formulation since the vector dimension $nr + 1$ depends on the value r of its first component. This is logically consistent, but the resulting definition of the feasible decision set $\mathcal{U} \subseteq \mathbb{R}^{nr+1}$, as well as the definition of the RS $\mathbf{f} : \mathbb{R}^{nr+1} \rightarrow \mathbb{R}^m$, become particularly tricky. To overcome this difficulty, we consider separately the decision r concerning the number of mixers

and the decisions $\mathbf{Z}_1, \dots, \mathbf{Z}_r$ about the positioning. Two alternative formulations can be considered: (1) the number \bar{r} of mixers to be installed is a priori fixed and only the mixer positions are to be decided; and (2) the positions $\mathbf{Z}_1, \dots, \mathbf{Z}_r$ of mixers are empirically defined by an a priori given function $\mathbf{Z}(\cdot)$ of r , i.e., $|\mathbf{Z}_1, \dots, \mathbf{Z}_r| = \mathbf{Z}(\cdot)$, and the number r is the unique decision variable. Formally, equation (3) can be reformulated by substituting (2b) for one of the following expressions:

[32] Formulation F1: $\mathbf{u} = |\mathbf{Z}_1, \dots, \mathbf{Z}_{\bar{r}}|$ with \bar{r} given;

[33] Formulation F2: $\mathbf{u} = r$ with $|\mathbf{Z}_1, \dots, \mathbf{Z}_r| = \mathbf{Z}(\cdot)$, where $\mathbf{Z}(\cdot)$ is an a priori given function.

[34] We refer to the associated problems as problems P1 and P2.

[35] Neither one nor the other of these problems is equivalent to equation (3) and therefore they yield a suboptimal solution. However, when formulation F1 is adopted, by repeatedly solving the associated problem P1 for $\bar{r} = 1, \dots, R$ and considering the nondominated points of the union set of the Pareto frontiers, one obtains the solution to equation (3). To improve the approximation obtained using formulation F2, the lake bathymetry is divided into a fixed number p of macro areas and decides the number of mixers $r_i (i = 1, \dots, p)$ to install in each area, i.e., $\mathbf{u} = \mathbf{r}$. The quality of the approximation improves with the number of macro areas considered (i.e., with the reduction of the surface of each macro area).

[36] The suboptimal solution $(\mathbf{r}^*, \mathbf{Z}(\mathbf{r}^*))$ to equation (3), obtained by solving problem P2, can be refined by subsequently solving problem P1 for $\bar{r} = \sum_{i=1}^p r_i^*$, that is problem

P3. Therefore, an improved approximate solution to equation (3) can be computed with the following steps.

2.3.1. Cascade Approach

[37] Step 1. Solve problem P2 over p macro areas using the global approach, thus obtaining the set of Pareto-optimal alternatives. The DM chooses the BCA among these alternatives, which is specified by the optimal number r_i^* of mixers in each macro area and the associated positions $\mathbf{Z}_i(r_i^*) i = 1, \dots, p$.

[38] Step 2. Solve problem P1 with the local approach, starting from the BCA obtained in step 1 and assuming

$$\bar{r} = \sum_{i=1}^p r_i^*.$$

2.4. Problem Formulation Versus RS Approximation

[39] Depending on whether the problem is formulated as P1, P2, or P3, one of the two RS approaches seems to be intuitively more appropriate. The local approach explores the feasible alternative set starting from a given alternative (usually the BCA alternative) and follows a path iteratively constructed in the neighborhood of the alternative currently preferred by the DM. When the feasible set is large and/or a good starting alternative is not available, the relative simplicity of the linear approximation is overwhelmed by the computational cost of the many iterations required for exploring the alternative set and hence the global approach seems to be more appropriate. This might be the case with problem P1. With problem P2 there is no a priori reason to prefer one approach over the other, while for problem P3 the local approach is most suitable since the starting alternative

is generated by solving an optimization problem (i.e., problem P1), and thus it should represent a “good” starting point. To verify these a priori conjectures, the two RS approaches were comparatively tested using a real world application.

3. Case Study

3.1. Study Site

[40] The Googong Reservoir is located in New South Wales, Australia (see Figure 3), and is one of five sources supplying Canberra’s water. The reservoir has a full-supply volume of $1.21 \times 10^6 \text{ m}^3$, a surface area at full supply of $\sim 3.5 \text{ km}^2$, an average depth of 35 m, and a maximum depth of 50 m, and is used for potable water supply and recreational purposes. It has a history of low to medium levels of cyanobacteria, namely *Anabaena*, which has the potential to produce taste and odor compounds and, at high concentrations, neurotoxins. An artificial destratification system was installed to address this problem in March 2007 (the BAU alternative), consisting of two pairs of 5 m diameter downward-pointing, surface-mounted impellers located as shown in Figure 3a. The mixers are surrounded by a draft tube (Figure 3b) reaching a depth of $\sim 10 \text{ m}$ and produce $\sim 3 \text{ m}^3 \text{ s}^{-1}$ per mixer.

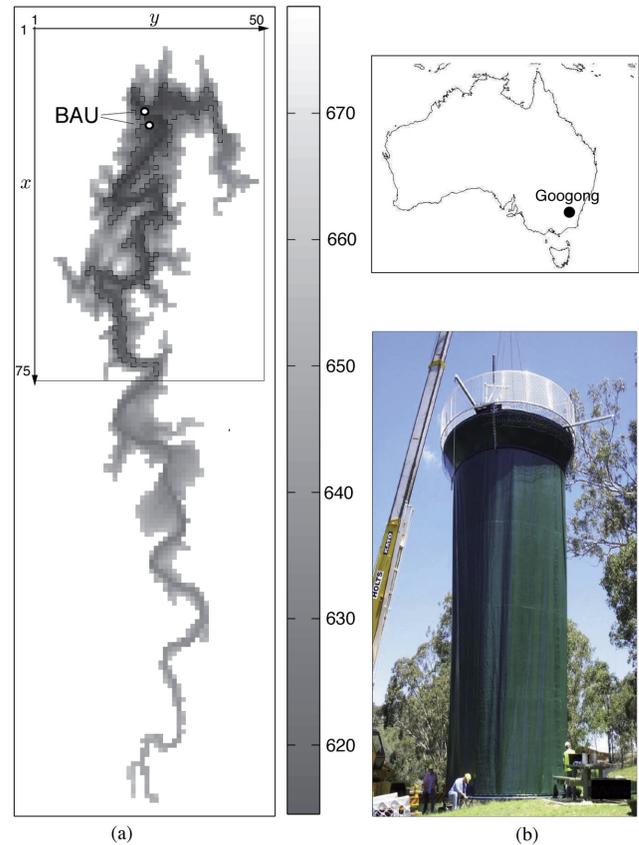


Figure 3. (a) The Googong reservoir grid bathymetry (60 × 60 m) and the part of the lake (delimited by a dashed line) where the positioning of mixers is feasible. Cartesian coordinates are used to locate mixers on the model grid. The location of the two pair of mixers installed in 2007 (BAU alternative) is shown. (b) A picture of one of the draft tube mixer installed.

3.2. The Indicators

[41] The management concern relates to *Anabaena* and manganese concentrations exceeding some reference value. We will consider the average annual number of days in which the concentration of chlorophyll-a associated with *Anabaena* in the epilimnion exceeds a given threshold, i.e.,

$$y_A = \frac{365}{h} \sum_{t=1}^h \chi_t, \quad \chi_t = \begin{cases} 1 & \text{if } A_t > \bar{A} \\ 0 & \text{otherwise} \end{cases}, \quad (4)$$

where A_t is the average concentration of chlorophyll-a in the epilimnion (temporal mean on day t) and \bar{A} is the critical threshold ($2 \mu\text{g}$ chlorophyll-a/L as stated in the Australian Drinking Water Guidelines) and h is the total number of days in the simulation horizon.

