

1 **PLANNING THE OPTIMAL OPERATION OF A**
2 **MULTI-OUTLET WATER RESERVOIR WITH WATER**
3 **QUALITY AND QUANTITY TARGETS**

4 Andrea Castelletti¹, M. ASCE; Hiroshi Yajima²; Matteo Giuliani³;
 Rodolfo Soncini-Sessa⁴; and Enrico Weber⁵

5 **ABSTRACT**

6 The integration of quality and quantity issues in the management of water resources
7 systems is key to meet society’s long-term needs for freshwater while maintaining essential
8 ecological services and economic benefits. Current water management practices are mostly
9 targeted towards quantitative uses and quality is usually addressed separately as an in-
10 dependent problem. One of the reasons for the lack of integration lies in the inadequacy
11 of optimization techniques nowadays available to cope with the large, distributed, simu-
12 lation models adopted to characterize the coupled ecological and biochemical processes in
13 water bodies. In this paper we propose a novel approach based on the conjunctive use of a
14 batch-mode Reinforcement Learning algorithm and a 1D coupled hydrodynamic-ecological
15 model to design the optimal operation of a multipurpose water reservoir accounting for both
16 quantity and quality targets. We consider up to five operating objectives, including both
17 in-reservoir and downstream water quality parameters, and design efficient operating poli-
18 cies conditioned upon not only the current storage but also water characteristics, such as

¹Assistant Professor, Dept. Electronics and Information, Politecnico di Milano, P.za Leonardo da Vinci, 32, 20133 Milano, Italy. E-mail: castelle@elet.polimi.it.

²Associate Professor, Dept. Management of Social Systems and Civil Engineering, Tottori University, Koyama, Tottori 680-8552, Japan. E-mail: yajima@cv.tottori-u.ac.jp.

³PhD student, Dept. Electronics and Information, Politecnico di Milano, P.za Leonardo da Vinci, 32, 20133 Milano, Italy. E-mail: giuliani@elet.polimi.it.

⁴Professor, Dept. Electronics and Information, Politecnico di Milano, P.za Leonardo da Vinci, 32, 20133 Milano, Italy. E-mail: soncini@elet.polimi.it.

⁵Research Associate, Dept. Electronics and Information, Politecnico di Milano, P.za Leonardo da Vinci, 32, 20133 Milano, Italy. E-mail: weber@elet.polimi.it.

19 temperature and total suspended solids at different depths. The approach is applied to a
20 real world case study in Japan consisting of a water reservoir, Tono Dam, equipped with
21 a selective withdrawal structure and used for flood protection, power generation, irrigation
22 and recreational purposes. Results show that a potential control over in-reservoir and down-
23 stream water quality can be gained without impairing the hydraulic capacity of the reservoir
24 by effectively exploiting - through the operating policy - the operational flexibility provided
25 by the selective withdrawal structures.

26 **Keywords:** selective withdrawal systems; reservoir operation; water quality; reinforcement
27 learning; optimization

28 INTRODUCTION

29 Dealing with a scarce resource, traditional reservoirs operating strategies are generally
30 designed to meet only the quantitative demand in river basins, e.g., agricultural supply,
31 hydropower production, and flood control. Water quality rarely competes with these water
32 uses and is usually addressed separately as an independent problem assuming the primary
33 quantity target forms constraints on the management options. As a consequence, most of the
34 lakes in the world suffer water quality deterioration and sedimentation to the point where the
35 primary storage functions are being impaired (ILEC 2005). Poor understandings of aquatic
36 and riparian ecosystems and the lack of quantification of social and environmental objectives
37 are some of the possible reasons for the inadequate development of water-quality based man-
38 agement strategies, taking advantage of the close relationships and synergies between quality
39 and quantity (Dhar and Datta 2008). Another, more technical reason is the lack of mathe-
40 matical tools capable of effectively combining the complexity of the distributed parameter,
41 physically-based models usually adopted to describe hydrodynamics and biochemical condi-
42 tions of water bodies (for a review, see Ostfeld and Salomons (2005)) and the computational
43 burden of rational, optimization based, decision-making (Castelletti et al. 2011).

44 As reported in Dortch (1997), three are the options available to affect in-reservoir and
45 downstream water quality in water systems: *i*) pre-treatment or control of reservoir inflows

46 (e.g., upstream settling basins, some type of watershed control or land management activ-
47 ity); *ii*) in-reservoir management or treatment techniques (e.g., destratification, aeration,
48 dilution, etc.); *iii*) management of reservoir outflows (e.g., controlling the outflow rate, out-
49 flow location and timing or releases treatments). In this paper, we focus on the third option
50 and consider the case in which the control on the outflow is performed by equipping the reser-
51 voir with a multilevel intake, or selective withdrawal system (SWS), that allows releasing
52 water at different depths with different physico-chemical properties (Bohan and Grace 1973;
53 Davis et al. 1987; Smith et al. 1987). The obtained flexibility in the selection of the outlet
54 offers advantages for more continuously meeting water quality goals, especially in reservoirs
55 affected by seasonal stratification, or for responding to short term events (Gelda and Effler
56 2007). SWSs were originally designed starting from the 1950s and 1960s to control the water
57 temperature in the receiving river using data collected from the reservoir of interest (Nece
58 1970). Then, numerical models have been developed to assist the design and evaluation of
59 these infrastructures in order to predict the effects on water quality and design water quality
60 targeted operation in new and existing reservoirs (Vermeyen et al. 2003).

61 Water quality objectives can be broadly divided into in-reservoir and downstream ob-
62 jectives. Most of the works proposed in the literature consider only downstream objectives
63 and one single water quality parameter at a time. The effects of different release strategies
64 for single-outlet dams have been widely studied for the control of the outflow salinity (Or-
65 lob and Simonovic 1982; Dandy and Crawley 1992; Nandalal and Bogardi 1995), turbidity
66 (Ikebuchi and Kojiri 1992), and specific ecological parameters as BOD, DO, or TOC (Dhar
67 and Datta 2008; Westphal et al. 2003). Multi-outlet reservoirs studies concentrated on the
68 problem of maintaining a desired discharge temperature (Fontane et al. 1981; Hanna et al.
69 1999; Sherman 2000; Hanna and Saito 2001; Gelda and Effler 2007; Baltar and Fontane
70 2008), which is key to preserve fish habitat (Vermeyen 1999) but is also an important factor
71 affecting irrigation and pollution control (Fontane et al. 1981). The control of downstream
72 temperature is not the only controllable variable and SWSs can be also exploited in order to

73 influence sedimentation and ecological discharge parameters (Hayes et al. 1998; Kerachian
74 and Karamouz 2006).

75 So far, relatively few studies have considered in-reservoir quality targets and, to the au-
76 thors' knowledge, optimization-based approaches are nearly unexplored. Ferris and Lehman
77 (2007) identified the causal factors affecting algal bloom and modeled this phenomenon to
78 test the possibility of influencing algal growth with few ad-hoc alternatives. A similar ap-
79 proach was developed by Lee and Guy (2012) to assess the room for reducing in-reservoir
80 sedimentation through outflow management. Simplifications such as the steady-state as-
81 sumption were introduced to control eutrophication (Kuo et al. 2008; Huang et al. 2012).
82 Alternative approaches are based on heuristic optimization algorithms, like Genetic Algo-
83 rithms (GAs): Khan et al. (2012) developed a Reservoir Optimization-Simulation with Sedi-
84 ment Evacuation model based on GAs to design management strategies considering sediment
85 evacuation in addition to classical quantity objectives as irrigation supply and hydropower
86 generation. Chaves and Kojiri (2007) adopted a Stochastic Fuzzy Neural Network approach
87 to obtain a quasi-optimal solution considering both quantity and quality objectives.

88 A general methodology for designing the optimal operation of SWSs considering multiple
89 water quantity and quality objectives, both in-reservoir and downstream, is still needed. In
90 this work, a novel approach is proposed for designing Pareto-optimal operating policies for
91 SWS reservoirs. The approach exploits the feedback between the selection of outlet loca-
92 tions for water quantity demand and the water quality patterns within the reservoir (Gelda
93 and Effler 2007), making it possible to satisfy downstream objectives and, simultaneously,
94 affect in-reservoir water quality. In particular, we adopt a Reinforcement Learning (RL) ap-
95 proach combined with a 1D coupled hydrodynamic-ecological model of the lake (DYRESM-
96 CAEDYM) to design Pareto-optimal operating policies conditioned upon some key informa-
97 tion on the current conditions of the lake as captured by the model. Unlike simulation-based
98 optimization methods (e.g., Implicit Stochastic Optimization (Labadie 2004) or Parametriza-
99 tion Simulation and Optimization (Koutsoyiannis and Economou 2003)) that do not offer any

100 performance guarantee and proof of convergence, the proposed approach is an approximation
101 of traditional Dynamic Programming (DP) and as such ensures some anticipated favorable
102 properties of the policy obtained (Powell 2007). More precisely, we use a batch-mode al-
103 gorithm, called fitted Q-iteration (Ernst et al. 2005; Castelletti et al. 2010), that allows
104 learning the operating policy offline on a sample data-set constituted of observational data
105 and/or the outputs of simulated experiments. The approach offers two important features
106 that make it particularly suitable for the high dimensional problem here considered: first,
107 the use of simulation to estimate quantities of interest, thus avoiding model-based computa-
108 tions that would make the approach inapplicable in combination with a high fidelity model
109 of the water quality processes (i.e. the so called curse of modeling (Tsitsiklis and Van Roy
110 1996)); second, it adopts a non-parametric function approximation (Ernst et al. 2005) of the
111 value function and thus considerably mitigates the curse of dimensionality associated with
112 DP based or derived approaches.

113 The approach is demonstrated on Tono Dam (Japan), an artificial reservoir constructed
114 for flood protection, power generation, supplying agricultural water downstream and recre-
115 ational purposes. Due to the region’s climate, the lake is characterized by prolonged periods
116 of stratification that negatively impact the water quality both in-reservoir and in the lake’s
117 outflow. With the purpose of conjunctively controlling water quality and quantity, the dam
118 was equipped with a SWS. In this study, the SWS operation is optimized with respect to
119 five objectives and the Pareto frontier of the problem is computed. The operating policies
120 corresponding to the extremes of the front are first analyzed in order to gain insight on the
121 strategy adopted by the algorithm for each single objective separately. Then a compromise
122 policy that simultaneously considers all the objectives is analyzed.

123 The paper is organized as follows: the Tono Dam case study is firstly described, followed
124 by the presentation of the methodology. Results and discussion are then reported, while
125 final remarks along with issues for further research are presented in the last section.

126 **SYSTEM DESCRIPTION**

127 **Tono Dam**

128 Tono Dam is located at the confluence of Kango and Fukuro rivers (Figure 1a), in the
129 western part of Japan. The construction works were completed in 2011. With a height of
130 75 m (Figure 1b), the dam forms an impounded reservoir of $12.4 \times 10^6 \text{ m}^3$ (gross capacity),
131 with a surface area of 0.64 km^2 and fed by a 38.1 km^2 catchment.