[42] As for manganese, any increase in its concentration in the hypolimnion can be either because of poor (or null) oxygenation in the hypolimnion and/or driven by the inflow. We will consider the average annual number of days in which the concentration of manganese in the layer just above the sediments exceeds a given threshold, i.e.,

$$y_M = \frac{365}{h} \sum_{t=1}^h \chi_t, \quad \chi_t = \begin{cases} 1 & \text{if } Mn_t > \overline{Mn} \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where Mn_t is the average concentration of manganese in the benthic layer (temporal mean on day t) and \overline{Mn} is the critical threshold (0.065 mg L^{-1} as stated in the Australian Drinking Water Guidelines).

[43] Finally, the installation cost should be taken into proper consideration. Assuming such cost is linearly proportional to the number of mixers installed, we may consider the total number of mixers as a proxy of the monetary indicator, that is

$$y_C = r, \quad (6)$$

where r is the number of mixers installed.

3.3. The Process-Based Model

[44] The 3-D coupled hydrodynamic-ecological model ELCOM-CAEDYM [Hodges et al., 2000; Hipsey et al., 2007] has been used to simulate the seasonal cycles in the Googong reservoir and event driven dynamics (wind and inflow). The Estuary and Lake Computer Model (ELCOM) is a three-dimensional hydrodynamics model used for predicting the velocity, temperature, and salinity distribution in natural water bodies, subjected to external environmental forcing such as wind stress and surface fluxes. The hydrodynamic simulation method solves the unsteady, viscous Navier-Stokes equations for incompressible flow, using the hydrostatic assumption for pressure. Simulated processes include baroclinic and barotropic responses, rotational effects, wind stresses, surface thermal forcing, inflows, outflows, and transport of salt, heat, and passive scalars. The hydrodynamic algorithms in ELCOM are based on the Euler-Lagrange method for advection of momentum with a conjugate-gradient solution for the free-surface height. The Computational Aquatic Ecosystem Dynamics Model (CAEDYM) was configured to simulate the major biogeochemical

processes influencing water quality, including primary production, nutrient and metal cycling, and oxygen dynamics. The ELCOM-CAEDYM model (simply model \mathcal{M} in what follows) was calibrated and validated against available field data for temperature, dissolved oxygen, nutrients, metals and chlorophyll-a, with results reported by Hillmer et al. [2008].

[45] For the purpose of producing sample data for RS approximation, the coupled ELCOM-CAEDYM model was run on a $60 \text{ m} \times 60 \text{ m}$ grid bathymetry with 1 m vertical grid resolution (Figure 3a) using a simulation time step of 2 min . One-year long simulation on a AMD Athlon(tm) 64 3.5 MHz required $\sim 7 \text{ d}$. The output required by the indicators has been sampled with two different time steps, every 12 h for the manganese and every 3 h for cyanobacteria which are much more sensitive to daylight variations.

4. Results

[46] The two response surface approaches were both tested for problems P1 and P2, while for problem P3 the local approach seemed the natural choice. With both approaches, the feasible goal method [Lotov et al., 2004] was adopted to solve the optimization problem. Given the academic nature of the study, when necessary, the role of the DM was played by a limnologist from CWR.

4.1. Problem P1: Optimizing the Location of a Fixed Number of Mixers

[47] Two mixers, as in the BAU alternative (Figure 3a), were considered with the purpose of verifying if any improvement of the BAU alternative was possible at no additional cost. Because of the irregular shape of the lake, it is convenient to describe mixer positions in curvilinear coordinates along the riverbed (the dark gray cells in Figure 4), giving a set of 121 options for the two mixers positions (u_1 and u_2) on the basis of the model spatial resolution. To avoid mutual mixer interference, the locations of the two mixers were made to differ by more than one grid point.

4.1.1. Global Approach

[48] Given the high computational burden associated with each simulation run, the initial alternatives were selected as those already simulated in previous studies [Castelletti et al., 2010a], which have the mixers positioned on the riverbed (left-bottom panel in Figure 4). Three iterations were carried out.

[49] At step 2 of each iteration, the approximation of the RS was obtained using radial basis functions (with spread values changing with the iteration from 2.6 to 2 for y_A and from 2.5 to 1.9 for y_M). y_C was not considered since the number of mixers is fixed. The Pareto frontier obtained at step 3 was visualized as an interactive decision map (IDM). The IDM obtained at the end of the third iteration (the procedure ends up in step 1 of the fourth iteration) is shown in Figure 5. Alternatives (44,95) (i.e., one mixer at location 44, the second at location 95 along the river channel) point A and (18,54) point B were considered to be “interesting” by the DM. The second alternative (B) dominates (i.e., performs better for both indicators) the BAU alternative while the first (A), although not dominating the BAU, yields a remarkable improvement in the algal indicator y_A (from

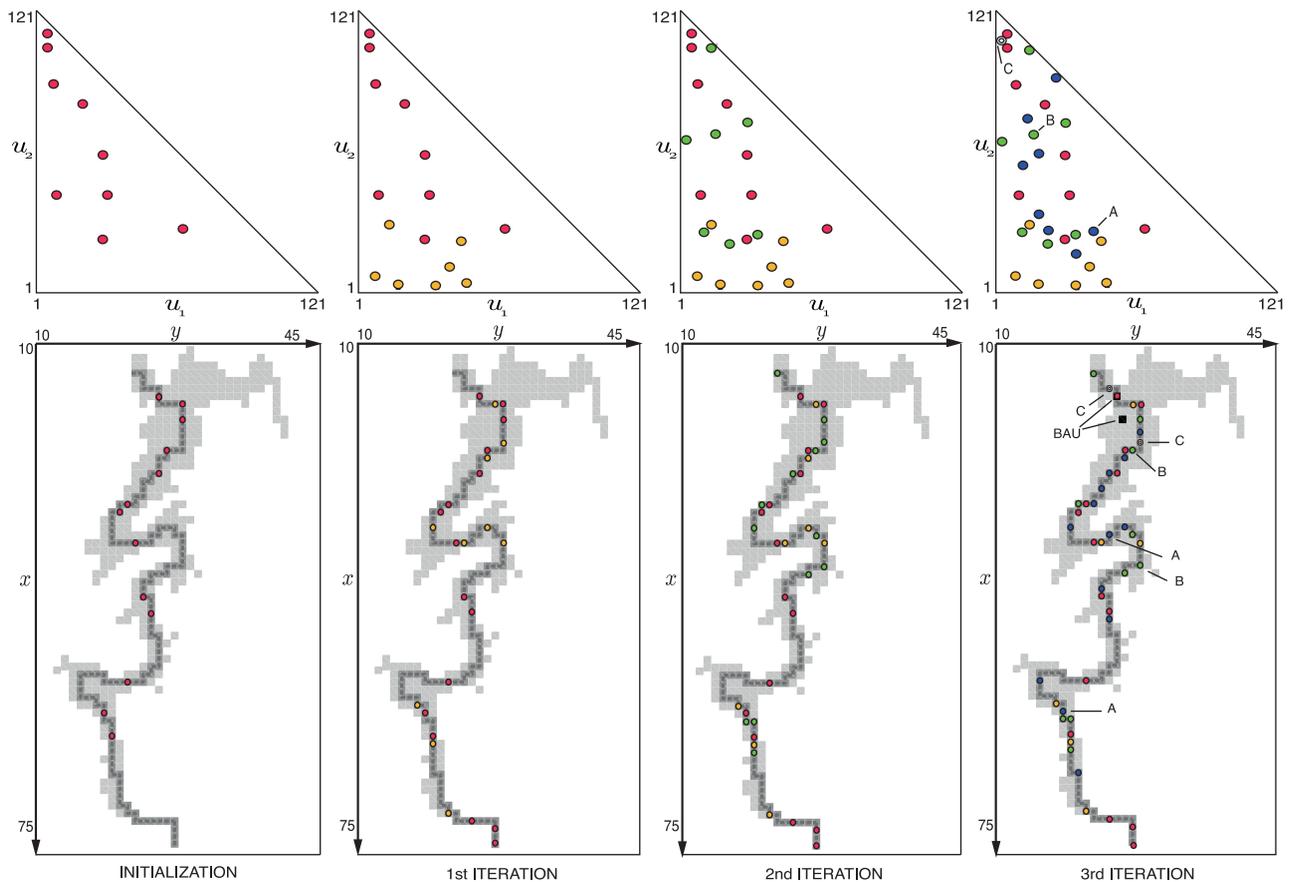


Figure 4. Problem P1/global. The alternatives considered at the (red circles) initialization step and (yellow, green, and blue circles, respectively) the subsequent three iterations of the LP-gRS procedure, represented (top) in the feasible decision set and (bottom) on the lake map. Positions A and B are the “interesting” alternatives obtained with the global approach. Position C is the one found by the local approach. The 121 dark cells indicate the riverbed and the mixer feasible positions. (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

180 to 48 d yr⁻¹) at the price of a small worsening of manganese indicator y_M (from 100 to 108 d yr⁻¹).