132 The reservoir was primarily constructed to support irrigated agriculture, for flood control
133 and recreational purposes, and is also connected to a small hydropower plant (1.1 MW
134 installed capacity). Due to the region's climate, the lake is characterized by prolonged
135 periods of stratification that produce negative effects on water quality both in-reservoir
136 and in the lake's outflow. With the purpose of conjunctively controlling water quality and
137 quantity, the dam was equipped with a selective withdrawal system (SWS) constituted by
138 a rack of 15 vertically stacked siphons (Figure 1c) allowing to release water from the active
139 storage at different depths. Siphons are operated by inflating or deflating air, and blending
140 is allowed: the total amount of water released through the SWS is equally divided among
141 the active siphons. Floods are only controlled using a flood orifice gate at elevation 182.8
142 m a.s.l. that operates on the flood control volume. Selective release is not available in the
143 sediment storage, however two more siphon gates are equipped below 156 m to release water
144 in winter period (from December to March) or to release the minimum environmental flow
145 in particularly dry periods, when the water level drops below the lower bound of the active
146 storage. In normal conditions between April and October, the minimum environmental flow
147 is guaranteed through the top 15 siphons, when the level drops below the SWS lower limit,
148 the sediment outlet is activated.

149 **Social, economic and environmental issues**

150 While one of the main purposes of Tono dam operation is to provide water for irrigation,
151 the SWS might have an impact on several other water uses. We distinguish between *in-*
152 *reservoir* and *downstream* issues, the former being affected by level variations, the latter by
153 the release.

155 Too low reservoir *levels*, which can be generated in the attempt to release water to satisfy
156 agricultural water demand, can potentially reduce the recreational value of the lake. In order
157 to emphasize this recreational interest, the SWS management has to consider to keep the
158 lake level as close as possible to a reference level of 182.8 m a.s.l. as the normal high water
159 level. This, however, implies stocking a significant volume of water in the reservoir with
160 potentially negative effects both in-reservoir, e.g., boosting algal blooms, and downstream,
161 e.g., water shortages.

162 Odors and unattractive appearance of *algal blooms* can detract from the recreational value
163 of the lake affecting the quality of the water stored in the reservoir. The physical processes
164 driving the bloom of algae are particularly complex. However, thermal stratification has a
165 dominant role. Controlling the temperature profile is a mechanical way of controlling the
166 depth of nutrient load intrusion and therefore the algae bloom, which is basically sensitive to
167 the available light, is stimulated by an intrusion in the layer of shallow stratification (Yajima
168 et al. 2006). Moreover the temperature profile might vary as a consequences of withdrawing
169 at different levels (Gelda and Effler 2007). Generally, the deeper the withdrawal the more
170 the deepening of the thermocline. Yet, this implies releasing colder water with potentially
171 negative effects downstream and might affect sedimentation in the way explained below.

172 High levels of in-reservoir *sedimentation* can remarkably reduce the reservoir life by induc-
173 ing the rapid silting of the impoundment. Sedimentation is basically driven by the inflow and
174 re-suspension can be assumed as negligible considering the reservoir depth (Evans 1994). In
175 particular, inflow intrusion is governed by the in-reservoir temperature profile and the inflow
176 temperature because floods are more likely to intrude just above the thermocline (Yajima
177 et al. 2006). Therefore, to maximize sediment evacuation, the release should be set at the
178 depth at which the turbid inflow is intruding and then, if necessary, dynamically moved
179 to the deeper siphons to intercept the maximum concentration of suspended solids not yet
180 evacuated. Moreover, some recent studies (Yajima et al. 2006) have shown that using the

181 top siphon combined with the spillways leads the inflow to the shallower depth and facilitate
182 sediment flushing from the spillway. These ways of operating the SWS might have nega-
183 tive effects on the other sectors, like, for instance, the ecosystem downstream, which might
184 be damaged by too warm water. Also recreation could be affected, since by keeping the
185 thermocline in the shallow layer, algal blooms are more likely to occur as explained above.

186 *Downstream*

187 Farmers are interested in reducing the water supply deficit, which has a direct effect on
188 the seasonal harvest and, therefore, on the annual income, which is the criterion through
189 which the farmers judge the level of attractiveness of an operating policy (Hashimoto et al.
190 1982).

191 The riverine ecosystem downstream from the dam is potentially threatened by large de-
192 viations of the water *temperature* from the seasonal natural patterns that might negatively
193 affect faunal richness in both fishes and invertebrates (Hanna and Saito (2001) and refer-
194 ences therein). According to Fontane et al. (1981) and Baltar and Fontane (2008), a simple
195 and physically rooted criterion to reduce the effect of artificially induced temperature varia-
196 tions is to force the outflow temperature to be as closest as possible to the (natural) inflow
197 temperature.

198 **MATERIALS AND TOOLS**

199 Planning efficient operating rules for the SWS based solely on the indications reported be-
200 forehand might turn out particularly difficult in this absence of quantitative references from
201 the historical operation (the dam has been just constructed). Moreover, while potentially
202 effective strategies can be anticipated for most of the involved issues separately considered,
203 their interaction and the associated conflicts make it hard to empirically formulating ad-
204 equately balanced rules. In this study, we adopt a batch-mode Reinforcement Learning
205 approach to design Pareto-optimal feedback operating policies for the SWS. The operating
206 policy is computed by repeatedly solving a regression problem on a data-set of one-step
207 transitions of the reservoir system generated by multiple simulations of a physically-based

208 coupled hydrodynamic-ecological model of Tono Dam under different external driver and
 209 release decision scenarios. The procedure adopted is described in the flowchart of Figure 2
 210 and its building blocks are described next.

211 **Problem formulation**

212 Given the current system conditions as described by the state vector $\mathbf{x}_t \in \mathbb{R}^m$ (e.g.,
 213 storage, temperature, suspended solid), a daily feedback operating policy p for the SWS
 214 returns the volume $u_t^i = m_t(\mathbf{x}_t)$, $i = 1, \dots, n$, to be released over the time interval $[t, t + 1)$,
 215 i.e. the next 24 hours, from each of the n SWS outlets and for each day t . The problem of
 216 designing the set of Pareto-optimal policies can be formulated as an optimal control problem
 217 of a dynamic system evolving according to a model $\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1})$, controlled by a
 218 vector $\mathbf{u}_t \in \mathcal{U}_t(\mathbf{x}_t) \subseteq \mathbb{R}^n$ of n feasible decisions, and affected by l stochastic external drivers
 219 $\boldsymbol{\varepsilon}_{t+1} \in \mathbb{R}^l$ (e.g., inflow, wind, solar radiation, nutrient load), i.e.

$$p^* = \arg \min_p \boldsymbol{\lambda} \cdot \mathbf{J}(p) \quad (1)$$

220 in which $\mathbf{J}(p) = [J^1, \dots, J^q]$ is the vector objective function and $\boldsymbol{\lambda}$ is the vector of weights
 221 with $\sum_{k=1}^q \lambda^k = 1$ and $\lambda^k \geq 0 \forall k$. The k -th objective is formulated as the expected total
 222 discounted cost over an infinite horizon (see, for more details, Castelletti et al. (2008))

$$J^k = \lim_{h \rightarrow \infty} \mathbf{E}_{\boldsymbol{\varepsilon}_1, \dots, \boldsymbol{\varepsilon}_h} \left[\sum_{t=0}^{h-1} \gamma^t g_{t+1}^k(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1}) \right] \quad (2)$$

223 where $g_{t+1}^k(\cdot)$ is the k -th immediate cost function associated to each system transition and
 224 γ is a discount factor ($0 < \gamma \leq 1$).

225 By reformulating and solving the problem for different values of the weights $\boldsymbol{\lambda}$, a finite
 226 subset of the generally infinite Pareto-optimal policy set is obtained (Weighting Method).

227 **The DYRESM-CAEDYM model**

228 The characterization of the system dynamics $f_t(\cdot)$ involves the description of the main
229 hydrodynamic and ecological processes arising in the reservoir. In principle, a 3D spatially-
230 distributed model (e.g., ELCOM-CAEDYM (Yajima et al. 2006)) would be the best choice in
231 terms of accuracy and physically meaningful description of the involved processes. However,
232 the reservoir is being created by damming two rivers in a quite narrow and steep section of
233 their course and vertical phenomena are dominating. Therefore, a simpler 1D model might
234 be an acceptable surrogate. By working with a 1D model, a full characterization of the
235 spatial dynamics between the inlet and the outlet of the reservoir is lost, however the run-
236 to-real time ratio drops off to nearly 1/12300 from the 1/30 of an equivalent 3D model. Yet,
237 compared with the simple lumped models traditionally used for reservoir policy design, 1D
238 models still have a very high number of state variables, which constitutes the main limitation
239 for their inclusion within a classical optimization framework.

240 In this study we adopted the 1D DYRESM-CAEDYM model developed by the Centre for
241 Water Research at the University of Western Australia (Hipsey et al. 2006; Imerito 2007).
242 The model consists of two main components: a 1D hydrodynamic model (DYRESM-Dynamic
243 Reservoir Simulation Model), including a vertical distribution of temperature, salinity and
244 density in a reservoir, and an aquatic ecosystem model (CAEDYM-Computational Aquatic
245 Ecosystem Dynamics Model), which simulates a range of biological, chemical and physical
246 processes, expressing the variables that are commonly associated with water quality (such as
247 total phosphorus, total nitrogen, chlorophyll-a, etc.). The model is based on a Lagrangian
248 architecture that models the reservoir as horizontal layers of uniform properties (i.e. tem-
249 perature and water qualities). The thickness of the layers varies in time depending on the
250 water density profile. In our model, the minimum and the maximum thickness of a layer is
251 set to 1 m and 2 m, respectively, which correspond to allow the definition of more than 30
252 layers in the Tono Dam reservoir. Twenty-one state variables are defined for each layer, for
253 a total of nearly 600 state variables (including the level). Details on the model calibration

254 can be found in Yajima et al. (2006)

255 **Operating objectives**

256 According to the multi-objective nature of the problem, an immediate cost function $g_t^k(\cdot)$
 257 is defined for each sector of interest affected by the SWS operation:

258 - *Level*: the squared positive difference of lake level with respect to the reference level $\bar{h} =$
 259 182.8 m:

$$g_{t+1}^{lev} = (\max(\bar{h} - h_{t+1}, 0))^2 \quad (3)$$

260 - *Algal bloom*: the daily average hourly maximum concentration of chlorophyll-a (Chl-a) in
 261 the see-through layer:

$$g_{t+1}^{algae} = \frac{1}{24} \sum_{\tau=1}^{24} \max_{z_\tau \in z_E} (chla_\tau(z_\tau)) \quad (4)$$

262 where $chla_\tau$ is the Chl-a concentration [$\mu\text{g/L}$] at the τ -th hour of day t , z_τ is the depth
 263 with respect to the lake surface, z_E is the see-through layer depth set at 7 m below
 264 water surface, where the thermocline is generally formed in summer.

265 - *Sedimentation*: the daily volume of sediment expelled with the release, which has to be
 266 maximized in order to reduce the silting of the reservoir and increase its expected life:

$$g_{t+1}^{sed} = TSS_{t+1}^{out} \quad (5)$$

267 where $TSS_{t+1}^{out} = \sum_{i=1}^N tss_{t+1}^i r_{t+1}^i + tss_{t+1}^{spill} r_{t+1}^{spill}$ with r_{t+1}^i [m^3/day] being the volume of
 268 water released from the i -th siphon of the SWS, tss_{t+1}^i the average total suspend solid
 269 concentration [g/m^3] in the corresponding layer, tss_{t+1}^{spill} the average total suspend solid
 270 in the layer of the spillway, and r_{t+1}^{spill} the actual release from the corresponding layer.

271 - *Irrigation*: the squared water daily deficit with respect to the agricultural water demand
 272 w_t :

$$g_{t+1}^{irr} = \beta_t (\max(w_t - (r_{t+1} - q_{t+1}^{MEF}), 0))^2 \quad (6)$$

273 where r_{t+1} is the total release from the dam (including SWS and spillway), q_{t+1}^{MEF}
 274 is the minimum environmental flow, and β_t is a time-varying coefficient taking into
 275 consideration the different relevance of the water deficit in different periods of the years.
 276 In particular, the immediate cost is elevated to the square to favour operating policies
 277 that reduce severe deficits in a single time step, while allowing for more frequent, small
 278 shortages, which cause less damage to the crop. This ensures that vulnerability is a
 279 minimum (Hashimoto et al. 1982).

280 - *Temperature*: the squared difference between the inflow and outflow temperature (as in
 281 Fontane et al. (1981) and Baltar and Fontane (2008)):

$$g_{t+1}^{temp} = (T_{t+1}^{out} - T_{t+1}^{in})^2 \quad (7)$$

282 where T_{t+1}^{out} is the average temperature in a section just downstream of dam outlet
 283 and $T_{t+1}^{in} = \frac{T_{t+1}^K a_{t+1}^K + T_{t+1}^F a_{t+1}^F}{a_{t+1}^K + a_{t+1}^F}$ with T^K and T^F being the average temperature [°C] of the
 284 inflow respectively in the Kango and Fukuro rivers, and a_{t+1}^K and a_{t+1}^F the corresponding
 285 flows.

286 **Batch-mode Reinforcement Learning**

287 To solve Problem (1) in this work we adopt a batch-mode Reinforcement Learning (RL)
 288 algorithm called fitted Q -iteration (Ernst et al. 2005; Castelletti et al. 2010). The fitted
 289 Q -iteration (FQI) combines RL concepts of off-line learning and functional approximation of
 290 the value function, from which the policy is derived, using tree-based regression. The optimal
 291 operating policy is determined on the basis of experience samples previously collected from
 292 the system or simulations thereof, i.e. a variety of system conditions experienced by the
 293 system under different combinations of release decisions and external driver realizations
 294 with the associated resulting immediate costs. Strictly, such experience is represented as a

295 finite data-set \mathcal{F} of tuples of the form $\langle t, \mathbf{x}_t, \mathbf{u}_t, t + 1, \mathbf{x}_{t+1}, g_{t+1} \rangle$, where

$$g_{t+1} = \sum_{k=1}^q \lambda^k g_{t+1}^k \quad (8)$$

296 The underlying idea of FQI is to replace the recursive solution of the Bellman equation by
 297 DP with a sequence of non-linear regressions over the data-set \mathcal{F} (see, for further details,
 298 Castelletti et al. (2010)) with the purpose of obtaining an approximation of the optimal,
 299 but uncomputable, DP solution.

300 *State reduction*

301 Although FQI alleviates the curse of dimensionality, it can handle no more than few
 302 dozens of state variables, while the 1D DYRESM-CAEDYM model embeds several hundreds
 303 (about 600). To combine the 1D model and FQI, a reduction of the state vector dimension
 304 is unavoidable. The original state vector \mathbf{x}_t is transformed into a smaller vector $\tilde{\mathbf{x}}_t \in \mathbb{R}^{\tilde{m}}$,
 305 with $\tilde{m} \ll m$, such that $\tilde{\mathbf{x}}_t$ is still significant in conditioning the release decision, but makes
 306 the control problem computationally tractable. The FQI will then work on a new data-
 307 set $\tilde{\mathcal{F}}$, containing the reduced state vector $\tilde{\mathbf{x}}_t$ instead of \mathbf{x}_t . Both formal (see the review by
 308 (Castelletti et al. 2012)) and empirical (expert based) approaches can be adopted to perform
 309 such reduction: the former are mainly based on Dynamic Emulation Modeling (DEMO)
 310 (ibidem), the latter exploit domain knowledge to identify the most interesting variables to
 311 be considered.

312 **Setting the Experiments**

313 In this section we describe the main assumptions made and the modeling solutions
 314 adopted to apply the above methodology to the Tono Dam case study.

315 *Decision Variables*

316 The SWS was planned to allow releasing at different depths and, possibly, blending water
 317 volumes with different physicochemical characteristics. This is reflected in our model by two
 318 decision variables: u^{-3} is the volume to be releases at 3 meters below the water surface, u^{-13}

319 at 13 meters. In both cases, the decision is defined with respect to the water body surface
 320 (see Figure 1b). These water depths should correspond, respectively, to the epilimnium and
 321 the hypolimnium of the stratified reservoir (Yajima et al. 2006). The vector of the decision
 322 variables $\mathbf{u}_t = [u_t^{-3}, u_t^{-13}]$ is defined over a feasibility set $\mathcal{U}_t(\mathbf{x}_t)$ that takes into account
 323 which outlets are available given the current storage, the physical constraints imposed by
 324 the siphons, and the SWS characteristics (Figure 1c). Precisely, each siphon cannot convey
 325 more than 7.353 m³/s, while the maximum flow rate allowed through the SWS outlet is
 326 13.780 m³/s. The water volume released through each siphon cannot be freely decided, but
 327 depends on the total amount released from the SWS, which is hydraulically equally divided
 328 among the open siphons. When more than one siphon is opened, each siphon cannot be
 329 operated at the maximum capacity.

330 *State Variables*

331 Reasonably, not all the state variables in the 1D DYRESM-CAEDYM model are equally
 332 relevant in the causal network linking the release decisions, and thus the operating policy,
 333 and the objectives. Some of them have little or no effect in conditioning the policy and
 334 can be removed. In this study, we used an expert-based approach to reduce the original,
 335 large dimensional state vector to a lower order vector $\tilde{\mathbf{x}}_t$ including only 5 state variables, all
 336 having a direct effect on the value of the objectives. These are the reservoir's level h_t , the
 337 temperature T_t^i and the total suspended solid TSS_t^i in the 1D model layer corresponding to
 338 the controlled outlet (with $i=-3; -13$). Observe that, since the controlled siphons depends
 339 on the position of the lake surface (Figure 1b), also the corresponding state variables are
 340 defined according to the same moving reference.

341 *Learning data-set*

342 The learning data-set $\tilde{\mathcal{F}}$ was constructed by running simulations of the 1D DYRESM-
 343 CAEDYM model over the period 1995-2006 (calibration period) under 100 different release
 344 scenarios generated pseudo-randomly with the aim of exploring the reduced state-decision
 345 space as more homogeneously as possible. The resulting data-set $\tilde{\mathcal{F}}$ is composed by 437,800

346 tuples, a dimension which can be hardly managed by the FQI algorithm as its computational
347 complexity grows more than linearly with the number of tuples $\#\tilde{\mathcal{F}}$ (i.e. $\#\tilde{\mathcal{F}} \cdot \log(\#\tilde{\mathcal{F}})$).
348 A resampling was then performed by trial-and-error to reduce the number of tuples thus
349 obtaining a reduced data-set of 4378 tuples which represents a good compromise between the
350 computational requirements and space-decision space exploration. The resampling procedure
351 preserves the decision frequencies of the large data-set and, therefore, also the homogeneous
352 exploration of the state-decision space.

353 *Benchmark*

354 Since Tono Dam was under construction at the time this paper was being prepared, there
355 was not historical reference for the resulting Pareto-optimal operating policies. However, a
356 release scenario is available for the period between 1990 and 1999. In a former study, the dam
357 construction authority (Japanese Ministry of Land, Infrastructure and Tourism) concluded
358 that 7 m depth withdrawal between April and October and 28 m depth withdrawal (the
359 lower of the two siphon gates in the sedimentation zone) between November and next year
360 March is the desirable SWS operation for Tono Dam. Therefore this operational scenario is
361 assumed as a benchmark to evaluate the Pareto-optimal operating policies computed in this
362 paper. Since it is partially overlapping the calibration period (1995-2006), only the first 5
363 years from 1990 to 1994 (validation period) have been considered in the comparison. The
364 performances of the benchmark with respect to the five considered objectives are reported in
365 Table 1. Actually, two objectives are reported for the *irrigation* sector: the first one, which
366 is considered in the optimization, is the daily quadratic water deficit defined in eq. (6);
367 since the physical meaning of this immediate cost is hardly interpretable, a second objective
368 defined as the daily water deficit along the year is reported to support policy evaluation, but
369 it is not considered in the optimization. It is worth noting that the benchmark operation
370 assumes the *irrigation* as the main objective and, indeed, the daily water deficit is nearly
371 insignificant ($0.028 \text{ m}^3/\text{s}$). Consequently the room for further improvements of this objective
372 is limited.

373 RESULTS

374 In this section the operating policies obtained by solving Problem 1 for 50 different
375 combinations of the weights are evaluated through simulation (dashed path in Figure 2)
376 over the period 1990-1994 (validation period). Results are evaluated in three steps: first,
377 the approximate 5D Pareto-front is analyzed to explore trade-offs, conflicts, and correlations
378 among the objectives; second, the extremes points of the front, i.e. the policies obtained
379 by setting to zero all the weights in eq. (8) but one, are considered to assess the individual
380 SWS operation strategies; third, the Pareto front is explored in order to find out possible
381 interesting compromise solutions between the five conflicting objectives.

382 5-D Pareto front

383 The 5-D Pareto front is represented through projections in the space of the *algae* and
384 *sedimentation* objectives in Figure 3 and Figure 4. In particular, Figure 3 shows that the
385 best performing alternatives with respect to *algae* and *sedimentation* (in the bottom-left
386 part of the figure) negatively impact on the *level* objective (very small circles): the opti-
387 mal operation for the first two criteria tends to release large amount of water to flush out
388 both nutrients and sediments producing a drawdown of lake level. As anticipated, *algae*
389 and *sedimentation* objectives are only partially conflicting (bottom-left part of the figure):
390 sediment evacuation is maximized by keeping the thermocline in the shallow layer, which is
391 a favorable condition for algal blooms too. The best alternatives for *algae* performs fairly
392 well also with respect to the *temperature* objective (blue circles) and, therefore, these two
393 objectives seem to be not in conflict. Also observe that all the Pareto-optimal alternatives
394 significantly outperform the performance of the benchmark with respect to the *algae* objec-
395 tive. Moreover, the benchmark is poorly performing with respect to *sedimentation* and *level*
396 objective. Finally, the benchmark performance on the *temperature* objective is intermediate
397 with respect to the range of variability obtained with the Pareto optimal alternatives.