4.1.2. Local Approach

[50] The local approach aims at improving an existing solution and therefore the BAU alternative would be the natural starting point. However, since this is not on the riverbed, the nearest alternative on the riverbed was considered. At the first iteration of the local procedure, new alternatives were generated (steps 1–4) in the neighborhood \mathcal{D} (Figure 6) of the BAU. In step 4, among the alternatives that dominate the BAU alternative and have integer coordinates, the DM chose point C (yellow circle) in Figure 7a as the tentative new BCA. In the termination test, the dominance was confirmed by the simulation (green circle) and therefore this alternative became the new BCA. In iteration 2, the RS approximation suggested only one alternative that dominates point C (the yellow circle in Figure 7b). However, in the termination test this dominance was not confirmed (the green circle) and the alternative C resulted as the final BCA. This performed slightly better (see point C in Figure 5) than the alternative B obtained with the global approach and located the two mixers not far from the BAU (bottom right-hand panel of Figure 4).

4.2. Problem P2: Optimizing the Number of Mixers With a Predefined Positioning Rule

4.2.1. Global Approach

[51] This case was developed and described by *Castelletti et al.* [2010a]. In order to reach a good compromise between accuracy and computational burden, the lake was divided in three macro areas (Figure 8), each one having the same volume (one third of the lake volume), which gives a three-dimensional decision vector \mathbf{u} (the components are the number of mixers in each macro area). Additionally, we assumed that the volume of water influenced by each mixer be the same, which thus determined the position of any mixers placed within a macro area. The maximum allowed number of mixers compatible with both the recreational use of the lake, cost, and logistical considerations was fixed to 10, which results in 286 possible alternatives. The DOE was based on a nearly uniform sampling of this set, which yielded 28 starting alternatives.

[52] The decision map obtained at the end of the procedure (terminated at step 1 of the fourth iteration) is shown in Figure 9. The alternatives (0, 2, 0) (two mixers located in the second macro area, point D) and (0, 2, 2) (two mixers located in both the second and third macro area, point E)

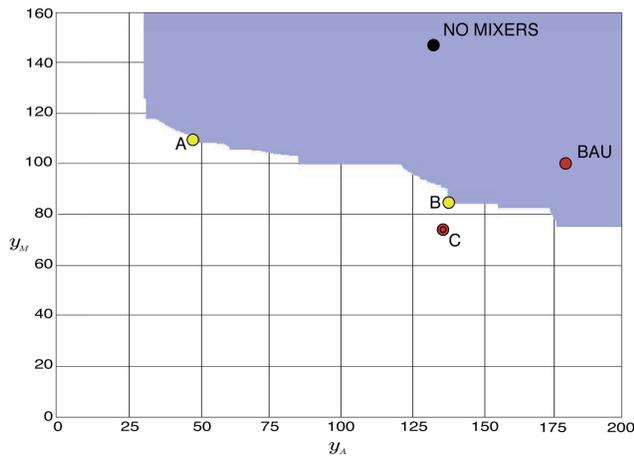


Figure 5. Problem P1/global. The decision map at the end of the procedure. The red and black circles are the performance, computed via simulation of model \mathcal{M} , respectively, of the BAU alternative and the no-mixer alternative; the yellow circles are the performances, computed with the RS approximation, of the “interesting” alternatives chosen by the DM. These coincide with the values obtained via simulation. The red concentric circle is the simulated performance of the BCA obtained through the local approach. (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

are two of the “interesting” alternatives [see *Castelletti et al., 2010a*]. Note that the BAU solution (two mixers in the first macro area) can be improved with the first alternative (two mixers in second macro area) with a reduction of $\sim 20\%$ in each of the two indicators y_A and y_M by simply changing the mixers’ position. By increasing the number of

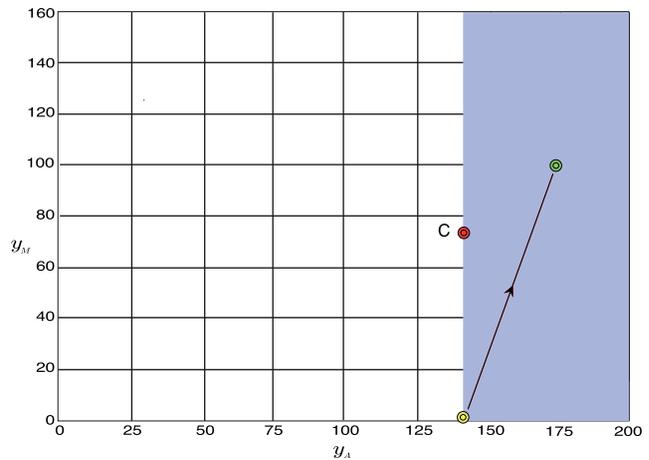
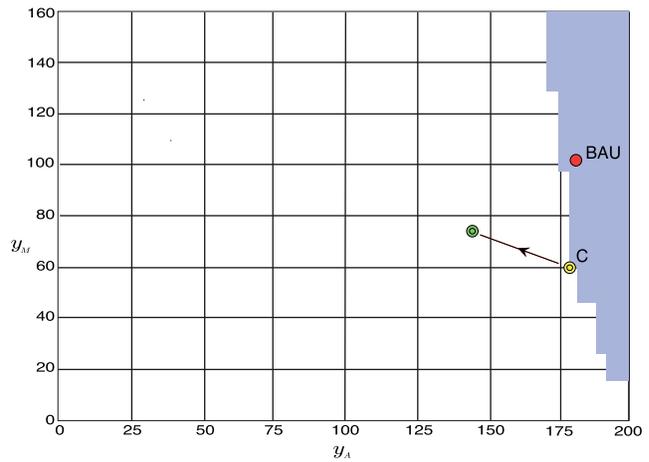


Figure 7. Problem P1/local. The decision map at (a) Iteration 1 and (b) Iteration 2. The green concentric circles are obtained via simulation of model \mathcal{M} , while the yellow circles are computed with the RS approximation. (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

mixers to four, the second alternative produces an even greater improvement of $\sim 50\%$.

4.2.2. Local Approach

[53] This case was developed by *Castelletti et al. [2010b]*. As in problem P1/local, the BAU alternative was chosen as the starting point, but the component of the decision vector was the number of mixers within each macro area. The procedure ended in four iterations and the succession of BCAs and associated performance are reported in Figure 10. The BCA eventually obtained is one of the alternatives found with the global approach to problem P2 (point E in Figures 8 and 9).