398 The performances of the *irrigation* objective are reported in Figure 4 against *sedimen-*
399 *tation* and *algae* (this latter can be considered as representative also of the *temperature*, as

400 explained above). The best performing alternatives on *algae (temperature)* and *sedimenta-*
 401 *tion* correspond to high values of irrigation deficit (yellow circles), meaning that the first
 402 three objectives are also conflicting with *irrigation*. Some interrelations exist between *irri-*
 403 *gation* and *level* in the top-right part of the figure corresponding to negative performances
 404 on *algae (temperature)* and *sedimentation*. So, two groups of conflicting alternatives can
 405 be identified: on one side, possible compromises between *level* and *irrigation*, on the other
 406 side, compromises between *algae, temperature,* and *sedimentation*. A slight improvement in
 407 the objectives of the first group produces a significant worsening in the performances of the
 408 second group and vice-versa.

409 **Pareto front extremes**

410 In Table 1 we compare the optimal single-objective policies mapping into the extreme
 411 points of the Pareto front, obtained by setting to zero all the components of the weights
 412 vector ($\boldsymbol{\lambda} = |\lambda^{lev} \lambda^{algae} \lambda^{sed} \lambda^{irr} \lambda^{temp}|$), but the one corresponding to the objective considered.
 413 The associated policies are named accordingly (e.g., the extreme policy for level p^{lev} is
 414 obtained by setting $\boldsymbol{\lambda} = |1 0 0 0 0|$). The performance is evaluated as the improvement
 415 with respect to the benchmark, which is the current best available solution. Results show
 416 that the SWS operation has a considerable impact on all the water-related issues considered
 417 and all the optimal policies significantly outperform the benchmark. Not surprisingly, the
 418 room for improvement on the irrigation sector is quite limited since this objective was the
 419 primary target considered in designing the benchmark policy. The analysis of the release
 420 strategy adopted by the individual policies is useful to validate the behaviours we prefigured
 421 in the description of the main issues involved in the problem (see above). In what follows,
 422 we therefore evaluate the different policies by analyzing the temporal pattern of the main
 423 variables (Figures 5-9). We do not analyze p^{irr} in detail as it is almost equivalent to the
 424 benchmark.

425 *Level*

426 The improvement obtainable with policy p^{lev} with respect to the benchmark is significant:
427 a high and constant lake level is highly conflicting with the release of water for other uses,
428 particularly for irrigation in dry years, which is the target of benchmark operation. The
429 system controlled with this policy behaves in a quite easily interpretable way (Figure 5b):
430 the optimal policy p^{lev} tries to keep the lake at the constant level of 182.8 m a.s.l. (37.8 m
431 from the bottom) by keeping the release at a minimum.

432 *Algal bloom*

433 The improvement obtained by policy p^{algae} over the benchmark is remarkable (a daily
434 average of nearly 4.7 $\mu\text{g/L}$ of Chl-a), meaning that the SWS operation might positively
435 impact algal blooms. Without constraints on the reservoir level, blooms are controlled by
436 increasing the release in spring/summer (Figure 6b), when algal blooms are more likely
437 to occur, thus flushing away water volumes with high Chl-a concentration. The resulting
438 reservoir levels are generally lower than those produced by the benchmark (Figure 6a) and
439 follow the natural inflow pattern. In other words, the reservoir capacity is not exploited and
440 the reservoir follows a river-like behaviour.

441 *Sedimentation*

442 The improvement of policy p^{sed} on the benchmark is less significant than with policy p^{algae}
443 and p^{lev} . Yet, the SWS operation seems to affect also the silting of the reservoir. Again, the
444 behaviour of the system controlled by policy p^{sed} follows the inflow dynamics: without any
445 constraint on lake level or penalty on wasting water, the optimal policy suggests to release
446 the inflow in order to flush out the maximum amount of sediments. This is evident when
447 comparing the inflow and total release patterns (Figure 7b). Furthermore, releasing the
448 inflow produces two favourable conditions: first, low lake levels reduce the retention time of
449 the reservoir and, therefore, prevent in-reservoir sedimentation, as also observed by Lee and
450 Guy (2012). Second, since the sediments tend to intrude along the thermocline (Yajima et al.
451 2006), which means that the highest TSS concentration is found around the thermocline, the

452 optimal policy moves the temperature profile so as to have the thermocline in correspondence
453 of one of the two depths where the release can be performed. As shown in Figure 7c, the
454 optimal policy maintains, on average, the thermocline around 13 m depth, while with the
455 benchmark policy the thermocline is constantly deeper, generally deeper than 13 m from the
456 surface, and so is not able to release the same amount of sediments as p^{sed} because sediments
457 are trapped in the uncontrollable region of the reservoir. The TSS concentration profiles for
458 the benchmark and policy p^{sed} for a small flood event are shown in Figure 8a-b: after the
459 flood event of January 9-10, policy p^{sed} is able to release more sediment and to reduce the TSS
460 concentration in the lake by adopting a more effective release strategy than the benchmark.
461 The benchmark in January releases only from the outlet at 28 m depth according to the rule
462 defined by the dam construction authority. On the contrary, the strategy adopted by policy
463 p^{sed} (Figure 8c) first releases at -3 m in order to keep the thermocline shallower (Gelda and
464 Effler 2007) and, then, opens the -13 m siphon in order to actually release the sediments.

465 *Temperature*

466 Lake level dynamics under policy p^{temp} follow a nearly periodic pattern (Figure 9b), with
467 values constantly lower than the benchmark (Figure 9a) generated by higher releases. With
468 this strategy, the optimal policy is able to stabilize the thermocline between 5 and 10 m depth
469 (Figure 9c) and to exploit blending between the two controlled siphons to generate the same
470 temperature as the inflow. The benchmark follows a different and less effective strategy
471 by maintaining the lake at higher levels. As a consequence, in summer the thermocline
472 decreases at 13 m depth or deeper and blending can not be exploited to meet the target
473 temperature of the outflows since the water has the same temperature at the two controlled
474 depths. Probably, a further release decision variable at 7 m depth (which is the average
475 depth at which the thermocline is located during the stratification) could make it easier to
476 intercept the intruding inflow and, therefore, to further reduce the difference between the
477 inflow-outflow temperature. This will be the subject of subsequent research.

478 **Pareto front compromises**

479 In this section, we analyze one policy (policy p^5 in Table 2 and Figures 3-4) particularly
480 interesting in terms of balance of the different objectives, which thus constitutes a strong
481 candidate to be the final compromise solution in a real policy making context. Obviously,
482 this is a subjective evaluation by the authors and the real Decision Maker (DM) might
483 prefer different alternatives. However, the aim of this analysis is to show that there is room
484 to design compromise policies outperforming the benchmark and of practical interest. As an
485 example, in Table 2 we also report the performance of other two alternatives showing slightly
486 different trade-offs among the objectives but still good candidates as best compromise.

487 As shown in Figure 10, policy p^5 really represents a possible interesting compromise
488 solution among the five conflicting objectives: although there exist some better alternatives
489 for each objective, which however are dominated on the other objectives, policy p^5 balances
490 all the considered objectives. Indeed, the presence of multiple objectives allow to obtain
491 very good performances (bottom part of the figure) only on few, non-conflicting objectives
492 which are identified by horizontal lines. However, these solutions negatively impact on the
493 remaining objectives, graphically represented by oblique lines. It is worth noting that the
494 conflicts previously identified between the two groups of objectives (*algae*, *sediment* and
495 *temperature* on one side and *level* and *irrigation* on the other one) are particularly evident
496 in Figure 10 looking at the high number of crossing lines between *temperature* and *level*.
497 Among the set of alternatives, policy p^5 represents instead a possible compromise as it
498 is almost an horizontal line, meaning that the satisfaction of the five objectives is almost
499 equivalent.

500 The SWS operating strategy of policy p^5 (Figure 11) is a mix of the different Pareto
501 extreme policies which tries to consider all the objectives and, depending on the period of
502 the year, it focuses on different objectives as shown in Figure 12. This compromise strategy
503 is also evident looking at Figure 13: following p^{algae} , p^{sed} and p^{temp} , policy p^5 activates a
504 drawdown cycle in summer by increasing the release, while it stores water in the winter and

505 early spring period to satisfy the *level* and *irrigation* objectives.

506 **CONCLUSIONS**

507 Despite the recent progress in the design of optimal planning and management strategies
508 for water resources systems, most of the studies reported in the domain literature deals with
509 quantitative objectives only, e.g., agricultural supply, hydropower energy production, and
510 flood control. However, a really sustainable operation should also consider water quality
511 targets. This paper illustrates a novel approach to design optimal operating policies for
512 water reservoirs equipped with multiple outlet release schemes which optimize quantity and
513 quality objectives both in-reservoir and downstream.

514 We combined a 1D physically-based description of the hydrodynamic and ecological pro-
515 cesses taking place in the lake with a batch-mode Reinforcement Learning algorithm to
516 design quasi optimal release strategies conditioned upon an augmented state including not
517 only the current storage but also water characteristics, such as temperature and total sus-
518 pended solid at different depths. The use of a batch-mode approach makes it possible to
519 combine simulation experiments conducted with a high fidelity physically-based model and
520 the guarantees on the policy optimality property offered by Dynamic Programming family
521 methods, which is particularly useful in a complex, many-objective context with no historical
522 reference for the operation.

523 The application to Tono Dam case study shows that the operation designed with the
524 proposed approach outperforms the current best available solution on all the objectives in-
525 dependently considered but also produces compromise policies that considerably improve the
526 water quality objectives in-reservoir and downstream at the cost of a very small, practically
527 negligible, reduction of the irrigation supply. For example, with the examined compromise
528 policy the improvements in the algae, sedimentation and temperature objectives with respect
529 to the benchmark are equal to 42%, 14% and 5% respectively, while the worsening of the
530 irrigation objective is equal to 15%. However, given the strong conflict between the level
531 objective and the water quality interests, favoring these latter produces a significant decrease

532 in the level objective performance, i.e. from 35.9 m² to 120.9 m².

533 Future research will concentrate on increasing the number of controlled siphons to im-
534 prove operation flexibility: allowing the release at more different depths should make it
535 possible to open the siphons at the point with the maximum sediment concentration, as
536 well as to do blending depending to the thermocline position. Further improvements on the
537 methodological ground can be achieved by: *i*) substituting the original high fidelity model
538 for a lower order dynamic emulator (Castelletti et al. 2012); *ii*) adopting Active Learning ap-
539 proaches (Rachelson et al. 2011) to improve the simulation-based exploration of the system
540 behavior and generate an equally informative data-sample with lower dimensionality and so
541 lower associated computational cost, but also potentially improved performance; *iii*) using
542 projections method (e.g., principal component analysis (Galelli et al. 2011)) to aggregate
543 interrelated objectives thus allowing a more dense Pareto front approximation and the use
544 of visualization technique to jointly explore the decision and the objective space (Kollat and
545 Reed 2007).