4.3. Problem P3: Optimizing the Location of the Optimal Number of Mixers

[54] The two alternatives $D = (0, 2, 0)$ and $E = (0, 2, 2)$ obtained above, are two suboptimal solutions to the *original multiobjective problem* (equation (2)). They can be refined by subsequently solving problem P1, which assumes the number \bar{r} of mixers is fixed to \bar{r}^* (2 and 4 for D and E, respectively). Therefore, the decision vector \mathbf{u} has

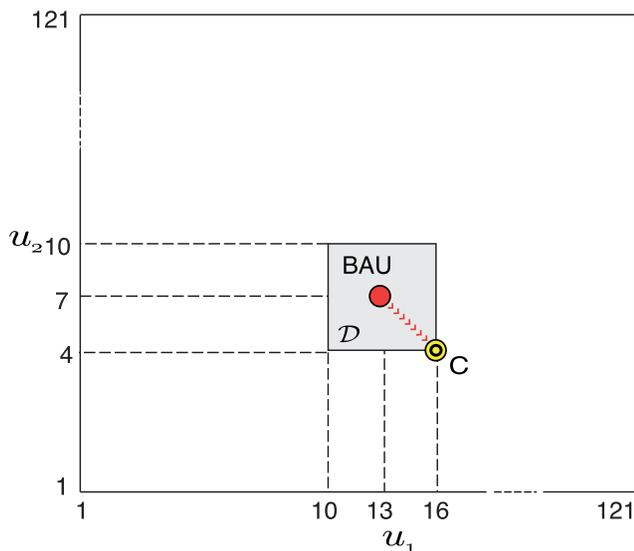


Figure 6. Problem P1/local. The new BCA (point C) chosen by the DM at iteration 1 starting from the BAU alternative (initialization) and given the neighborhood \mathcal{D} (shaded area), in which the new BCA was sought for.

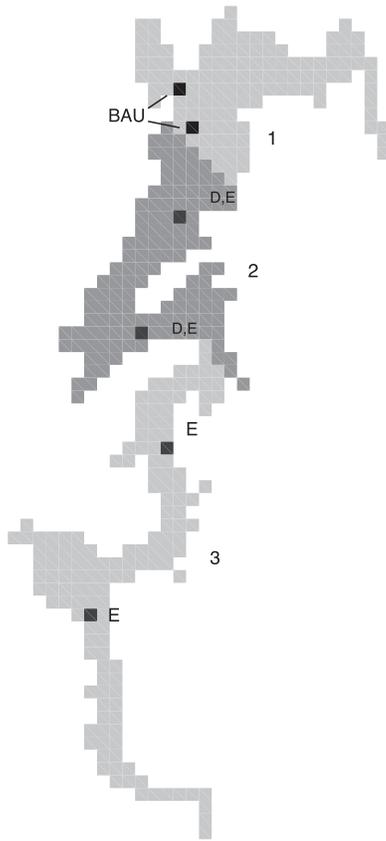


Figure 8. Problem P2/global. The three macro areas considered (different gray levels) and the positions of the mixers in the BAU alternative and the two “interesting” alternatives (D and E) discussed by *Castelletti et al.* [2010a].

two components (the locations of the mixers along the river bed) when refining D and four components when dealing with E. Since we presumably start from a good solution and we want to seek small improvements, the local approach appears to be the natural choice. Starting from D, the final BCA (alternative F in Figure 11) was found at the third iteration since no improvement was possible at the fourth. Starting from E, which is the common solution to problem P2 with both the global and local approach, the procedure ended after five iterations with alternative G (Figure 12).

[55] In both cases, the improvement with respect to the initial alternative is quite remarkable. In particular, alternative G results in the manganese indicator going to zero. The relatively even spread of the mixers across the reservoir in alternative G appears to sufficiently raise oxygen levels across the reservoir such that manganese release from the sediment (the major source in this system) is effectively controlled.

5. Discussion

[56] From the comparative analysis of the results of the above problems (Figure 13) we can draw some general considerations and suggestions about whether and when to adopt the global or the local approach, and which decision variable definition (i.e., which problem) is more preferred.

5.1. Global or Local?

[57] In principle, by repeatedly solving problem P1 for different numbers of mixers and using the global approach, the entire Pareto frontier of the original multiobjective problem is obtained. This would make this combination of problem/approach the most potentially interesting choice.

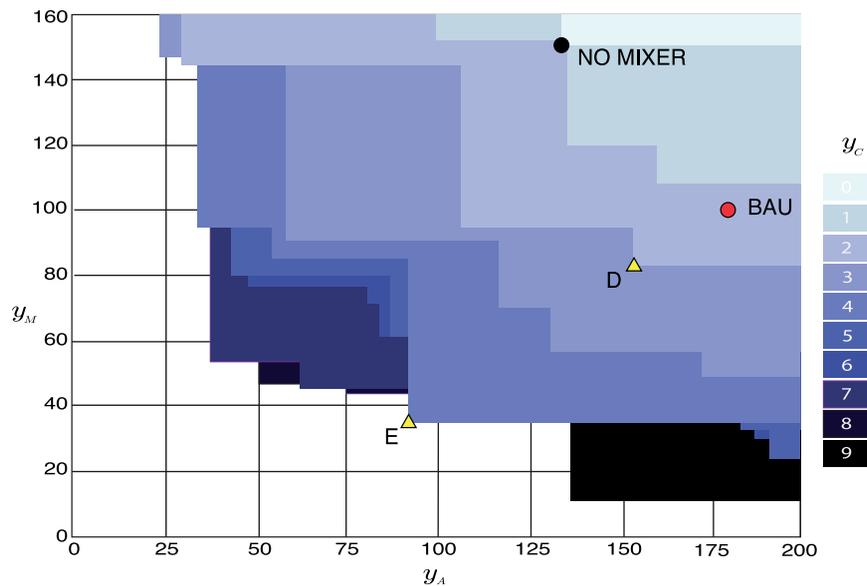


Figure 9. Problem P2/global. The decision map at the end of the procedure. The red and black circles are the performance, computed via simulation of model \mathcal{M} , respectively, of the BAU alternative and the lake without mixers; the yellow triangles are the performances, computed with the RS approximation and confirmed via simulation, of the “interesting” alternatives discussed by *Castelletti et al.* [2010a]. (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

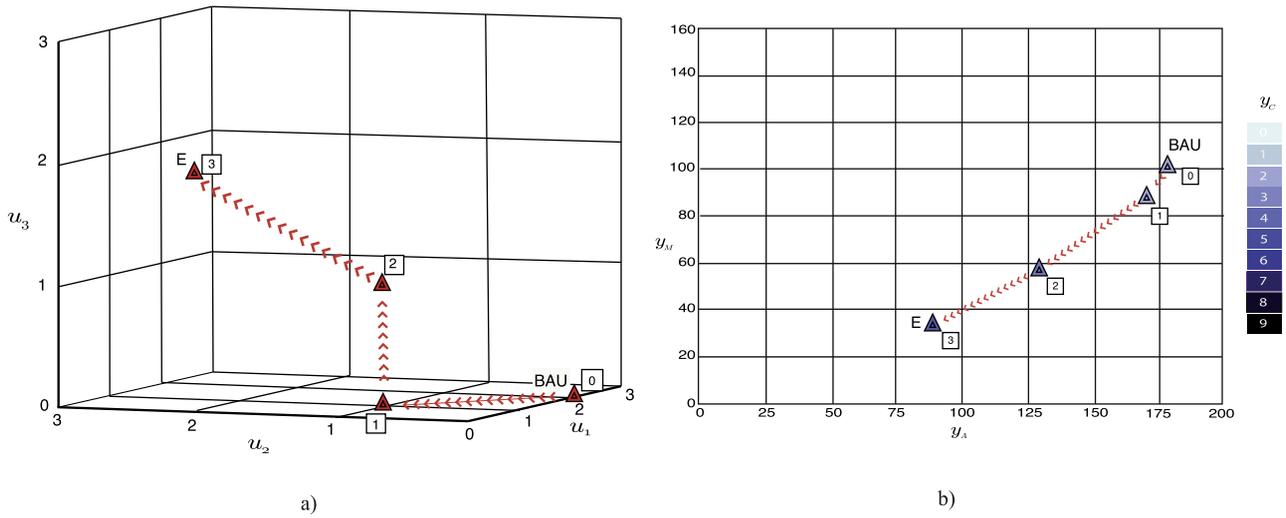


Figure 10. Problem P2/local. (a) The sequence of BCAs (red triangles) iteratively chosen by the DM in the three iterations performed with the local approach starting from the BAU. (b) The associated performance computed via simulation of model \mathcal{M} . (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