546 **ACKNOWLEDGEMENT**

547 The work was funded by the Tono Dam construction office, Japan Ministry of Land,
548 Infrastructure, Transport and Tourism. Matteo Giuliani was partially supported by *Fon-*
549 *dazione Fratelli Confalonieri*. The authors would like to thank Giovanni Garbarini and
550 Alessandra Galli for their contribution in developing the numerical analysis.

551 **REFERENCES**

- 552 Baltar, A. and Fontane, D. (2008). “Use of Multiobjective Particle Swarm Optimization in
553 Water Resources Management.” *Journal of Water Resources Planning and Management*,
554 134(3), 257–265.
- 555 Bohan, J. and Grace, J. (1973). “Selective withdrawal from man-made lakes.” *Report No.*
556 *H-73-4*, U.S. Army Engineer Waterways Experiment Station, Vicksburg, Mississippi.
- 557 Castelletti, A., Antenucci, J. P., Limosani, D., Quach Thi, X., and Soncini-Sessa, R. (2011).
558 “Interactive response surface approaches using computationally intensive models for mul-

559 tiobjective planning of lake water quality remediation.” *Water Resources Research*, 47,
560 W09534.

561 Castelletti, A., Galelli, S., Ratto, M., Soncini-Sessa, R., and Young, P. C. (2012). “A general
562 framework for Dynamic Emulation Modelling in environmental problems.” *Environmental*
563 *Modelling & Software*, 34, 5–18.

564 Castelletti, A., Galelli, S., Restelli, M., and Soncini-Sessa, R. (2010). “Tree-based rein-
565 forcement learning for optimal water reservoir operation.” *Water Resources Research*,
566 46(W09507).

567 Castelletti, A., Pianosi, F., and Soncini-Sessa, R. (2008). “Water reservoir control under
568 economic, social and environmental constraints.” *Automatica*, 44(6), 1595–1607.

569 Chaves, P. and Kojiri, T. (2007). “Deriving reservoir operational strategies considering water
570 quantity and quality objectives by stochastic fuzzy neural networks.” *Advances in Water*
571 *Resources*, 30(5), 1329–1341.

572 Dandy, G. and Crawley, P. (1992). “Optimum operation of a multiple reservoir system in-
573 cluding salinity effects.” *Water Resources Research*, 28(4), 979–990.

574 Davis, J., Holland, J., Schneider, M., and Wilhelms, S. (1987). “SELECT: a numerical,
575 one-dimensional model for selective withdrawal.” *Report No. E-87-2*, U.S. Army Engineer
576 Waterways Experiment Station, Vicksburg, Mississippi.

577 Dhar, A. and Datta, B. (2008). “Optimal operation of reservoirs for downstream water
578 quality control using linked simulation optimization.” *Hydrological Processes*, 22, 842–853.

579 Dortch, M. (1997). “Water quality considerations in reservoir management.” *Water resources*
580 *update*, 108, 32–42.

581 Ernst, D., Geurts, P., and Wehenkel, L. (2005). “Tree-Based Batch Mode Reinforcement
582 Learning.” *Journal of Machine Learning Research*, 6, 503–556.

583 Evans, R. (1994). “Empirical evidence of the importance of sediment resuspension.” *Journal*
584 *of Hydrobiologia*, 284(1), 5–12.

585 Ferris, J. and Lehman, J. (2007). “Interannual variation in diatom bloom dynamics: Roles

586 of hydrology, nutrient limitation, sinking, and whole lake manipulation.” *Water Research*,
587 41(12), 2551–2562.

588 Fontane, D., Labadie, J., and Loftis, B. (1981). “Optimal control of reservoir discharge
589 quality through selective withdrawal.” *Water Resources Research*, 12(6), 1594–1604.

590 Galelli, S., Giuliani, M., and Soncini-Sessa, R. (2011). “Dealing with many-criteria problems
591 in water resources planning and management.” *18th IFAC World Congress*, Milan, Italy.

592 Gelda, R. and Effler, S. (2007). “Simulation of Operations and Water Quality Performance
593 of Reservoir Multilevel Intake Configurations.” *Journal of Water Resources Planning and
594 Management*, 133(1), 78–86.

595 Hanna, R. B. and Saito, L. (2001). “Simulated Limnological Effects of the Shasta Lake
596 Temperature Control Device.” *Environmental Management*, 27(4), 609–626.

597 Hanna, R. B., Saito, L., Bartholow, J. M., and Sandelin, J. (1999). “Results of Simulated
598 Temperature Control Device Operations on In-Reservoir and Discharge Water Tempera-
599 tures Using CE-QUAL-W2.” *Lake and Reservoir Management*, 15(2), 87–102.

600 Hashimoto, T., Stedinger, J., and Loucks, D. (1982). “Reliability, resilience, and vulnerability
601 criteria for water resource system performance evaluation.” *Water Resources Research*,
602 18(1), 14–20.

603 Hayes, D., Labadie, J., Sanders, T., and Brown, J. (1998). “Enhancing water quality in
604 hydropower system operations.” *Water Resources Research*, 34(3), 471–483.

605 Hipsey, M., Romero, J., Antenucci, J., and Hamilton, D. (2006). *Computational Aquatic
606 Ecosystem Dynamics Model: CAEDYM v2.3 Science Manual*. Centre for Water Research,
607 University of Western Australia.

608 Huang, Y., Huang, G., Liu, D., Zhu, H., and Sun, W. (2012). “Simulation-based inexact
609 chance-constrained nonlinear programming for eutrophication management in the Xiangxi
610 Bay of Three Gorges Reservoir.” *Journal of Environmental Management*, 108, 54–65.

611 Ikebuchi, S. and Kojiri, T. (1992). “Multiobjective reservoir operation including turbidity
612 control.” *Water Resources Bulletin*, 28(1), 223–231.

- 613 ILEC (2005). *Managing Lakes and their Basins for Sustainable Use: A Report for Lake*
614 *Basin Managers and Stakeholders*. International Lake Environment Committee Founda-
615 tion, Kusatsu, Japan.
- 616 Imerito, A. (2007). *Dynamic Reservoir Simulation Model: DYRESM Science Manual*. Centre
617 for Water Research, University of Western Australia.
- 618 Kerachian, R. and Karamouz, M. (2006). “Optimal reservoir operation considering the wa-
619 ter quality issues: A stochastic conflict resolution approach.” *Water Resources Research*,
620 42(W12401).
- 621 Khan, N., Babel, M., Tingsanchali, T., Clemente, R., and Luong, H. (2012). “Reser-
622 voir Optimization-Simulation with a Sediment Evacuation Model to Minimize Irrigation
623 Deficits.” *Water Resources Management*, 1–21.
- 624 Kollat, J. and Reed, P. (2007). “A framework for Visually Interactive Decision-making and
625 Design using Evolutionary Multi-objective Optimization (VIDEO).” *Environmental Mod-*
626 *elling & Software*, 22(12), 1691–1704.
- 627 Koutsoyiannis, D. and Economou, A. (2003). “Evaluation of the parameterization-
628 simulation-optimization approach for the control of reservoir systems.” *Water Resources*
629 *Research*, 39(6), 1170–1187.
- 630 Kuo, J., Hsieh, P., and Jou, W. (2008). “Lake eutrophication management modeling using
631 dynamic programming.” *Journal of Environmental Management*, 88(4), 677–687.
- 632 Labadie, J. (2004). “Optimal operation of multireservoir systems: State-of-the-art review.”
633 *Journal of Water Resources Planning and Management*, 130(2), 93–111.
- 634 Lee, C. and Guy, F. (2012). “Assessing the potential of reservoir outflow management to
635 reduce sedimentation using continuous turbidity monitoring and reservoir modelling.” *Hy-*
636 *drological Processes*, n/a–n/a.
- 637 Nandalal, K. and Bogardi, J. (1995). “Optimal operation of a reservoir for quality-control
638 using inflows and outflows.” *Water Science and Technology*, 31(8), 273–280.
- 639 Nece, R. (1970). “Register of Selective Withdrawal Works in United States.” *Journal of the*

640 *Hydraulics Division*, 96(9), 1841–1872.

641 Orlob, G. and Simonovic, S. (1982). “Reservoir operation for water quality control.” *Ex-*
642 *perience in Operation of Hydrosystems*, T. Unny and E. McBean, eds., Water Resources
643 Publications, Littleton, Colorado, 263–285.

644 Ostfeld, A. and Salomons, S. (2005). “A hybrid genetic-instance based learning algorithm
645 for CE-QUAL-W2 calibration.” *Journal of Hydrology*, 310, 122–142.

646 Powell, W. (2007). *Approximate Dynamic Programming: Solving the curses of dimensional-*
647 *ity*. Wiley, NJ.

648 Rachelson, E., Schnitzler, F., Wehenkel, L., and Ernst, D. (2011). “Optimal sample selection
649 for batch-mode reinforcement learning.” *Proceedings of the 3rd International Conference*
650 *on Agents and Artificial Intelligence (ICAART 2011)*, Rome, Italy.

651 Sherman, B. (2000). “Scoping Options for Mitigating Cold Water Discharges from Dams.”
652 *Report no.*, CSIRO Land and Water, Canberra, AUS.

653 Smith, D., Wilhelms, S., Holland, J., Dortch, M., and Davis, J. (1987). “Improved description
654 of selective withdrawal through point sinks.” *Report No. E-87-2*, U.S. Army Engineer
655 Waterways Experiment Station, Vicksburg, Mississippi.

656 Tsitsiklis, J. and Van Roy, B. (1996). “Feature-Based Methods for Large Scale Dynamic
657 Programming.” *Machine Learning*, 22, 59–94.

658 Vermeyen, T. (1999). “Summary of the Shasta dam temperature control device and how it
659 is working.” *Water Operation and Maintenance Bulletin*, 187, 9–17.

660 Vermeyen, T., DeMoyer, C., Delzer, W., and Kubly, D. (2003). “A Survey of Selective With-
661 drawal Systems.” *Report No. R-03-03*, Denver Technical Service Center Water Resources
662 Services Water Resources Research Laboratory, Denver, Colorado.

663 Westphal, K., Vogel, R., Kirshen, P., and Chapra, S. (2003). “Decision support system for
664 adaptive water supply management.” *Journal of Water Resources Planning and Manage-*
665 *ment*, 129(3), 165–177.

666 Yajima, H., Kikkawa, S., and Ishiguro, J. (2006). “Effect of selective withdrawal system

667 operation on the long-and short-term water conservation in a reservoir.” *Annual Journal*
668 *of Hydraulic Engineering, JSCE*, 50, 1375–1380. (in Japanese).