However, numerical results do not seem to support this a priori guess (Figure 13). In fact, the approximate Pareto frontier (black continuous line) obtained from problem P1/global is partly dominated, at least for the two mixer case, by points C and F calculated with the LP-/RS procedure. Second, the recursive resolution of problem P1 with the global procedure for different numbers of mixers \bar{r} is a computational challenge as the number of feasible alternatives grows slightly less than exponentially with \bar{r} , and thus the number of simulation runs necessary to generate an acceptable frontier approximation also grows quite rapidly. Generally, when the number of mixers is the same, the alternatives determined by the local approach always perform better than those obtained with the global approach, except for solution A (Figure 13). This should be expected as the

local procedure is always initialized with a “good” alternative. The local approach is clearly sensitive to the quality of the initial alternative, as proved by its different behavior while starting from E rather than D. However, a good starting alternative is, at least in principle, always obtainable at a relatively low cost either by running an expert-based “what-if” analysis or by few iterations of the global procedure. On the other hand, the global approach has the clear advantage of providing an approximation of the entire Pareto frontier of the original problem, while the local approximation is reliable only in a small neighborhood of the final BCA. This can be easily seen in Figure 14, where the decision map of problem P2 solved through the LP-/RS procedure (dashed line) provides an acceptable approximation of the Pareto frontier of the original problem only in a small

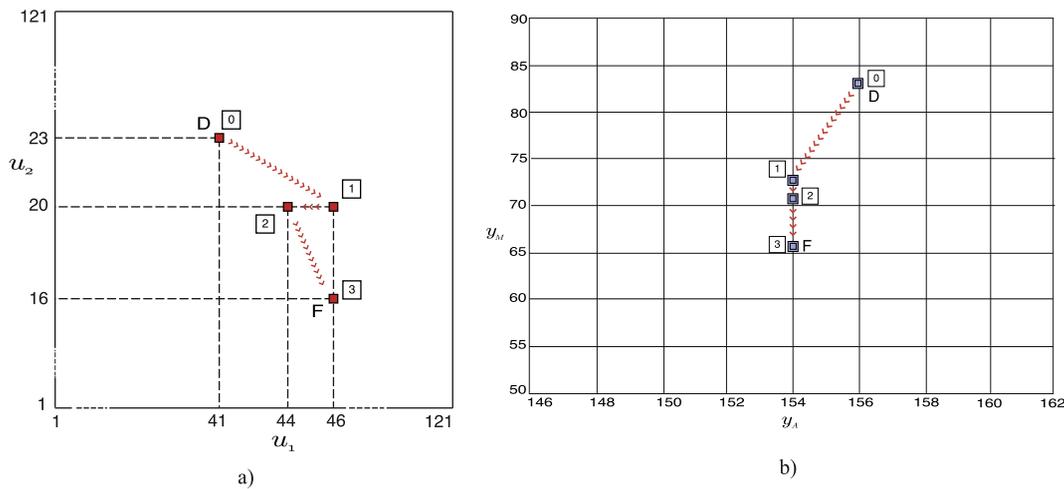


Figure 11. Problem P3/local. (a) The sequence of BCAs iteratively chosen by the DM in the four iterations performed with the local approach starting from alternative D (one of the solutions to problem P2). (b) The associated performances computed via simulation of model \mathcal{M} . (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

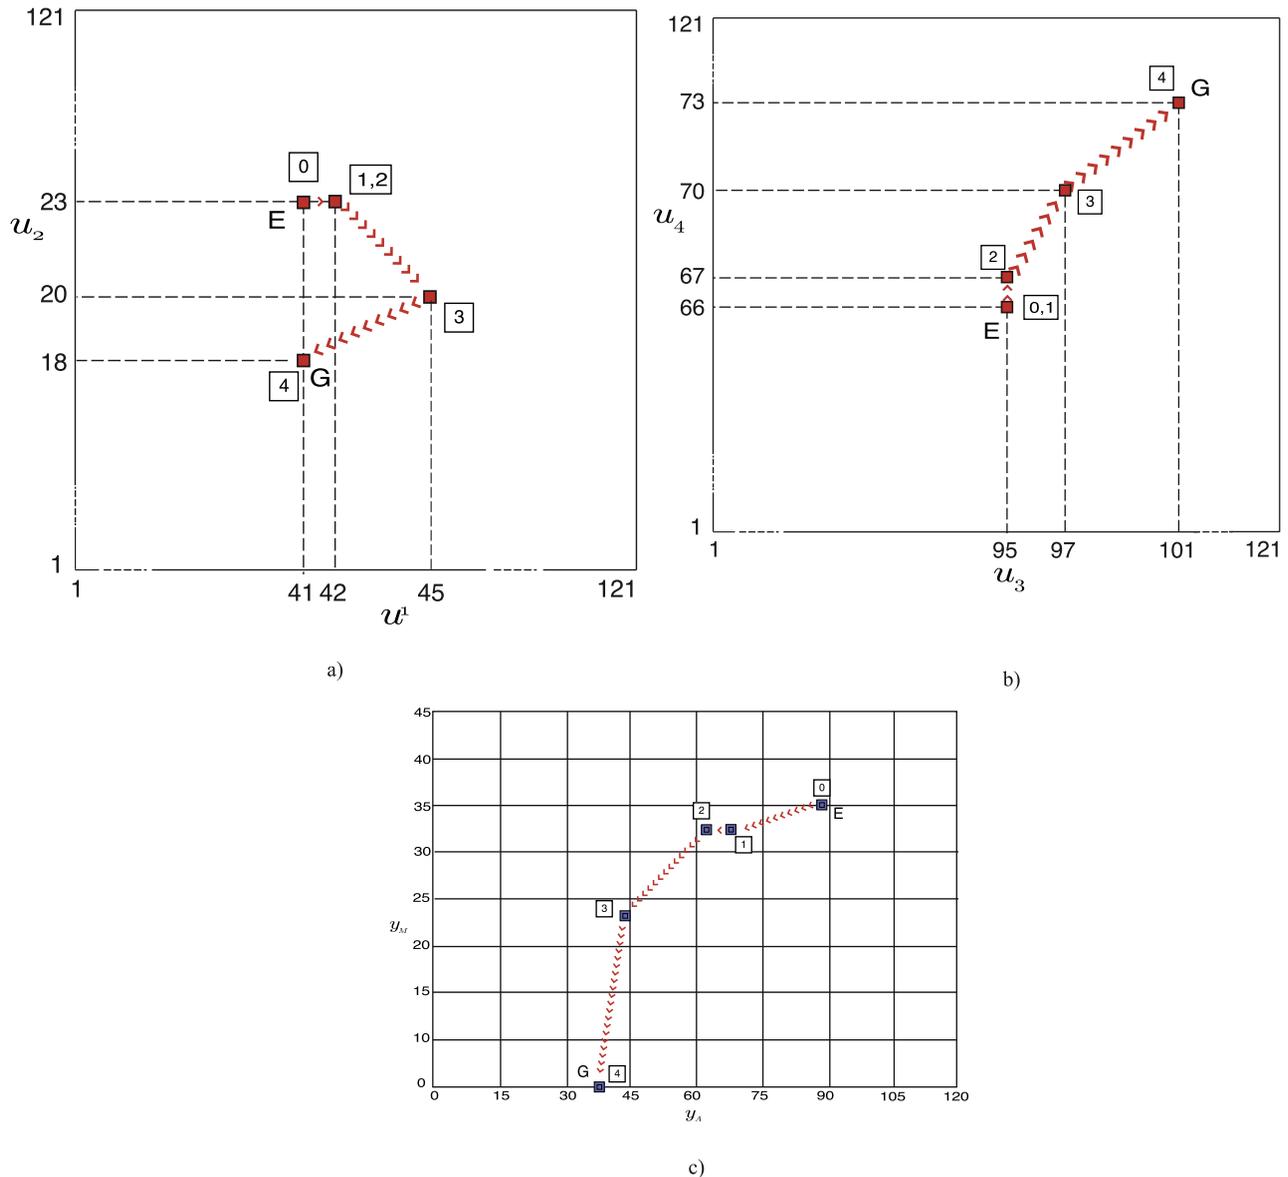


Figure 12. Problem P3/local. (a), (b) The sequence of BCAs iteratively chosen by the DM in the five iterations performed with the local approach starting from alternative E (one of the solutions to problem P2). (c) The associated performances computed via simulation of model \mathcal{M} . (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

neighborhood of point C (the final BCA), but is definitely poor for any value of y_A lower or greater than 137. Moreover, the definition of the neighborhood \mathcal{D} , wherein, at each iteration, the local procedure seeks an improvement in the BCA, appears to be critical. Large \mathcal{D} is allowed to explore a wide range of variability of the decision vector in a few iterations, whereas a small \mathcal{D} would entail a considerable number of iterations. However, with a too large \mathcal{D} , the linear local approximation is no longer valid and the tentative BCAs is refused in the termination test (Figure 15a) thus the procedure could end up with very poor solutions. With a too small \mathcal{D} , the opposite occurs (Figure 15b) but the solution of the problem requires a considerably larger number of iterations. Similar to the gradient-based search method, to be effective the local procedure requires a balance of

computational burden, which is reduced by large \mathcal{D} , and accuracy, which improves with small \mathcal{D} . As far as the global approach is concerned, the approximation is expected to be good around the sampled points but might be unreliable in other parts of the decision space (unexplored or poorly explored areas). However, if the DM selects the interesting points in one of these regions, at the subsequent iteration the RS approximation is correspondingly improved.

[58] As far as the computational costs are concerned, the number of simulations of the ELCOM-CAEDYM model required by the global approach to obtain an approximation of the entire Pareto frontier was generally greater than the time required by the local one to determine an approximate BCA. For example, the global procedure required 31 simulation runs on problem P1 and 42 on problem P2. With the

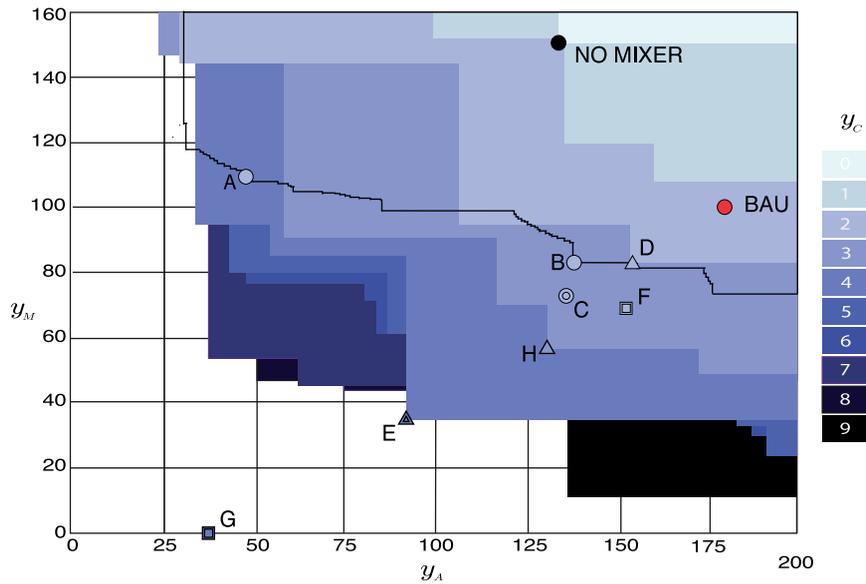


Figure 13. Summary of the results obtained with the different combinations of problems and formulations. Circles, triangles, and squares indicate the alternatives obtained from problem P1 (B and C), P2 (D and H), and P3 (F and G), respectively. The simple border is used for alternatives computed with the global approach and the double border for those obtained with the local approach. The colored decision map is for problem P2/global (see Figure 9), while the black continuous line is the decision map of problem P1/global (see Figure 5). The colors (also of circles/triangles/squares) are used for the number of mixers associated with each solution, accordingly to the color map on the right. (For interpretation of the references to color in this figure legend, the reader is referred to the online version of this article.)

local approach, this reduced to six simulations for point C, 18 for E, 15 for F, and 24 for G. However, the time required to find the final BCA for either approach was 10 to 20 times faster than using an exhaustive traditional “what-if” analysis. For example, problem P2 required nearly 240 and 90 d for the global and local approach, respectively (on an Intel Xeon 3.16 GHz QuadCore with 16GB Ram), while an exhaustive analysis would have taken about 5.5 yr. These figures have only relative value, as they will vary significantly with both the complexity of the problem under study (e.g., number and shape of objectives and constraints) and the subjective judgment of the DM. The estimates are also based only on the machine time and do not include the time required by the interaction with the DM. In the global approach this can be largely reduced, especially for the early iterations, by substituting the human DM for a model that mimics his preference. The underlying idea is that the DM should be mainly attracted by the alternatives associated to high trade-offs, i.e., those mapping into the knee of the Pareto frontier. However, though many multiobjective decision methods prioritize the search for such points [see *Rachmawati and Srinivasan, 2006* and references therein], this a priori assumption on the DM preference structure might not reflect his actual way of taking decisions. Conversely, the involvement of the DM from the very beginning, as required by the local approach, does not imply any assumption on the DM preference, even if the less-accurate Pareto frontiers shown at first iterations can constitute a (misleading) reference point that unduly influences all subsequent evaluations (the so-called *anchoring bias* [*Tversky and Kahneman, 1974*]). To reduce this bias effect, techniques exist to iden-

tify promising solutions that may have been neglected by the DM [e.g., *Babbar-Sebens and Mukhopadhyay, 2009*]. Finally, both approaches require the DM to evaluate a high number of alternatives, possibly leading to the problem of user fatigue [e.g., *Singh et al., 2010* and references therein].

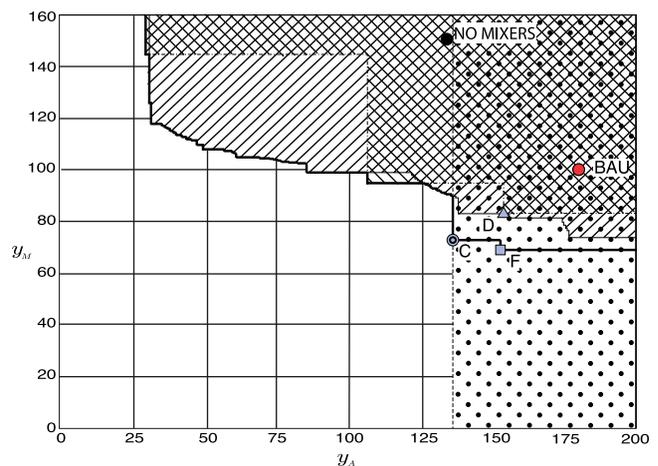


Figure 14. The decision maps of problem P1/global (continuous thin line, see Figure 5) and problem P2/global for $\bar{r} = 2$ (dot-dashed line, see Figure 9). The unknown Pareto frontier is, by definition, not dominated by the union of these maps and by points C and F obtained with the local approach and $\bar{r} = 2$ (see bold line). The dashed line is the decision map of problem B/local (see also Figure 7b).