669 **List of Tables**

670	1	Performances of the Pareto front extreme policies. The gray shaded objective was not directly considered in the optimization.	29
671			
672	2	Performances of three interesting compromise alternatives. The gray shaded objective was not directly considered in the optimization.	30
673			

TABLE 1. Performances of the Pareto front extreme policies. The gray shaded objective was not directly considered in the optimization.

Sector	Description	Unit of Measure	Benchmark	Policy	Improvement
Level	daily quadratic positive difference w.r.t. the reference level of 182.8 m a.s.l.	$(\text{m})^2$	35.9338	0.0195	35.9143
Algae	daily average maximum concentration of Chl-a in the see-through layer	$\mu\text{g/L}$	6.8133	2.1293	4.684
Sedimentation	daily volume of sediment expelled with the release	g/day	$2.8102 \cdot 10^6$	$2.8699 \cdot 10^6$	$5.97 \cdot 10^4$
Irrigation	daily quadratic water deficit modulated by β	$(\text{m}^3/\text{day})^2$	$6.9408 \cdot 10^7$	$6.5449 \cdot 10^7$	$3.959 \cdot 10^6$
Irrigation (2)	daily water deficit along the year	m^3/s	0.0280	0.0270	0.001
Temperature	daily average quadratic difference of temperature between inflow and outflow	$(^\circ\text{C})^2$	4.3818	2.5143	1.8675

TABLE 2. Performances of three interesting compromise alternatives. The gray shaded objective was not directly considered in the optimization.

Policy [weights]	Level (m) ²	Algae μg/L	Sedimentation g/day	Irrigation (m ³ /day) ²	Irrigation (2) m ³ /s	Temperature (°C) ²
p^5 [.33 .13 .0001 .2099 .33]	120.9102	3.9351	$3.2093 \cdot 10^6$	$7.9648 \cdot 10^7$	0.0374	4.1692
p^{5a} [.35 .10 .0001 .1999 .35]	151.2584	3.3330	$2.7721 \cdot 10^6$	$7.7229 \cdot 10^7$	0.0348	3.6647
p^{5b} [.31 .15 .0001 .2299 .31]	111.6356	4.1838	$3.0261 \cdot 10^6$	$7.5715 \cdot 10^7$	0.0346	4.2691
Best case	0.0195	2.1293	$3.8376 \cdot 10^6$	$4.0183 \cdot 10^7$	0.0177	2.2275
Worst case	322.6452	6.8133	$2.4509 \cdot 10^6$	$14.2593 \cdot 10^7$	0.0504	12.4755

674 **List of Figures**

675 1 Tono Dam location in Western Japan (panel a), the main characteristics of
676 the reservoir with the decision variables adopted in this study (panel b), and
677 the schematization of the SWS structure (panel c). 33

678 2 Schematization of the procedure adopted. The black line is the optimization
679 workflow, the dashed line is evaluation via simulation of the optimal operating
680 policy. 34

681 3 Projection of the 5D Pareto front in the objectives space *algae-sedimentation*;
682 circles size is proportional to the logarithm of the *level* objective (the bigger
683 the circle, the better is the alternative); colors provide the *temperature* ob-
684 jective. The *irrigation* objective is not represented. Policy p^5 represents a
685 possible compromise alternative. 35

686 4 Projection of the 5D Pareto front in the objectives space *algae-sedimentation*;
687 circles size is proportional to the logarithm of the *level* objective (the bigger
688 the circle, the better is the alternative); colors provide the *temperature* ob-
689 jective. The *temperature* objective is not represented. Policy p^5 represents a
690 possible compromise alternative. 36

691 5 Water surface level (black line), inflow (blue line), and the total actual release
692 (red line) produced by the benchmark in panel (a) and policy p^{lev} in panel (b)
693 over the validation period. 37

694 6 Water surface level (black line), inflow (blue line), and the total actual release
695 (red line) produced by the benchmark in panel (a) and policy p^{algae} in panel
696 (b) over the validation period. 38

697 7 Water surface level (black line), inflow (blue line), and the total actual release
698 (red line) produced by the benchmark in panel (a) and policy p^{sed} in panel
699 (b) over the validation period. Panel (c) reports the thermocline depth for
700 the same alternatives over the validation period (the dotted lines represent
701 the average depth). 39

702 8 TSS concentration profiles for a small flood event in January 1990 for the
703 benchmark in panel (a) and policy p^{sed} in panel (b). The black line represents
704 the thermocline and the green lines represent the positions of the controlled
705 outlets (3 and 13 m depth). Panel (c) shows the outlets operating strategy
706 for policy p^{sed} in the same period. 40

707 9 Water surface level (black line), inflow (blue line), and the total actual release
708 (red line) produced by the benchmark in panel (a) and policy p^{temp} in panel
709 (b) over the validation period. Panel (c) reports the thermocline depth for
710 the same alternatives over the validation period (the dotted lines represent
711 the average depth). 41

712 10 Representation of the performances of policy p^5 (red line) with respect to all
713 the other alternatives. For illustration purposes the objectives are standard-
714 ized (zero mean and unit standard deviation). 42

715 11 Water surface level (black line), inflow (blue line), and the total actual release
716 (red line) produced by the benchmark in panel (a) and policy p^5 in panel (b)
717 over the validation period. 43

718	12	Comparison of the cyclostationary means (over the validation period) of the	
719		immediate costs $g_t^k(\cdot)$ between policy p^5 and the single-objective policies. . .	44
720	13	Cyclostationary mean (over the validation period) of the lake levels for differ-	
721		ent policies.	45

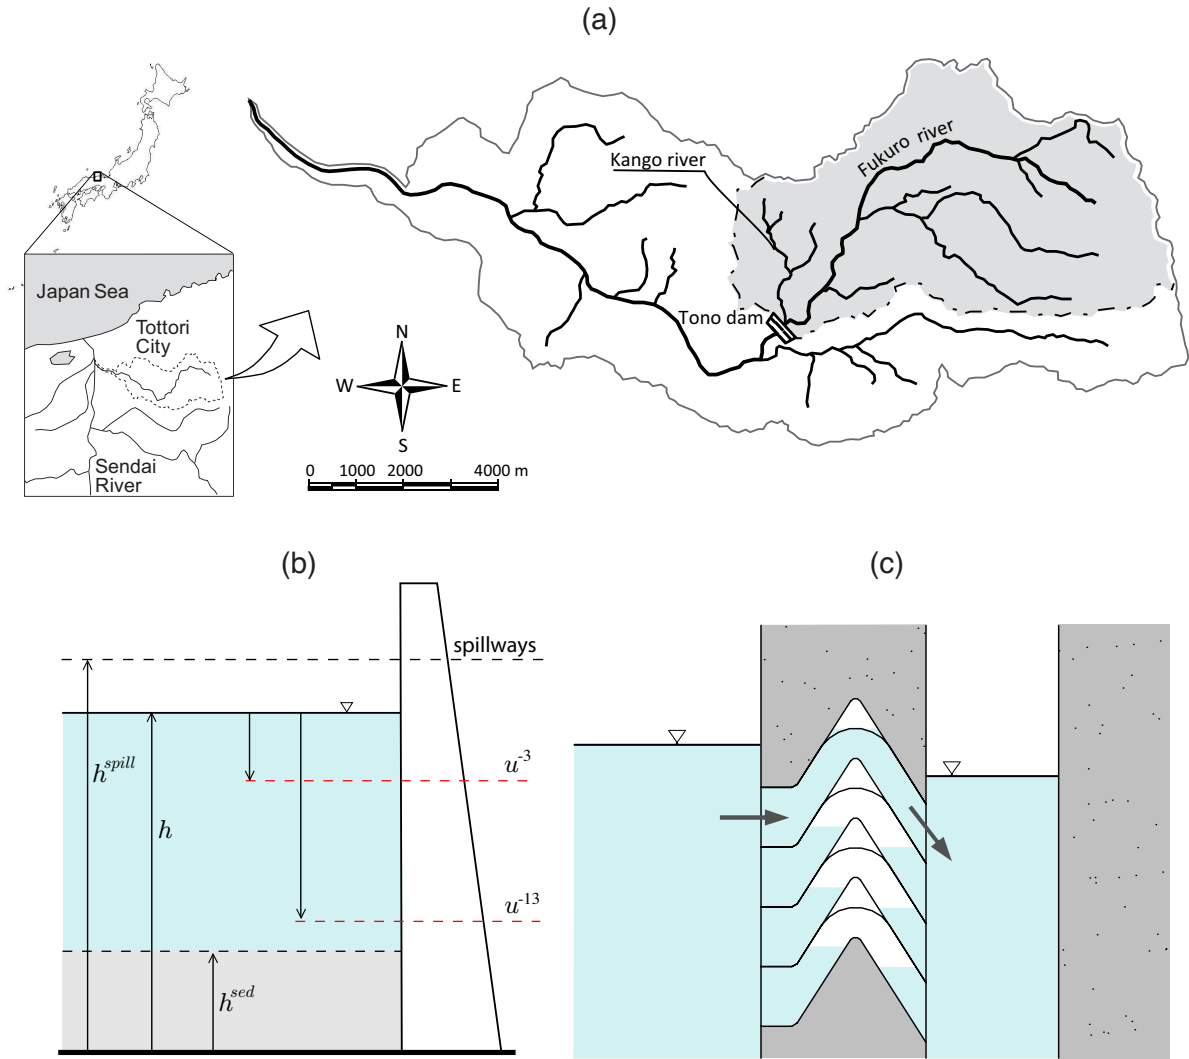


FIG. 1. Tono Dam location in Western Japan (panel a), the main characteristics of the reservoir with the decision variables adopted in this study (panel b), and the schematization of the SWS structure (panel c).

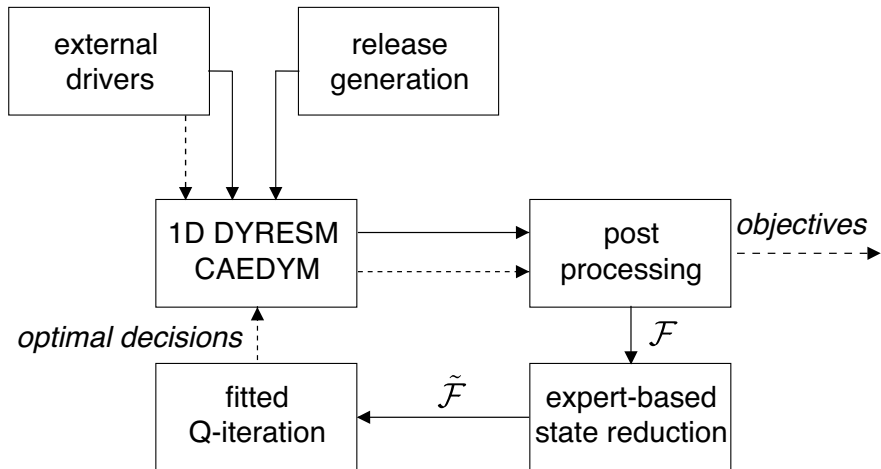


FIG. 2. Schematization of the procedure adopted. The black line is the optimization workflow, the dashed line is evaluation via simulation of the optimal operating policy.