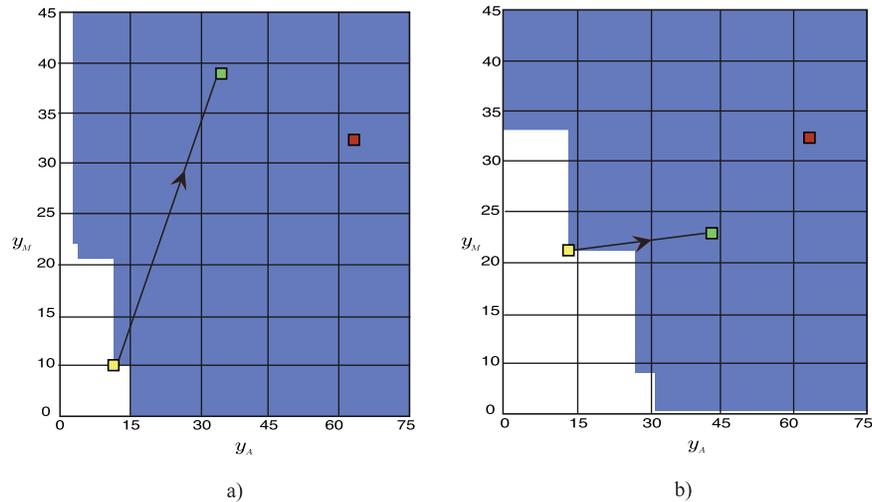


Figure 15. Problem P3/local starting from E. The effect of different sizes of the set \mathcal{D} on the acceptance of the tentative new BCA in the termination test: (a) when \mathcal{D} is defined by varying the position of each mixer of up to ± 4 cells, the tentative new BCA is refused; (b) On the contrary, when the variation is only up to ± 3 cells, the tentative BCA is accepted.

5.2. On the Definition of the Decision Variables

[59] The final mixer positions obtained with the different combinations of problem formulations and solution approaches permits some comments. Moving the mixers upstream results in a positive effect on the cyanobacteria concentration (compare the positions of A and B or F and B in Figure 16a and their performances in Figure 13). Increasing the distance between two mixers improves the manganese indicator (compare the BAU alternative with C or D and F). As a logical consequence, the installation of more mixers should improve both indicators as the same time. This is confirmed by alternatives E and G (Figures 13 and 16b). Moreover, the big improvement associated with G seems to be connected to the even spatial distribution of the mixers. This intuition is supported by the fact that the alternative corresponding to three mixers evenly spaced (point H in Figure 16b) is among the alternatives temporary considered by the DM in problem P2, both global and local, and eventually discarded in favor of alternative E (see Figure 13). If the DM were unwilling to pay for one additional mixer, alternative H would have been chosen as the final BCA. Interestingly, solution G is not among the feasible solutions to problem P2 because, with three macro areas and the positioning rule adopted, a solution with four evenly spaced mixers is not feasible. This suggests that the decision vector formulated in terms of the coordinates of a given number of mixers (formulation F1 in what follows, see section for its formal definition) should be preferred over formulating the decision vector as the number of mixers to be located accordingly to a predefined positioning rule (formulation F2). On the other hand, formulation F1 implies a large feasible alternative set and thus a very large number of simulations are required at each iteration of the global procedure or many iterations are necessary for the local procedure to converge. Therefore, the cascade combination of the two approaches and the two formulations appears to be the best option: The global approach is used to determine a good initial BCA using formulation F2 (i.e., solving

problem P2); the local approach is subsequently run with formulation F1 (i.e., solving P3) to refine such solution. This is the novel cascade approach we propose as outlined above.

6. Concluding Remarks

[60] In this paper we compared the global and local approach to RS approximation on a real example of remediation in a water supply reservoir. Besides the operational, case-study-specific value of the approach, the paper's aim was to provide some general guidelines about the selection of the more appropriate combination of approach and problem formulation, and about the definition of the decision variables. These can be summarized in the following points.

6.1. Local Approach

[61] *pros*: When the starting point is sufficiently good, the final BCA is generally better than that found with the global approach.

[62] *pros*: Generally, it is computationally less intensive than its global twin.

[63] *cons*: The Pareto frontier approximation is reliable only in a neighborhood of the final BCA, which basically means that it provides only one Pareto-efficient solution.

[64] *cons*: Choosing the set \mathcal{D} , wherein the linear approximation is acceptable, is not straightforward and might require several experiments. This results in an increase of the computational costs, which are, however, unavoidable in order to not end up with a poor solution.

[65] *cons*: The quality of the final BCA strongly depends upon the quality of the initial alternative available (usually the BAU alternative).

6.2. Global Approach

[66] *pros*: Large ranges of variation of the decision vector can be easily explored.

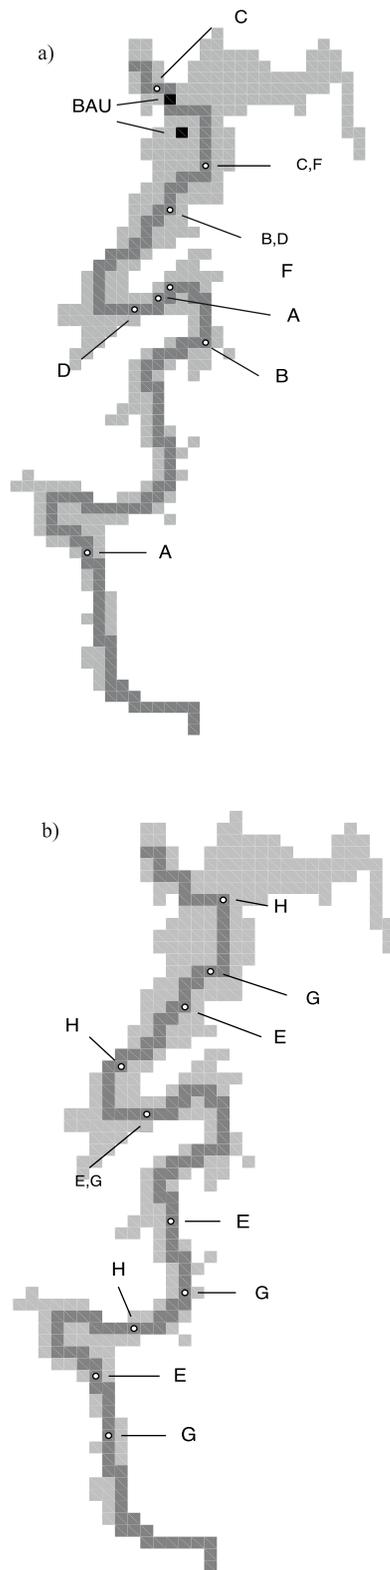


Figure 16. The mixer positions according to the BCAs determined with the different problems/approaches, with (a) two mixers and (b) three/four mixers.

[67] *pros*: It provides an approximation of the entire Pareto frontier, which is, by construction, more accurate around the objective values the DM is more interested in (i.e., those associated to the “interesting” alternatives).

[68] *cons*: Generally, the computational burden is higher than with the local approach and grows rapidly with the dimensionality of the decision vector.

[69] *cons*: The only way to mitigate the computational burden is to adopt some mixer positioning rules, however, this often conditions the final BCA in an unpredictable way.

6.3. Formulation F1

[70] *pros*: The final BCA does not depend upon any a priori positioning rules, which might be problem dependent and require some expert domain interaction.

[71] *cons*: The mixer number is given and the approach must be reiterated to explore the mixer number dimension. This might be computationally very intensive.

6.4. Formulation F2

[72] *pros*: It is computationally more convenient.

[73] *cons*: The a priori positioning rule can condition the final BCA excluding potentially interesting solutions.

[74] In conclusion, the best option appears to be the cascade approach: formulation F2 (i.e., problem P2) and the global approach first and formulation F1 (i.e., problem P3) with the local approach for refinement.