FIG. 3. Projection of the 5D Pareto front in the objectives space algae-sedimentation; circles size is proportional to the logarithm of the level objective (the bigger the circle, the better is the alternative); colors provide the temperature objective. The irrigation objective is not represented. Policy p^5 represents a possible compromise alternative.

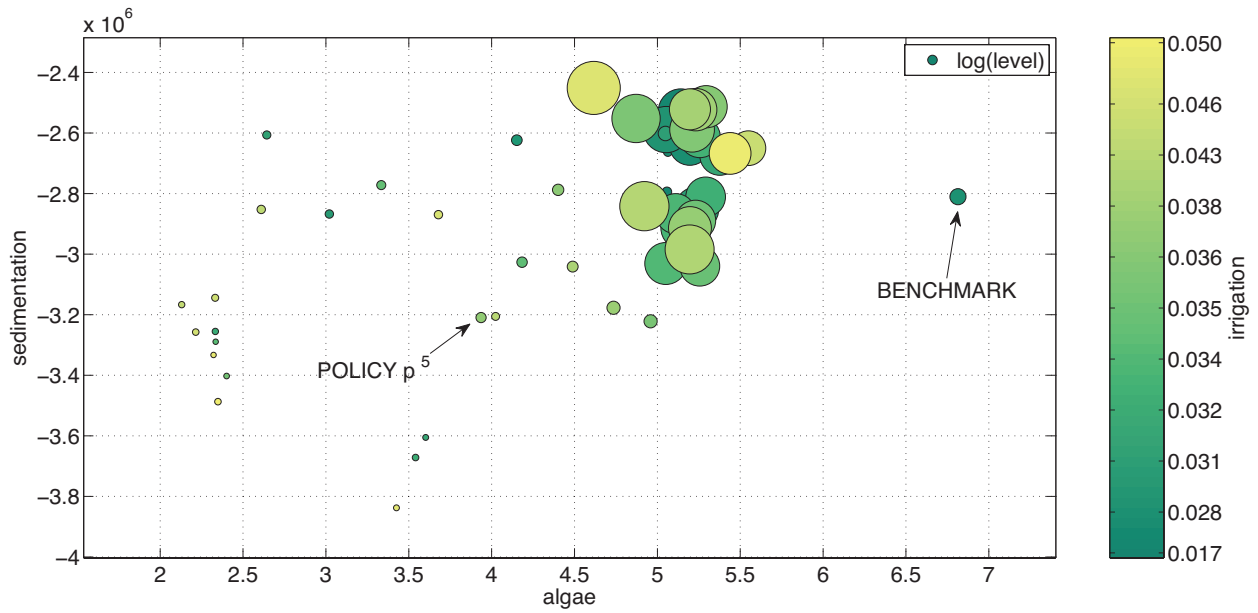


FIG. 4. Projection of the 5D Pareto front in the objectives space algae-sedimentation; circles size is proportional to the logarithm of the level objective (the bigger the circle, the better is the alternative); colors provide the temperature objective. The temperature objective is not represented. Policy p^5 represents a possible compromise alternative.




FIG. 5. Water surface level (black line), inflow (blue line), and the total actual release (red line) produced by the benchmark in panel (a) and policy p^{lev} in panel (b) over the validation period.

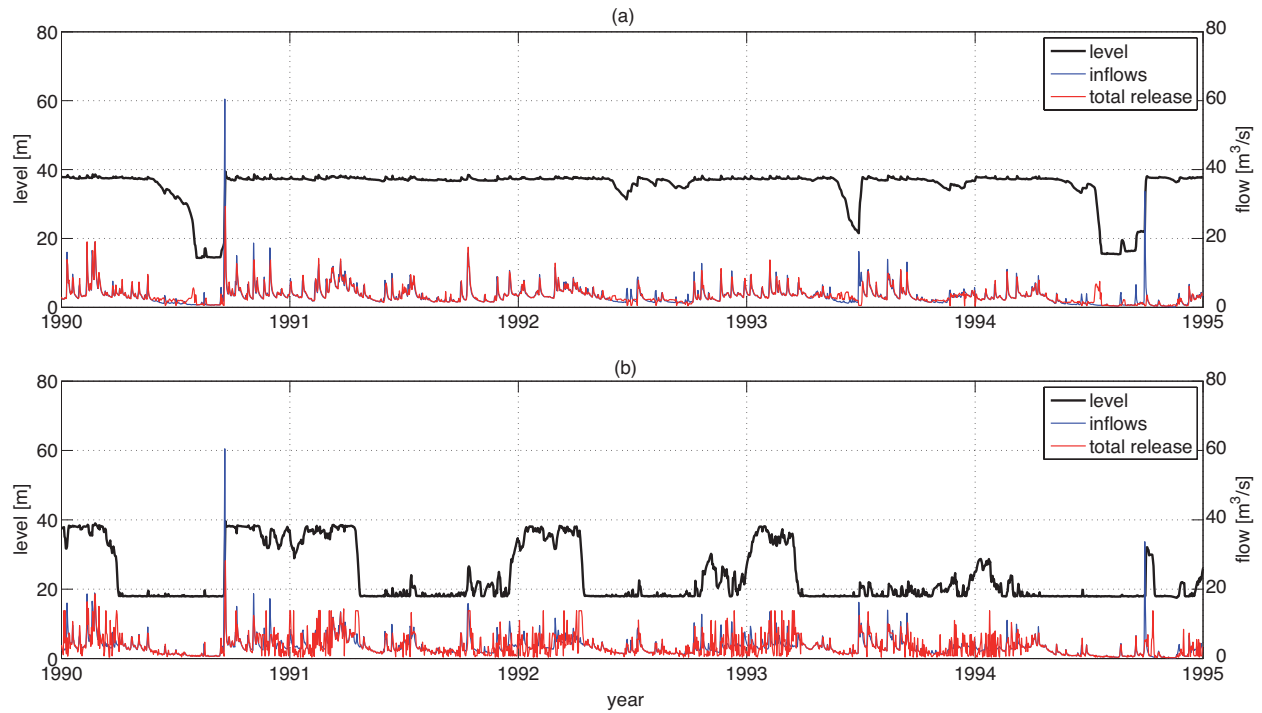


FIG. 6. Water surface level (black line), inflow (blue line), and the total actual release (red line) produced by the benchmark in panel (a) and policy p^{algae} in panel (b) over the validation period.




FIG. 7. Water surface level (black line), inflow (blue line), and the total actual release (red line) produced by the benchmark in panel (a) and policy p^{sed} in panel (b) over the validation period. Panel (c) reports the thermocline depth for the same alternatives over the validation period (the dotted lines represent the average depth).

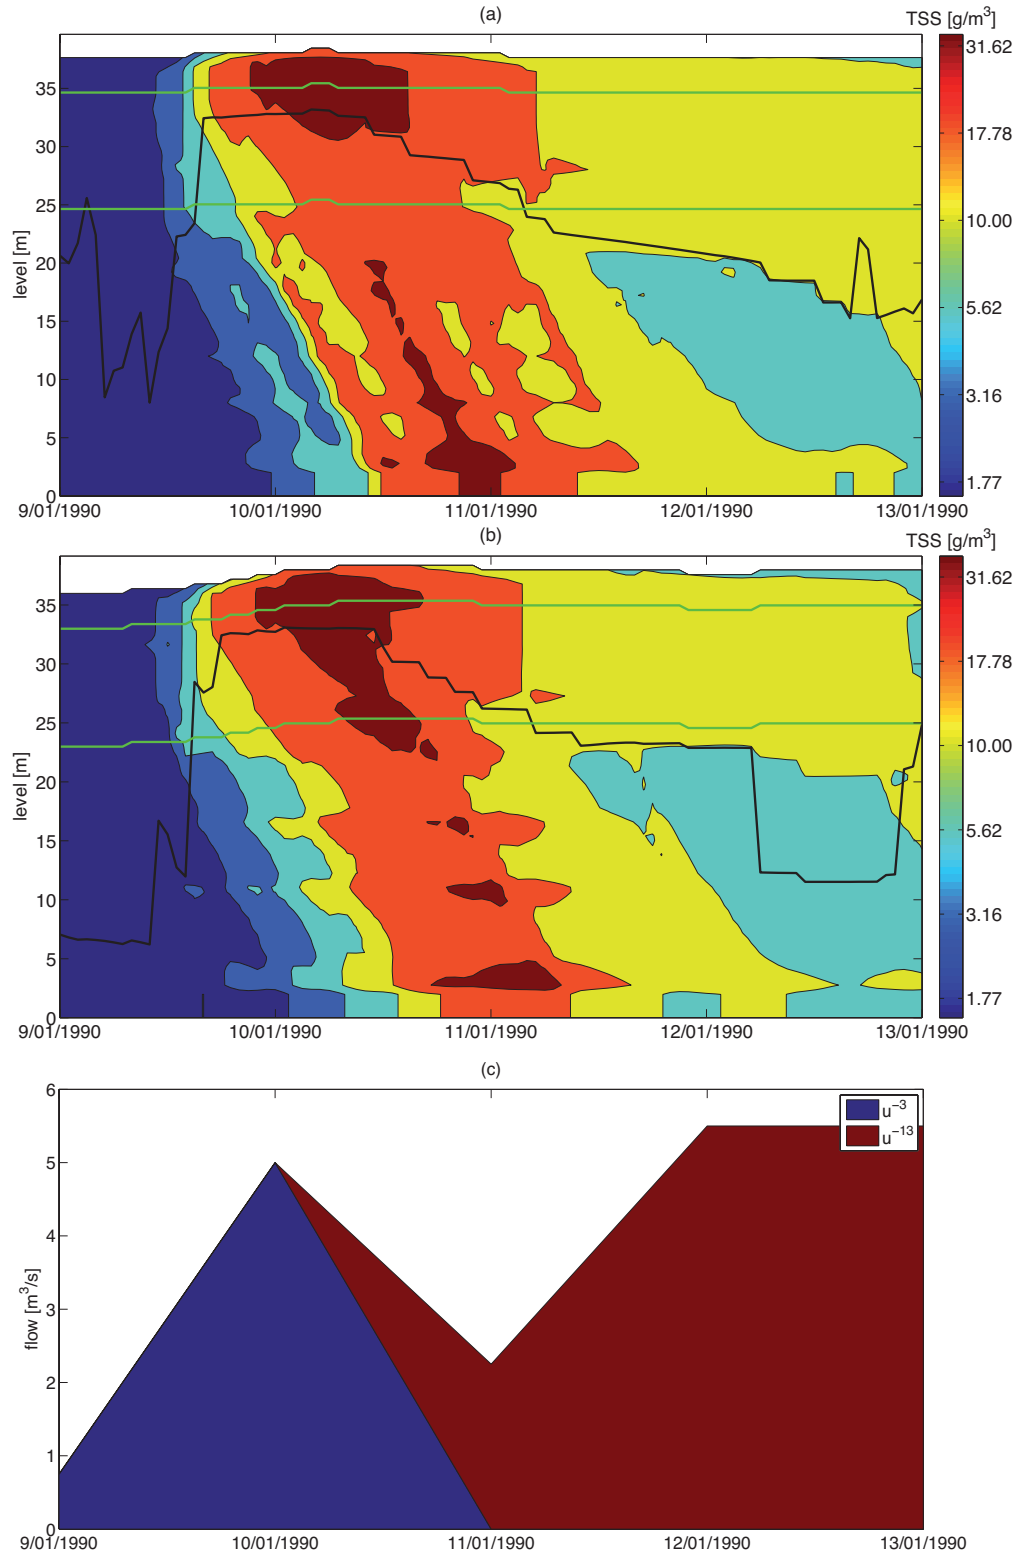


FIG. 8. TSS concentration profiles for a small flood event in January 1990 for the benchmark in panel (a) and policy p^{sed} in panel (b). The black line represents the thermocline and the green lines represent the positions of the controlled outlets (3 and 13 m depth). Panel (c) shows the outlets operating strategy for policy p^{sed} in the same period.




FIG. 9. Water surface level (black line), inflow (blue line), and the total actual release (red line) produced by the benchmark in panel (a) and policy p^{temp} in panel (b) over the validation period. Panel (c) reports the thermocline depth for the same alternatives over the validation period (the dotted lines represent the average depth).

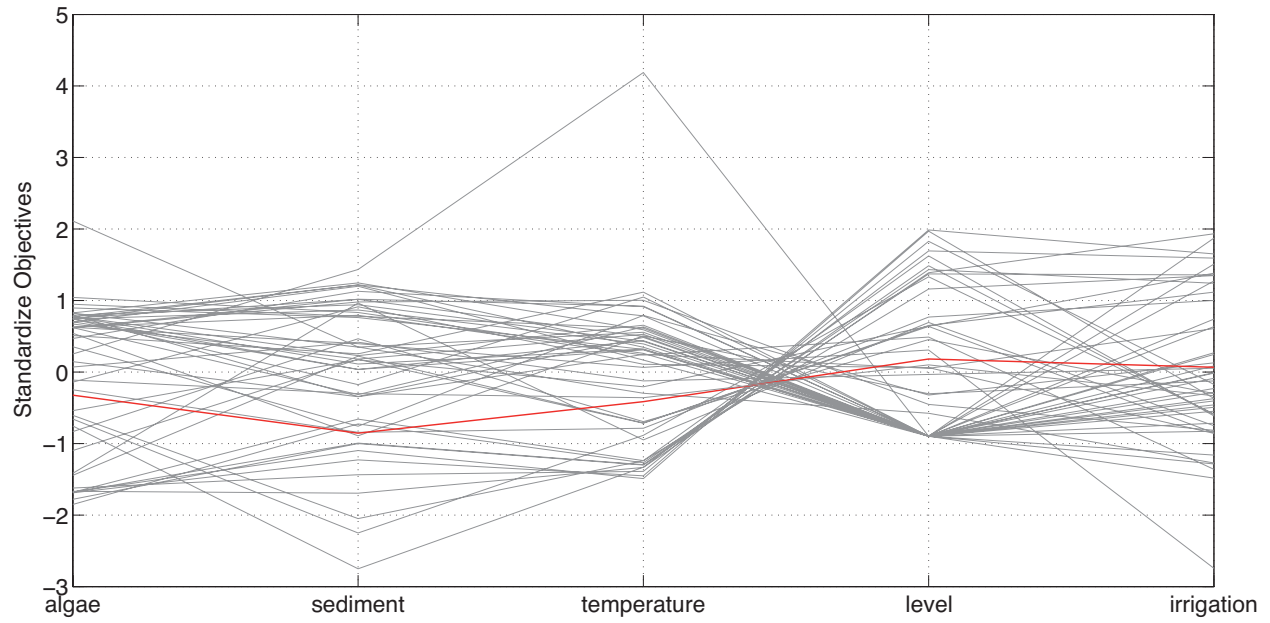


FIG. 10. Representation of the performances of policy p^5 (red line) with respect to all the other alternatives. For illustration purposes the objectives are standardized (zero mean and unit standard deviation).




FIG. 11. Water surface level (black line), inflow (blue line), and the total actual release (red line) produced by the benchmark in panel (a) and policy p^5 in panel (b) over the validation period.

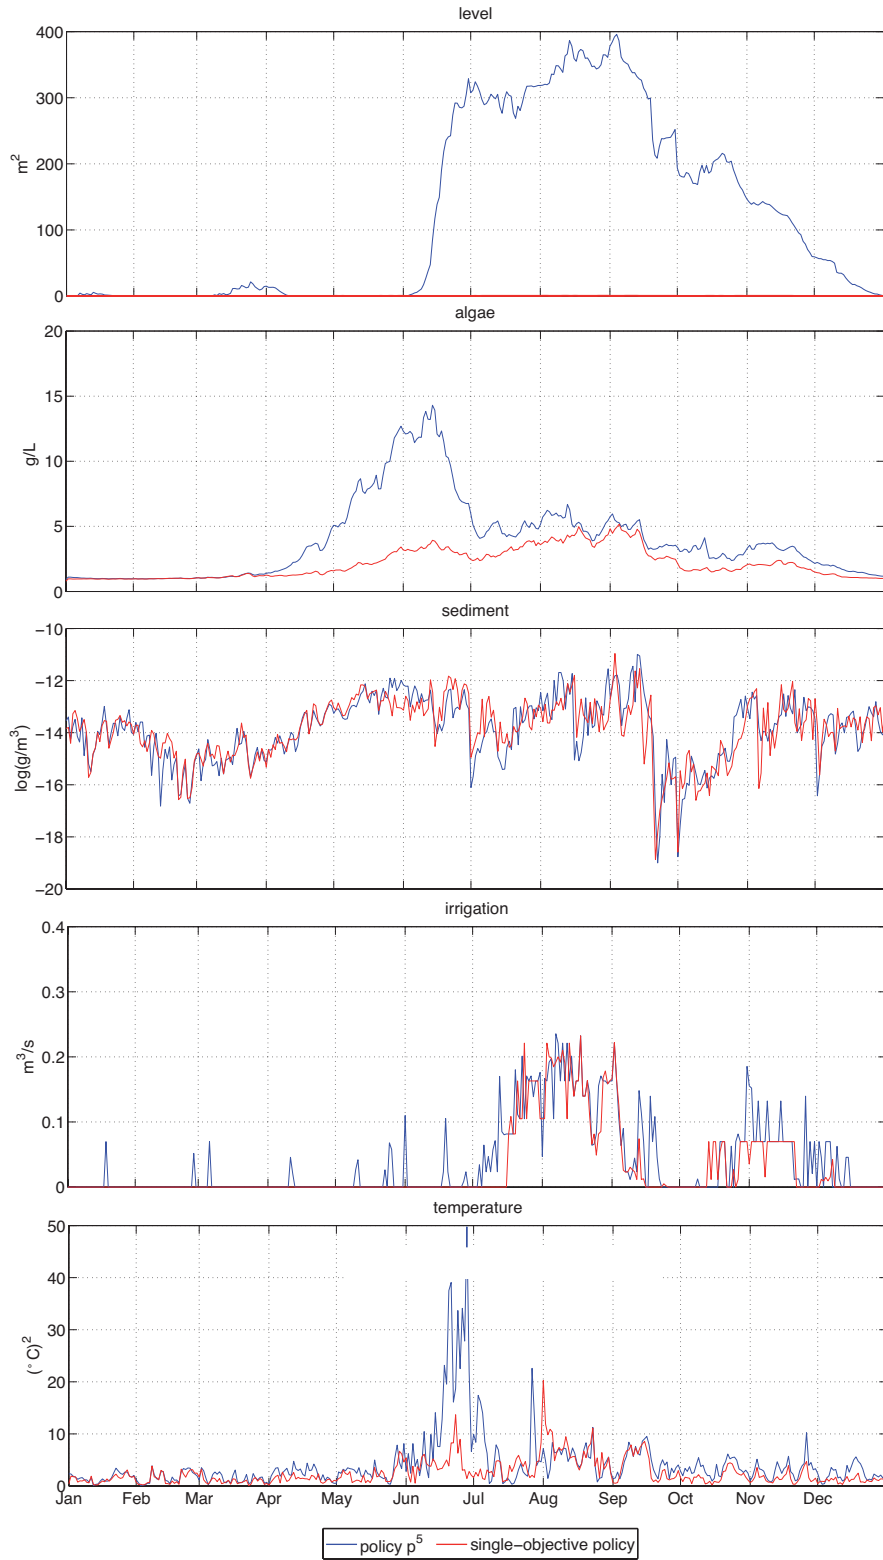


FIG. 12. Comparison of the cyclostationary means (over the validation period) of the immediate costs $g_t^k(\cdot)$ between policy p^5 and the single-objective policies.

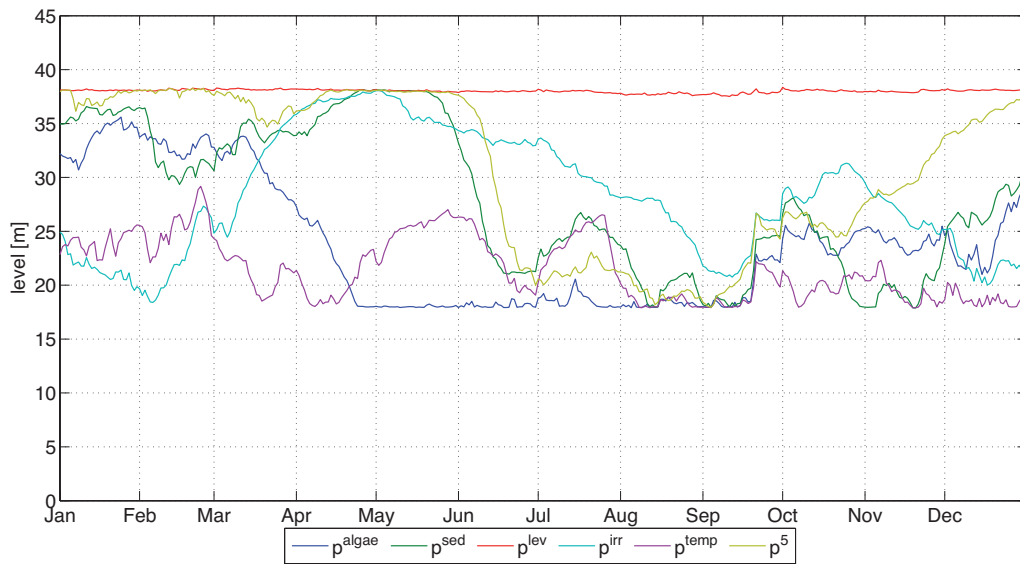


FIG. 13. Cyclostationary mean (over the validation period) of the lake levels for different policies.