[75] **Acknowledgments.** The ELCOM-CAEDYM calibration used in this study was funded by ActewAGL under the project “Murrumbidgee to Googong Raw Water Transfer Project—Hydrodynamic Modeling.” The CWR reference 2408-AC was created from this paper.

References

- Babbar-Sebens, M., and S. Mukhopadhyay (2009), Reinforcement Learning for Human-Machine Collaborative Optimization: Application in Ground Water Monitoring, in *Proceedings of IEEE Systems, Man and Cybernetics Conference*, pp. 3563–3568, Nice, France.
- Branke, J., K. Deb, K. Miettinen, and R. Slowinski (Eds.) (2008), *Multiobjective optimization, Interactive and evolutionary approaches, Lecture Notes in Computer Science*, vol. 5252, Springer, Berlin-Heidelberg, Germany.
- Broad, D. R., G. C. Dandy, and H. R. Maier (2005), Water distribution system optimization using metamodels, *J. Water Resour. Plann. Manage.*, 131(3), 172–180.
- Castelletti, A., and R. Soncini-Sessa (2006), A procedural approach to strengthening integration and participation in water resources planning, *Environ. Modell. Software*, 21(10), 1455–1470.
- Castelletti, A., F. Pianosi, and R. Soncini-Sessa (2008), Integration, participation and optimal control in water resources planning and management, *Appl. Math. Comput.*, 206(1), 21–33.
- Castelletti, A., F. Pianosi, R. Soncini-Sessa, and J. P. Antenucci (2010a), A multi-objective response surface approach for improved water quality planning in lakes and reservoirs, *Water Resour. Res.*, 45, W06502, doi:10.1029/2009WR008389.
- Castelletti, A., A. Lotov, and R. Soncini-Sessa (2010b), Visualization-based multi-objective improvement of environmental decision-making using linearization of response surfaces, *Environ. Modell. Software*, 25, 1552–1564.
- Castelletti, A., S. Galelli, M. Ratto, R. Soncini-Sessa, and P. C. Young (2011), A general framework for dynamic emulation modelling in environmental problems, *Environ. Modell. Software*, in press.
- Espinoza, F. P., and B. S. Minsker (2006), Development of the enhanced self-adaptive hybrid genetic algorithm (e-SAHGA), *Water Resour. Res.*, 42, W08501, doi:10.1029/2005WR004221.
- Hillmer, I., R. Alexander, and P. S. Yeates (2008), Murrumbidgee to Googong raw water transfer project—Hydrodynamic modelling, *Tech. Rep. WP2185*, Centre for Water Research, University of Western Australia, Crawley, WA.
- Hipsey, M. R., J. R. Romero, J. P. Antenucci, and J. Imberger (2007), The Computational Aquatic Ecosystem Dynamics Model (CAEDYM): A Versatile Water Quality Model for Coupling with Hydrodynamic

- Drivers, in *Proceedings of the 7th International Conference on Hydroinformatics*, pp. 526–533, Nice, France.
- Hodges, B. R., J. Imberger, A. Saggio, and K. B. Winters (2000), Modeling basin-scale internal waves in a stratified lake, *Limnol. Oceanogr.*, *45*(7), 1603–1620.
- Huisman, J., J. Sharples, J. M. Stroom, P. M. Visser, W. E. A. Kardinaal, J. M. H. Verspagen, and B. Sommeijer (2004), Changes in turbulent mixing shift competition for light between phytoplankton species, *Ecology*, *85*(11), 2960–2970.
- Jamshid Mousavi, S., and M. Shourian (2010), Adaptive sequentially space-filling metamodeling applied in optimal water quantity allocation at basin scale, *Water Resour. Res.*, *46*, W03520, doi:10.1029/2008WR007076.
- Jönk, K. D., J. Huisman, J. Sharples, B. Sommeijer, P. M. Visser, and J. M. Stroom (2008), Summer heatwaves promote blooms of harmful cyanobacteria, *Global Change Biol.*, *14*(3), 1365–2486.
- Kollat, J. B., and P. Reed (2007), A framework for visually interactive decision-making and design using evolutionary multi-objective optimization (VIDEO), *Environ. Modell. Software*, *22*(12), 1691–1704.
- Kourakos, G., and A. Mantoglou (2008), Remediation of heterogeneous aquifers based on multiobjective optimization and adaptive determination of critical realizations, *Water Resour. Res.*, *44*, W12408, doi:10.1029/2008WR007108.
- Lotov, A. V. (1989), *Methodology and software for interactive decision support*, *Lecture Notes in Economics and Mathematical Systems*, vol. 337, Generalized reachable sets method in multiple criteria problems, pp. 65–73, Springer, Berlin, Germany.
- Lotov, A., V. A. Bushenkov, and G. K. Kamenev (2004), *Interactive Decision Maps: Approximation and Visualization of Pareto Frontier*, 336 pp., Kluwer, Boston, MA.
- Lotov, A. V., L. V. Bourmistrova, R. V. Efremov, V. A. Bushenkov, A. L. Buber, and N. B. Brainin (2005), Experience of model integration and pareto frontier visualization in the search for preferable water quality strategies, *Environ. Modell. Software*, *20*(2), 243–260.
- Madetoja, E., H. Ruotsalainen, V.-M. Monkkonen, J. Hamalainen, and K. Deb (2008), Visualizing multi-dimensional Pareto-optimal fronts with a 3D virtual reality system, in *Proceedings of the International Multiconference on Computer Science and Information Technology*, pp. 907–913, Wisla, Poland.
- Morillo, S., J. Imberger, J. P. Antenucci, and D. Copetti (2009), Using impellers to distribute local nutrient loadings in a stratified lake: Lake Como, Italy, *J. Hydraul. Eng.*, *135*(7), 564–574.
- Myers, R., and D. Montgomery (1995), *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*, 728 pp., John Wiley & Sons, NY.
- Rachmawati, L., and D. Srinivasan (2006), A multi-objective genetic algorithm with controllable convergence on knee regions, in *Proceedings of IEEE Congress on Evolutionary Computation*, pp. 1916–1923, Vancouver, Canada.
- Rodriguez, H. G., J. Popp, C. Maringanti, and I. Chaubey (2011), Selection and placement of best management practices used to reduce water quality degradation in Lincoln Lake watershed, *Water Resour. Res.*, *47*, W01507, doi:10.1029/2009WR008549.
- Schindler, D. W., R. E. Hecky, D. L. Findlay, M. P. Stainton, B. R. Parker, M. J. Paterson, K. G. Beaty, M. Lyng, and S. E. M. Kasian (2008), Eutrophication of lakes cannot be controlled by reducing nitrogen input: Results of a 37-year whole-ecosystem experiment, *Proc. Natl. Acad. Sci. U. S. A.*, *105*(32), 11:254–11:258.
- Schladow, S. G. (1993), Lake destratification by bubble-plume systems: Design methodology, *J. Hydraul. Eng.*, *119*(3), 350–368.
- Singleton, V. L., and J. C. Little (2006), Designing hypolimnetic aeration and oxygenation systems: A review, *Environ. Sci. Technol.*, *40*(24), 7512–7520.
- Singh, A., B. S. Minsker, and P. Bajcsy (2010), Image-based machine learning for reduction of user-fatigue in an interactive model calibration system, *J. Comput. Civil Eng.*, *24*(3), 241–251.
- Tversky, A., and D. Kahneman (1974), Judgment under uncertainty: Heuristics and biases, *Science*, *185*, 1124–1130.
- Yan, S., and B. Minsker (2006), Optimal groundwater remediation design using an adaptive neural network genetic algorithm, *Water Resour. Res.*, *42*, W05407, doi:10.1029/2005WR004303.

J. P. Antenucci, Centre for Water Research, University of Western Australia M023, 35 Stirling Highway, Crawley, 6009 Western Australia, Australia.

A. Castelletti, D. Limosani, X. Quach Thi, and R. Soncini-Sessa, Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci, 32, 20133 Milano, Italy. (castelle@elet.polimi.it)